PLATTFORM INDUSTRIE 4.0

WORKING PAPER

Artificial intelligence and law in the context of Industrie 4.0

Imprint

Published by
Federal Ministry for Economic Affairs and Energy (BMWi)
Public Relations
10119 Berlin
www.bmwi.de

Text and editing
Plattform Industrie 4.0
Bertolt-Brecht-Platz 3
10117 Berlin

Design
PRpetuum GmbH, Munich

Status
April 2019

Print
BMWi

Illustrations
sdecoret – Adobe Stock (Title),
Patra Kongsirimongkolchai / EyeEm – Getty Images (p. 5),
Photographer is my life. – Getty Images (p. 8, p. 20),
Andriy Onufriyenko – Getty Images (p. 11),
Yuichiro Chino – Getty Images (p. 14, p. 31),
Westend61 – Getty Images (p. 17),
BlackJack3D – Getty Images (p. 24),
Caiaimage/Robert Daly – Getty Images (p. 28)

This publication as well as further publications can be obtained from:
Federal Ministry for Economic Affairs and Energy (BMWi)
Public Relations
E-mail: publikationen@bundesregierung.de
www.bmwi.de

Central procurement service:
Tel.: +49 30 182722721
Fax: +49 30 18102722721

Content

Introduction: What is this analysis about?

"Artificial intelligence (AI) refers to systems that display intelligent behaviour by analysing their environment and taking actions – with some degree of autonomy – to achieve specific goals. AI-based systems can be purely software-based, acting in the virtual world (e.g. voice assistants, image analysis software, search engines, speech and face recognition systems) or AI can be embedded in hardware devices (e.g. advanced robots, autonomous cars, drones or Internet of Things applications)." [1]

AI systems are designed to replicate human cognitive abilities through the use of algorithms. In simple terms, algorithms are mathematical commands that transform input data into output data.

One possible way of doing this is to have an algorithm that uses a pre-defined sequence of unequivocal and finite steps to search a data resource for patterns and connections which will help achieve a well-defined objective or resolve a specific user problem. This type of AI, which is restricted to the bounds defined by code, is called 'weak AI'. It is used for closed systems that are usually designed to support humans (e.g. automated diagnostic tools used in medicine, voice and face recognition, and for analysing the stock markets).

'Strong AI', by contrast, will self-develop after initial programming and a machine-learning and training phase. It uses deep-learning methods or artificial neural networks to develop cognitive abilities and, in the absence of a pre-defined objective or alternative solutions, sets its own cause-and-effect chains into motion. For instance, it will search the data resources available to it for new insights and respond to this information by self-developing the original algorithm. The decisions taken by strong AI are not the pre-defined product of unequivocal coding, but the result of an autonomous process.[2] 'Strong' AI is therefore characterised by the fact that it will leave behind the determined paths that it uses to begin with.

The following example from factory planning and control helps clarify the difference between 'weak' and 'strong' AI:[3]

The systems implemented in digital, interconnected factories (machines, human-machine-hybrids, robotic systems) are already aware of their own capabilities and objectives. They are able to interact autonomously with their environment without relying on a central control unit. A piece of production equipment will establish exactly what it is that the customer wants, either directly or using big data. It will know where to order the relevant parts, use IT and cloud systems to interact with service providers or suppliers and guide the materials supplied to the right place in the factory, using algorithms and the sensors fitted to the transport, unloading and logistics units. These are self-regulating machines and systems that are connected to one another vertically (within the company) and horizontally (across companies). A good example of a use case for this is robot-driven bottling strategies used in breweries to allow for customised production. A bottling machine will retrieve the customer's order from the brewery's online order service, order the ingredients needed to produce the right quantity of drink in the desired mixing ratio, ask the flexible system to mix the drinks and then call upon other autonomous production modules, such as the rinsing and bottling systems, the sealing and the labelling machines to do their work to create the finished product. If the machine finds itself short of ingredients, bottle parts or labels, it will order new ones, setting the necessary logistical processes in motion. The machine can even request any maintenance work it needs, using its sensor-driven cyber-physical systems. Guided by algorithms, the machine makes its way through a sequence of unequivocal and finite decisions. Soon, these distributed units that interact with one another will be able to optimise the relevant components, formulations and/or any transport and delivery routes relevant to the process themselves, right across the value chain (self-optimising machine). In our example, this would mean that the system will then be able to modify the list of ingredients, i.e. swap lemons for limes, if the background analysis of customers'

The term 'artificial intelligence' was coined by American computer scientist John McCarthy in 1956. He described AI as a field within computer science, dedicated to research into mechanisms used in intelligent human behaviour. At the time, his focus was on simulating game situations with the help of software used on a computing machine.

1 European Commission, High-Level Expert Group on Artificial Intelligence, 18 December 2018.

2 Ernst, JZ 2017, p. 1026, 1027.

3 Gesmann-Nuissl, InTeR 2018, Editorial.

preferences suggests that these now prefer a more sour taste and if 'lime' is one of the options available in the decision pattern.[4] Both variants of the system are a more advanced form of (supportive) factory automation, achieved by the addition of more commands. But the human code writer will always define the options available to the machine, which means that the system is an example of 'weak' AI.

As soon as a machine or system – e.g. the bottling installation – is able to take its own decisions that are no longer dependent on the initial command, but based upon large amounts of data collected by the machine itself and on the capabilities it has developed through machine learning – i.e. as soon as it can not only select the best of several options ('sour'= lemon or lime), but create new flavours independently, it qualifies as a self-learning machine. If the machine goes beyond identifying human preferences and also orders the necessary ingredients, adjusts its settings and control processes itself, conducts and monitors the production process autonomously, and modifies it if necessary, then this machine is an example of 'strong' AI. In this scenario, the artefacts are beyond human control, setting their own cause-and-effect chains in motion, defining their work and procedures themselves – including division of labour and processes involving other companies – and taking care of the final quality management, which is geared towards constant optimisation. These machines develop cognitive capabilities akin to human intelligence.

All experts agree that this type of 'strong' humanoid AI is a thing of the remote future. What we are seeing in real life today and what is likely to stay for the foreseeable future is 'weak' AI, which is associated with an enormous acceleration of the automation process and characterised by its supportive and (seemingly) smart behaviour.

4 *Philyra*, an AI software developed in cooperation with German perfumer Symrise and the IBM Thomas J. Watson Research Center, works in a similar way. It analyses data pertaining to the 1.7 billion fragrances in the Symrise database as to which have sold well on what markets, and uses this information to create new recipes. These are then sent straight to the lab or bottling robot which mixes the relevant ingredients following the recipe.

Artificial intelligence and legal personality

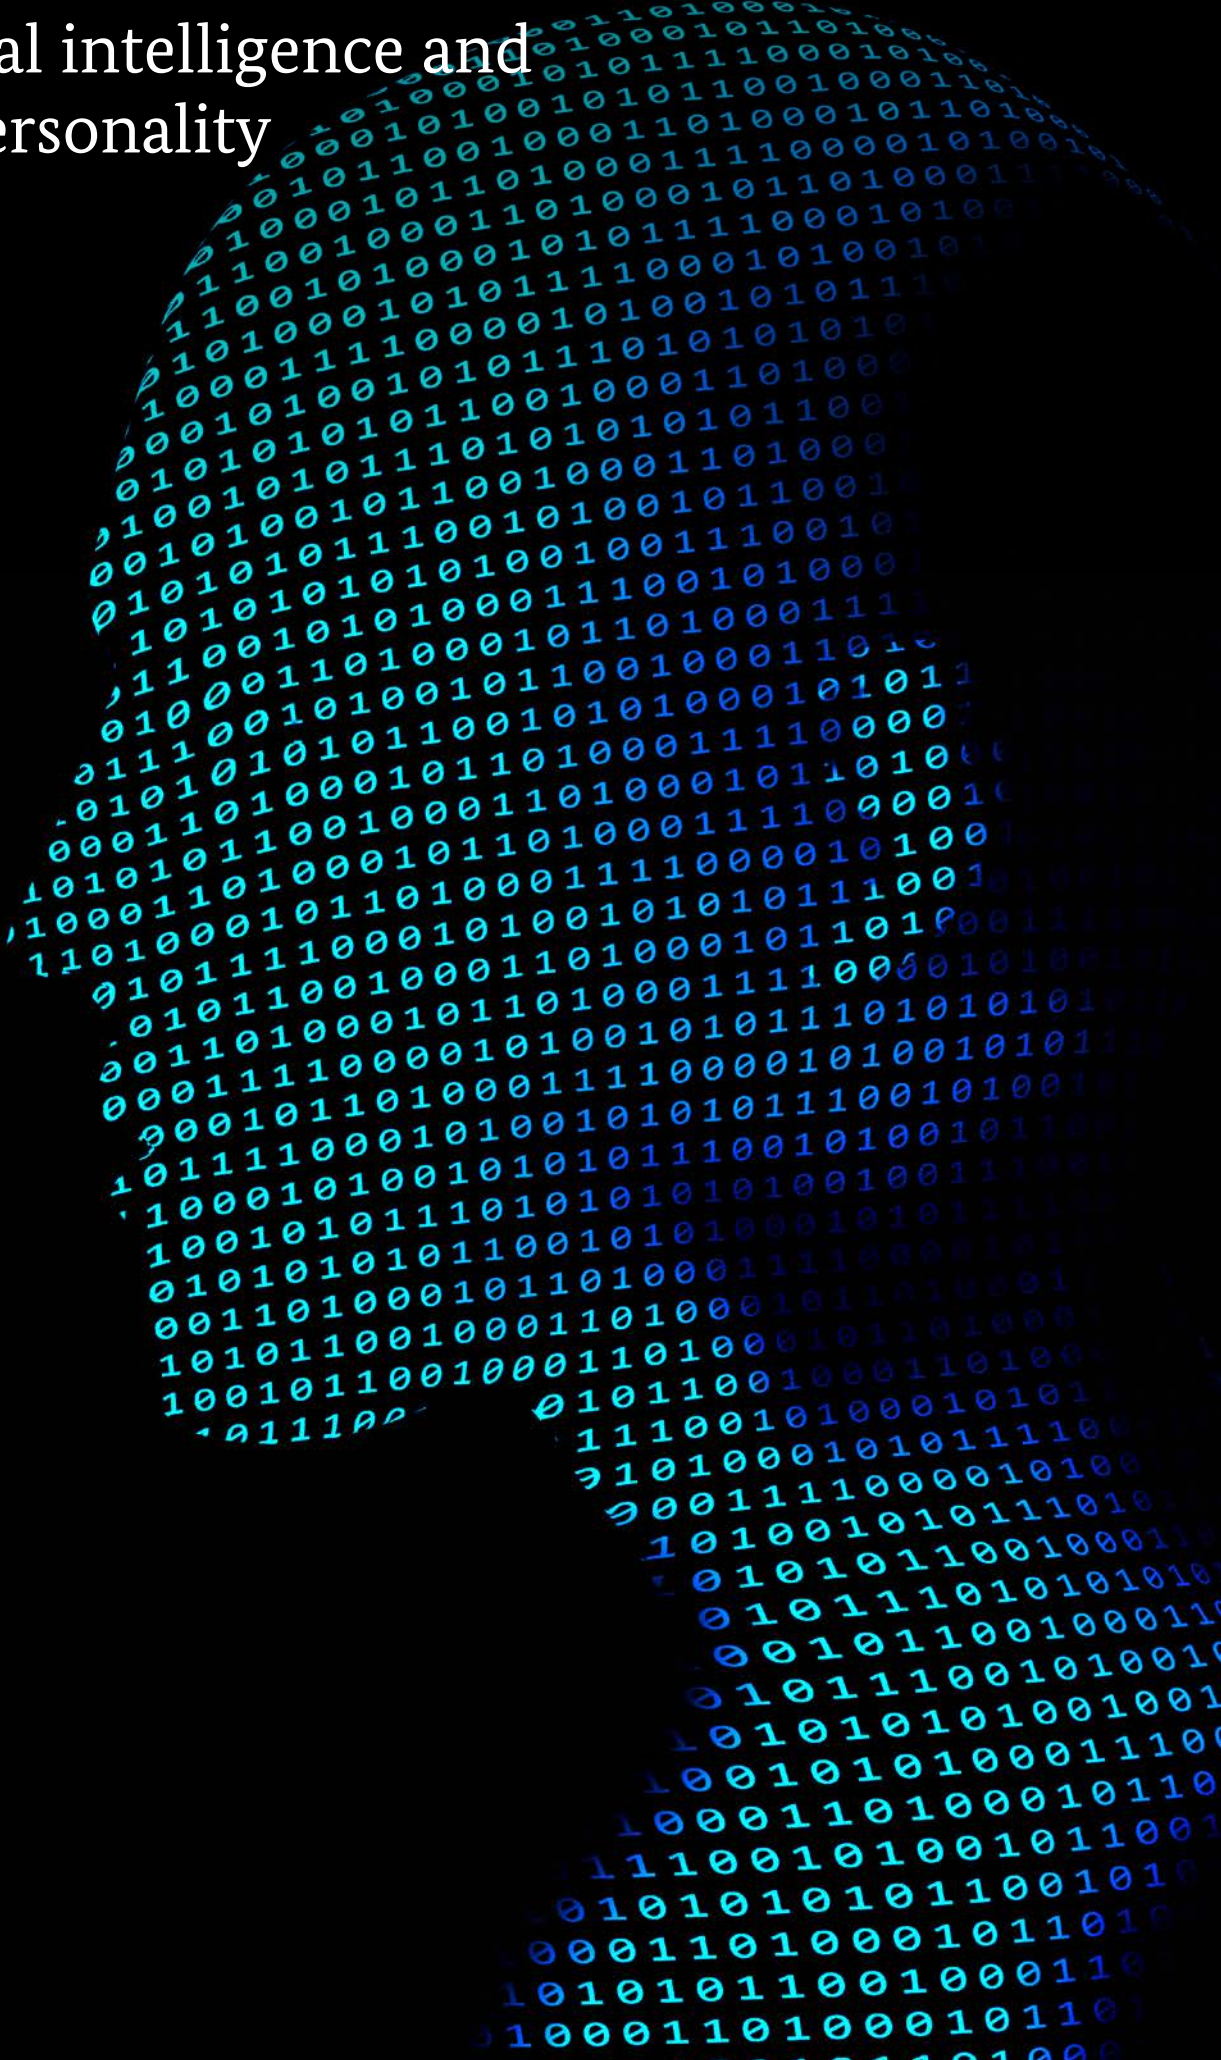

A: Factsheet

What is the issue at stake?

With automated and partly automated systems approaching human behaviour and capabilities (to a certain extent) and with interconnected systems creating a sense of high complexity and a lack of predictability that makes it hard to clearly assign and cleanly separate responsibilities and liabilities in cases of error or malfunction, there is a feeling that the legal system should be adjusted to take account of this new reality. One possibility that is being discussed is to make AI systems subjects to the law in their own right.[5] The European Parliament has therefore put forward a proposal for the creation of an 'electronic personality (ePerson)' for smart robots and autonomous systems.[6] It is questionable whether this move, which would make ePersons independent agents with their own rights and obligations, is really necessary, given the fact that it will remain possible to trace back the decisions taken by 'weak' AI to humans and their actions, as explained above.

Key questions

- Is there a need for a new type of legal personality that would make robotic systems agents complete with their own rights and obligations?
- Are there any gaps in the liability regime that would make it necessary for weak-AI systems to be given the status of a new type of legal personality?

B: Legal assessment

I. A note on what it means to be a 'legal person' under the existing law and on the closed system of 'legal personality':

In our system, a person that is subject to the law is an agent who has been declared capable of exercising rights and responsibilities (i.e. who has legal capacity). Under the current legal system, this applies to natural persons (i.e. humans that were born and are still alive) and legal persons (i.e. a collective consisting of persons and objects combined to form a permanent organisation serving a particular purpose). There are no other types of legal personality under our legal system.

That said, this is a reflection of political will and there is no reason it could not be changed. Early on, *Savigny* drew on Kant and the *zeitgeist* of his era when specifying the necessary characteristics of legal persons, resulting in an understanding whereby the definition of a 'subject to the law' had to be closely linked to that of a human being. Later, the legal definition of 'legal personality' moved away from this notion, towards a new understanding underpinned by a large number of legal and philosophical treatises[7]: the nature of a subject to the law is not pre-defined; all that matters is its ability to embody legal functions. As a consequence, it was then considered permissible for subjects to the law to be represented not only by natural persons, but also by social substrates of organisations, which were able to represent rights and obligations, even in the absence of a will, consciousness or emotions of their own. And even organisations without a personal substrate were accepted as representatives. In 1905, Radbruch summarised the situation rather succinctly:

> "To be a person is the result of a personification process undertaken by the legal order. All persons, be they natural or legal in nature, are creations of the legal order. Even natural persons are, strictu sensu, 'legal persons'."

5 *Bräutigam/Rücker*, E-Commerce, 14 B Rn. 4; Pieper, DSRITB 2016, 971, 978; *Mayinger*, Die künstliche Person, 2016, p. 65 f.

6 European Parliament resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics (2015/2103/INL), p. 28.

7 Et.al. Savigny, Gierke, Brinz&Hauriou, Radbruch.

Once it has been established that the creation of 'legal personality' comes down to a legal act (of attribution), it is definitely possible to have a debate about a potential new type of legal personality (ePerson) that might be introduced, although this would raise a number of additional questions to be answered.

II. Considering the issue of attributing responsibility and the legal consequences

Clearly, the first question that needs answering is whether there are any gaps when it comes to attributing responsibility/liability for decisions taken by 'weak' AI. Is it really necessary to introduce the concept of 'ePerson' for 'weak' AI? It is safe to assume that this will only be justified if it is established that there is indeed a problem with attributing legal responsibility for these actions, i.e. if there are no agents who can be held responsible for them. Whether or not this is the case needs to be explored separately for the fields of civil law, criminal law and public law. The focus of this work should be on the attribution of responsibility and liability in cases of technical failure and in cases where malfunctioning technical systems are placed on the market. This touches upon aspects of contract law (contracts concluded by automated or partly-autonomous systems; contractual liability), tort law (non-contractual liability as per Section 823 German Civil Code; product liability law, Road Traffic Act etc.) and product (safety) law.

As for the legal consequences, the question is to what extent an ePerson (if it can at all be found to be directly or indirectly responsible) could be held to account. What would be the legal consequences resulting from the introduction of an 'ePerson' and how would this affect the overall logic underpinning the legal system? Is it necessary for technical autonomy to go hand in hand with autonomous responsibility, i.e. would it be necessary for an ePerson to have recoverable assets attributed to it so as to allow for any damages to be paid? The ability to have property is inextricably linked to the ability to have legal capacity. The two are conditional on one another – as they are for the existing types of legal personality – which means that this would create another issue that would have to be resolved.

A more general point that would also require exploring is the question to what extent this type of legal personality would be coherent with constitutional and international principles. There are many legal texts, for instance, that establish a direct link between 'subjects to the law' and human dignity. Under the existing legal order, legal persons are expected to have attributes such as morality and a will of their own; these attribute play an important role and are presumed to guide the legal persons' behaviour. This is even the case for legal persons, on which legal capacity can be bestowed even in the absence of a personal substrate, as they have still been formed by natural persons joining forces. It would therefore be necessary to clarify how the relevant attributes could be implemented in the AI systems so that these can live up to the moral expectations upon which the legal system is predicated.

C: Recommendations for action

We currently do not see the need for an 'ePerson' to be introduced. All of the issues arising at this point in time can be resolved within the existing legal order (cf. the other factsheets). Current AI systems have not yet attained a degree of autonomy so extensive that it would no longer be possible to attribute their decisions to humans and their actions. The gaps in the current responsibility regime can be closed by amending existing provisions (cf. the other factsheets; this especially applies to the attribution of responsibility in the context of liability).

Artificial intelligence, access to data and data protection

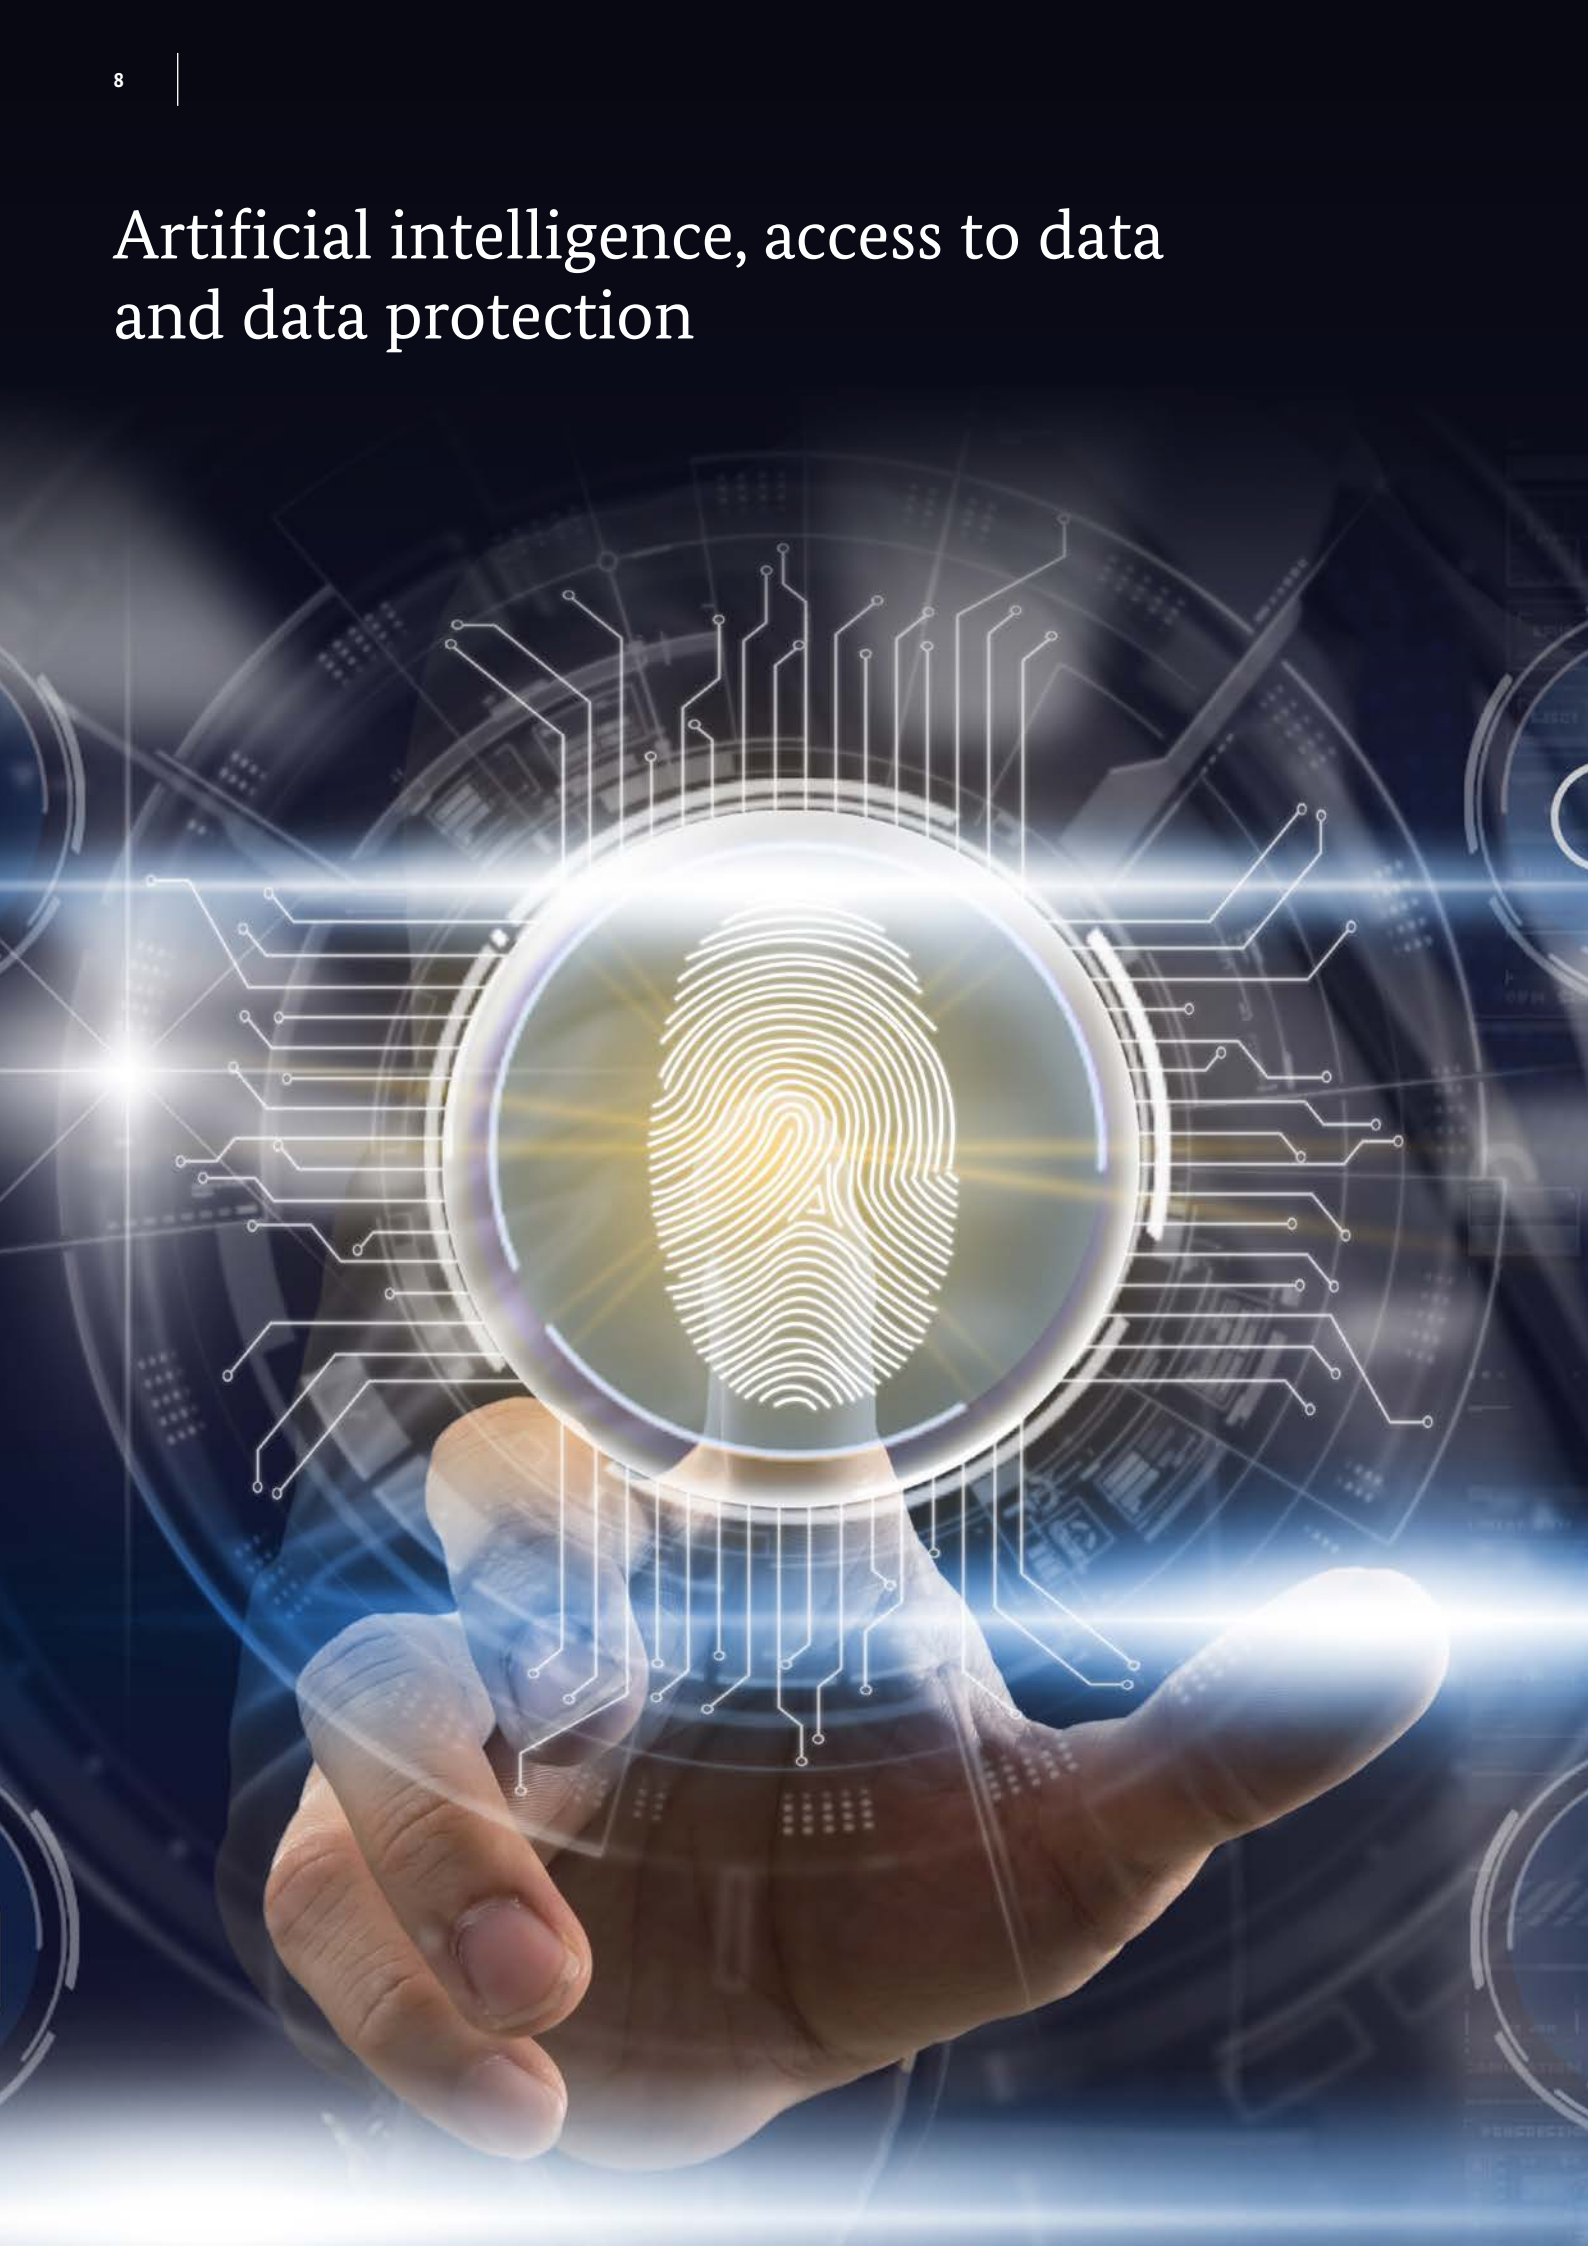

A: Factsheet

What is the issue at stake?

The mechanisms underpinning artificial intelligence are becoming increasingly important in the context of Industrie 4.0, especially with regard to quality assurance at the end of the manufacturing process, predictive maintenance, assisting robots, and also HR and continuing education. The ways in which humans and machines communicate with one another are also changing. Computer systems and robots now understand and 'speak' human language, for instance when they conduct error analyses for maintenance or support repair workers.

Pattern and speech recognition, machine learning and the development of AI applications all rely on the use of data. Data is nothing less than the essence of digitisation as a whole and AI in particular. What counts here is not only the availability of large amounts of data, but also its quality.

Much of the data used – even in the industrial context – is directly or indirectly personal in nature as it is linked to a specific, traceable natural person who interacts with the AI system. This means that the provisions of the General Data Protection Regulation (GDPR) apply – despite the fact that the natural person to whom the data can (theoretically) be traced back does not actually have a role in the processing of that data.

Other issues that need following up on in the context of AI besides the protection of personal data include transparency with regard to the algorithms used and their decisions, and access to the data that has been used to train AI. All recommendations given by AI-based systems should be transparent and understandable to the user. It must be guaranteed to the greatest possible extent that there are ways in which any incorrect or even discriminatory decisions taken by a system operating on an incomplete or one-sided database (machine bias) can be recognised and remedied.

What questions/challenges are there for Industrie 4.0?

The GDPR and the way it has been interpreted by the supervisory authorities impose what can be strict conditions for the use of data for creating and developing AI systems for Industrie 4.0. Among the questions that need resolving are the following:

- What are the advantages of pseudonomising datasets, in particular where different interests need to be weighed against each other?
- How can data that has been collected for a specific purpose (e.g. quality assurance) be used for other purposes (to train AI systems and optimise the supply chain) when the principle of 'purpose limitation' (Art. 5(1) lit. b GDPR) applies?
- How do the 'data minimisation' principle (Art. 5(1) lit. c GDPR) and 'data protection by design' (Art. 25(1) GDPR) impact the profitability of AI in the context of Industrie 4.0 applications?
- What about 'automated individual decision-making, including profiling' (Art. 22 GDPR) in the context of self-learning and autonomous systems?
- What would be useful amendments to the GDPR and its implementing acts to facilitate the use and improvement of AI systems for Industrie 4.0?
- What can be done to make AI systems more transparent and easier to explain?
- How can public-sector data be made more accessible, for instance to ensure that AI systems are trained with more solid data?

B: Legal assessment

I. Pseudonomisation of data sets

Pseudonomisation can be used as a special safeguard to minimise the risk of breaching data-protection law. Pseudonomisation makes it more difficult to attribute data to a specific or traceable person. The name of the person and/or all other identifying information is replaced with a code, e.g. a combination of numbers or letters. The allocation table used to generate the code must remain invisible so that only the pseudonym is present in the data being processed. This makes it possible for [personal] data to be put to separate, subsequent use, without compromising the rights of the persons affected.

However, this way of minimising risks will only work in actual practice if the supervisory authorities actually take account of the statutory exemptions that exist besides the one based on informed consent (e.g. data processing that is necessary for the purposes of legitimate interests (Art. 6(1) lit. f GDPR)). Unless this is the case, there will be no incentives for pseudonomisation or other actions that will help limit infringements of the general right to protection of personality. This is particularly true of Industrie 4.0, a field of application where much of the data being used is of a personal nature, but where this personal aspect of the data is not important or even without meaning for the purpose of data processing (e.g. for M2M communications).

When weighing different interests as per Art. 6(1) lit. f GDPR, pseudonomisation can be a means to guarantee that AI systems can access the data they need in order to gain new insights, whilst also ensuring better protection of the individual's privacy than would be the case if the person gave their informed consent.

Machine learning requires large amounts of data, and pseudonomisation is a suitable way of ensuring that innovative AI applications can be developed and that individual users can still be guaranteed a high level of data protection achieved by technical means.

Besides the GDPR, the future ePrivacy Regulation will also play a significant role, especially in regulating access to users' end devices. Here, too, it is important to provide for a secondary use of pseudonomised data instead of only allowing data to be used if users have given their consent and the data has been rendered anonymous. This is necessary so that the required amounts of data will be made available. The scope of the ePrivacy Regulation regarding non-personal data must be restricted to this effect, in order to make room for pro-innovation and competitive Industrie 4.0 applications in Germany. As for non-personal machine data used in an industrial context, it is true that this data must be kept confidential, but there is no need for an even higher level of protection to apply than the one stipulated in GDPR for personal data. Nor would this be helpful for innovation or conducive to the use of Industrie 4.0 in Germany.

II. Data collected for a specific purpose

This is another area where Plattform Industrie 4.0 holds the view that it would be helpful if the supervisory authorities were to approach the relevant provisions by weighing the different interests at stake against one another. The provision restricting the use of data to the specific purpose for which it was collected is designed to prevent personal data from being used in ways that are impossible to understand or trace. This can also be achieved through technical safeguards (such as pseudonomisation). This is why data processing practices that offer this kind of protection should be adequately taken into account during the assessment pursuant to Article 6(4), which is designed to ascertain whether the purpose for which the data is to be used is compatible with the initial purpose. Furthermore, it should be clarified in the interpretation of the GDPR that training AI systems is to be recognised as a sufficiently unequivocal and legitimate purpose.

III. The principles of data minimisation and data protection by design and their impact

It will certainly be necessary to create an evidence base for navigating the balance between AI systems' need for data and the principle of data minimisation. But it seems that one aspect that is generating a lot of interest is the potential – and not unresolvable – conflict of interest between the principle of data minimisation and AI systems' need for large amounts of good quality data. A sound balance is needed also to ensure the desired level of public acceptance and, by extension, the positive effects resulting from the use of AI.

At the same time, there can be no doubt that the bar used in a mere industrial context cannot be the same as for the use of AI by consumers. Once again, and on the basis of Recital 78 GDPR, it would seem necessary to come to the

conclusion that risk-minimising factors such as pseudonomisation can be used to bring actions in line with the protective purpose served by the principle of data minimisation.

IV. Automated individual decision-making, including profiling in the context of self-learning and autonomous systems

In principle, there is a particular interest in ensuring that users are able to control how their data is used. This also means educating users about data-driven business models and about how users of Industrie 4.0 can exercise their right to digital autonomy.

Irrespective of the relevant legal obligations, Plattform Industrie 4.0 holds that it is essential to be transparent vis-à-vis the users of AI applications. Any decision based solely on automated processing is only permissible if the user has been informed of this as part of a relevant process or in an agreement (cf. Article 22 GDPR).

However, this must not result in a situation whereby automated decision-making processes based on self-learning systems cannot be used as a result of Art. 22, even though they do not have any direct impact on the data subjects. This is another situation in which it would be important to ensure that protective action designed to afford adequate privacy protection for individuals is given a positive assessment.

V. Amendments to the GDPR and its implementing acts to facilitate the use and improvement of AI systems for Industrie 4.0

The 2020 evaluation of the GDPR ought to acknowledge the fact that data diversity and data protection are not mutually exclusive. It has to be recognised that the current version of the GDPR does not always strike the right balance between the individual's right to protection and other interests worthy of protection. Medical diagnostics, for instance, relies on a database that is as broad as possible. If data protection rules are interpreted in an overly repressive way, this also means that individuals cannot benefit from the relevant medical insights. There is now broad public consensus on the fact that a certain volume of data is needed to ensure that the digital transformation can continue to deliver positive change for all of society. For the purposes of Industrie 4.0, in particular, it is necessary to distinguish more clearly between cases whereby a person's general right to protection of personality may be violated and cases whereby the data may be of a personal nature, but only used to a marginal extent (as part of 'industrial' data processing).

Against this backdrop, the evaluation will have to call into question the dominant role of the principles of data minimisation and purpose limitation underlying the GDPR, and shift the focus towards transparency regarding (automated) decision-making. It is important to ensure that this is not understood to mean that every operation based on an algorithm must be transparent to the last detail, but rather to mean that it must be possible to explain it in principle.

VI. Transparency and explaining AI systems

Only if AI is used in as transparent a way as is possible, will it be possible to build people's trust in autonomous, decision-making systems. For transparency's sake, the GDPR provides for wide-ranging information requirements and a right for data subjects to have an automated decision reviewed by a human being.

In principle, customers ought to always know whether they are dealing with a person or an AI system. Users must also have clarity as to what customer data is used by the AI systems. Any blanket requirement for the use of AI implementations or trained AI systems to be disclosed ought to be avoided as it would heavily encroach on business secrets and IPR, and thereby hamper the development of AI.

There should be more support for research in the field of accountability and to make AI easier to understand, so as to spur the development of technical solutions to improve transparency. The same applies for the development of ways to raise awareness of the ethical boundaries of AI, something that is needed for both developers and users.

VII. Access to public-sector data

Machine-learning for AI systems is only possible if training data is available. Whether or not AI is successful much depends upon this. The data must be of good quality, credible, up-to-date and accessible in uniform formats that are also machine-readable. Ultimately, AI applications will never be better than the data that has been used to train them. One important source of training data that has so far remained largely untapped is the public authorities, which collect and store huge amounts of data.

We therefore welcome the efforts undertaken by the European Commission to promote the joint use of data in the public sector, for instance by amending the Public Sector Information Directive (PSI Directive).

Public-sector data ought to be placed on open data websites in a standardised and machine-readable format, so that it can be used by AI developers and users. The standards necessary for this ought to be developed and bottlenecks preventing access to public-sector data removed.

However, it would not be helpful to go even further and stipulate access rules for data held by the private sector. Any provision making it mandatory for companies to open up their data stocks would come with an inherent risk of promoting unwarranted data minimisation, especially for data collected by machines – a development that would slow down the development of digital technologies. Moreover, there are no indications of market failure that would warrant this kind of regulation. There already is successful cooperation between telecommunications companies and public-sector authorities, for instance for the development of analytical tools that draw on location data to allow for smart traffic management.

The principle of contractual freedom and also a voluntary approach to open data ought to prevail in AI as well. Companies must be free to decide who they want to be able to access their non-personal data and under what circumstances – be this as part of a B2B data partnership or

by way of a contractual agreement with public-sector authorities.

C: Options and recommendations for action

When weighing different interests as per Art. 6(1) lit. f GDPR, pseudonomisation can be a means to guarantee that AI systems can access the data they need in order to gain new insights.

Furthermore, it should be clarified in the interpretation of the GDPR that training AI systems is to be recognised as a sufficiently unequivocal and legitimate purpose.

On the basis of Recital 78 GDPR, it would seem necessary to come to the conclusion that risk-minimising factors such as pseudonomisation can be used to bring actions in line with the protective purpose served by the principle of data minimisation.

Irrespective of the relevant legal obligations, Plattform Industrie 4.0 holds that it is essential to be transparent vis-à-vis the users of AI applications. Nevertheless, it must still be possible to implement automated decision-making processes based on self-learning systems, wherever these decisions do not immediately affect the data subject.

The 2020 evaluation of the GDPR ought to acknowledge the fact that data diversity and data protection are not mutually exclusive. It needs to be recognised that the current version of the GDPR does not always strike the right balance between the individual's right to protection and other interests worthy of protection.

There should be more support for research in the field of accountability and to make AI easier to understand, so as to spur the development of technical solutions to improve transparency. The same applies for the development of ways to raise awareness of the ethical boundaries of AI, something that is needed for both developers and users.

Public-sector data ought to be placed on open data websites in a standardised and machine-readable format, so that it can be used by AI developers and users. The standards necessary for this ought to be developed and bottlenecks preventing access to public-sector data removed. However, it would not be helpful to go even further and stipulate access rules for data held by the private sector.

Liability

A: Factsheet

What is the issue at stake?

The development may still be in its infancy, but Industrie 4.0 is beginning to use IT systems based on artificial intelligence. Artificial Intelligence (AI) is a blanket term for computer applications that replicate smart behaviour and are capable of machine learning. AI comes in two forms: expert systems and neural networks. Expert systems draw on the knowledge of human specialists in designated fields, which has been translated into computer-based models and rules. All output of expert systems is pre-defined by these rules, which also means that it can be fully understood and explained. Neural networks are loosely based on the design of the human brain and consist of several layers of interconnected artificial neurons. The systems are then trained with algorithms and learn to recognise patterns and correlations in a given database. Decision algorithms then make it possible for the systems to apply the learning outcomes to new databases. The more complex and multi-layered the neural network, the more difficult it is to understand the processing and decision-making mechanisms used by the system. The number of parameters, variables, and interdependencies within a neural network is so immense that it is impossible to predict exactly which learning processes and output patterns it will produce; all estimates can only be given within a certain range of potential outcomes.

Self-learning systems can only function properly if they are based on a suitable neural network, if the learning and decision-making algorithms are coded flawlessly, and if a necessary amount of unbiased training data is available. Artificial intelligence is currently mainly used in fields where large amounts of data need to be quickly analysed for certain patterns that can be used for decision-making and/or recommendations for action. At present, the number of scenarios for using AI in industry is very limited, with most of these restricted to data analysis and process optimisation, where they are used to support human users.[8]

None of the legal analyses undertaken by Plattform Industrie 4.0 exploring Industrie 4.0 scenarios have found any regulatory gaps in the current liability regime. There was just one situation for which policy-makers were encouraged to explore the possibility of amending the Liability Act: damages occurring at an Industrie 4.0 manufacturing site. The results are to be reviewed and published in this factsheet focusing on Industrie 4.0 scenarios that rely on AI.

What questions/challenges are there for Industrie 4.0?

- What liability regimes are available? – some preliminary thoughts
- Does the Act on Liability for Defective Products apply to AI systems (in other words: is an AI system a product within the meaning of the Act on Liability for Defective Products)?
- Are manufacturers liable for the AI they use (liable within the meaning of producer's liability)?
- On whom is the burden of proof in cases of damage caused by AI (causality)?
- Do the findings of Working Group 4 stand up, considering this background?

B: Preliminary ethical considerations

Human beings develop AI systems and use them to serve their purposes. However, this must not result in a situation whereby liability for the use of AI and the legal consequences are shifted to the technology. For as long (natural or legal) persons are held liable for the consequences of the use of AI, it will be in their own best interest to control and monitor the AI system and its output. To the extent that there is public consensus that liability ought to continue to rest with human beings, even where AI is being used, it will be possible to uphold the existing principles of liability law (perhaps with some minor modifications).

8 A survey conducted by the Association of German Engineers (VDI) among its members found that only 7% of the responding SMEs and 12.6 % of the large companies surveyed use AI, and where so, mostly for data analysis (cf. "Künstliche Intelligenz – VDI-Statusreport Oktober 2018", p. 21).

C: Legal assessment

I. What liability regimes are available? – some preliminary thoughts

For the most part, manufacturers of products that make use of AI are subject to the liability regimes of product and producer's liability law.

Product liability law: Under the Act on Liability for Defective Products, a manufacturing company is held accountable whenever a product it has put on the market is defective and where this defect has resulted in the death or injury or illness of a person, or in damage to an object intended for private use or consumption. Liability under the Act on Liability for Defective Products arises mainly in the following cases:

- Design flaw (product design not compatible with technical knowledge):
- Manufacturing error (manufactured product does not comply with design specifications for that product);
- Instruction error (insufficient user instructions or lack of hazard warning).

Producers are only liable if the defect causing the damage was already present when the product was placed on the market.

Producer's liability: Producers are liable if a legal right protected under Section 823 (German Civil Code) is violated by a manufactured product. This only applies if the manufacturer is culpable of this violation. This is the case if the manufacturer has failed to comply with his/her obligations to provide for safety. The following is a list of the obligations to provide for safety recognised under German law:

- Organisational obligations (the manufacturer must organise his/her company in such a way that errors/defects are spotted and corrected in checks and controls);
- Obligation to instruct (the manufacturer must inform users about the correct way of handling the product and about any hazards);
- Product oversight obligation (the manufacturer must investigate any information received about errors/defects and hazards emanating from his/her products); and
- Preventative obligations (the manufacturer must remedy any hazards recognised, if necessary by removing the product from the market).

Under Section 823(1) German Civil Code, producers are liable for injuries to physical integrity, health and property and also for infringements of rights that are equivalent to the right to property, the right to the operations of an established commercial business, to the extent that the damage has been caused by a violation of the general right to protection of personality. Even Section 823(1) German Civil Code, however, does not stipulate for damages in cases of mere pecuniary loss.

The next section will explore how these legal principles would have to be applied in cases in which damages have been caused by products that make use of AI.

II. Does the Act on Liability for Defective Products apply to AI systems (in other words: is an AI system a product within the meaning of the Act on Liability for Defective Products)?

Product liability law applies to moveable goods. There is major controversy about whether AI systems, which are made up of computer algorithms, i.e. software, are to be classified as moveable goods within this meaning[9]. Currently, software is not classified as a product within the meaning of product liability law. This means that, for this reason alone, AI developers cannot be held liable as producers under the Act on Liability for Defective Products.

At the same time, product liability law only applies in cases where there has been damage to physical integrity, health, and property. For the AI coder to be held liable under the Act on Liability for Defective Products, the – immaterial –

9 Cf. et al. the Preliminary Preliminary Concept Paper for the future Guidance on the Product Liability Directive 85/374/EEC of 18 September 2018 and the Comments by the German Bar Association's Committee for Information Rights of November 2018

AI itself would have to be able to have an impact of its own on these legal rights 'in the real world'.

Given that immaterial AI can only indirectly cause violations of the legal rights cited above 'in the real world' – i.e. through 'tools' such as the Internet of Things or in the form of embedded software – the producer of the product will be the first to be held liable.

Whenever a producer uses AI (coded in-house or by external coders) to steer or control moveable goods that he/she places on the market (e.g. a robot), he/she is liable for any damage caused by that robot – irrespective of whether the AI used by the robot has been coded in-house or not.

If a producer is held liable for products that use AI that has been coded by an external coder, he/she may be entitled (under contract law) to make a recourse claim against the coder. However, much of the software used is open source (OSS), which is subject to a regime that excludes liability and is not rooted in German law, and usually developed by a large number of coders, which makes it impossible to attribute errors to specific individuals. As a result, any recourse claims are likely to be unsuccessful.

This liability regime is understandable given the speed of innovation in OSS, and it does not place producers that consciously decide to opt for OSS at un unfair disadvantage. Section 3(1) no. 2 Act on Liability for Defective Products also affords protection for any "use to which it could reasonably be expected that it [the product] would be put".

If the AI was produced by a company different from the producer of the physical product that "takes action" through the AI and causes the damage, and if the use of the object "by AI" was already factored in (e.g. presence of an add-on feature and/or AI coder asked to write code for that purpose by the producer of the physical product) and if this happened without any additional safety precautions being taken, product liability law also applies. The question is whether the producer of the physical product can claim joint and several compensation from the AI coder (pursuant to Section 5 Act on Liability for Defective Products) – a question that can only be answered once it has been established if product liability law is applicable to AI in its capacity as a type of software.

In cases whereby a physical product has been "hijacked" by AI and this could not have been anticipated (unlike in cases of 'foreseeable abuse'), the product would not be considered

to be defective. What remains is liability under the general provisions of tort law, just as for other illegal actions (including hacker attacks). Given the fast pace of digitisation, the legislator might want to consider widening the scope of the kind of mandatory security requirements that already exist for critical infrastructure to include other areas, which would also include a requirement for security updates to be made by operators or manufacturers of AI.

Some argue that it is all too easy for a manufacturer to exculpate themselves in cases in which a legal rights violation results from a decision taken by AI for which the AI has only acquired the necessary capacity as part of its learning after it was placed on the market. Section 1(2) no. 2 Act on Liability for Defective Products excludes liability for damage resulting from defects that only occurred after a product was placed on the market. However, it is possible to argue that the error consists not in the knowledge acquired as part of the learning process but in a flaw in the coding which made it possible for the AI to take the relevant decision. Others are exploring a teleological interpretation of the Act, which could mean that defects rooted in the learning

process of the AI are not covered by the 'later defect defence' (cf. e.g. Preliminary Concept Paper, para. 59ff).

Given that not only the coding process for the algorithms (construction of the AI) is prone to mistakes, but also the learning process, which may be conducted with insufficient or substandard data, it might be worth reconsidering the interpretation of 'defect' in the context of AI. There are also very few rules and best practices that are generally accepted and that could be used to test whether an algorithm or a stock of learning data is sufficient and adequate. However, this type of difficulty when interpreting legislation and setting relevant technical standards arises with every new technology. It is not specific to AI and has traditionally been resolved through the development of case law.

Finally, there is the question as to whether a product fitted with AI is to be considered to enter the market afresh as soon as there has been a software update, even if this update has only served to fix a bug. If this were the case, this would expand the scope of product liability beyond what is currently set out in the relevant legislation. This means that we need a very clear definition that makes a distinction between additional AI functionalities resulting in a new product and mere repair work. As soon as this issue has been resolved, the current criteria under product liability law can be used to achieve a solution that is legally unobjectionable.

III. Are manufacturers liable for the AI they use (liable within the meaning of producer's liability)?

Unlike with product liability, producer's liability (Section 823(1) German Civil Code) can only be triggered and the manufacturer held liable if the relevant damage has been caused with intent or through negligence.

The damage must be attributable to the manufacturer, i.e. it must have resulted from a violation on the part of the manufacturer. AI, however, is specifically designed to produce end results that cannot be foreseen in detail. It is doubtful whether a manufacturer can be accused of negligently causing specific damage if he was not able to foresee this damage and therefore unable to prevent it. However, this cannot be interpreted to mean that the manufacturer benefits from a blanket exemption from any liability as per Section 823(1) German Civil Code. At the most, the fact that specific damage cannot be foreseen can mean that the manufacturer is released from their obligation to warn the user about the potential damage in advance. The manufacturer continues to be obliged to carefully organise the coding process, monitor the product and eliminate hazards once they have been recognised, and to be liable for any damage caused by failure to perform these duties. Furthermore, a manufacturer that uses self-learning AI is certainly in a position to foresee that the AI may produce unforeseen outcomes and take unforeseen decisions, which in turn may cause new hazards and risks.

A solution for self-controlled vehicles will be based on the hazard liability rules enshrined in the German Road Traffic Act. This solution, however, is very limited in scope as it cannot even be used for self-controlling forklift trucks or transport vehicles used in production and logistics facilities where the German Road Traffic Act does not apply, resulting in a lack of a clear liability regime. Producer's liability, which requires an element of fault, does not deliver adequate outcomes here. There have recently been calls to introduce a form of strict liability analogous to Section 833 German Civil Code (strict liability for animal owners), so as to prevent the formation of a liability gap. Strict liability is based on the notion that whoever legally establishes and exercises a hazardous operation for their own benefit must pay for the damage incurred by another person as a result of the risk inherent in that operation, to the extent that this other person cannot prevent the damage.

It is true that there are some similarities with the example above. Whoever deploys a system that uses AI and thus engages in autonomous learning and decision-making, has made a conscious decision to use a system that comes with inherent risks for which they should be held liable.

Our advice is that the legislator should closely monitor these developments.

IV. Do the findings of Working Group 4 stand up, considering this background?

Working Group 4 first approached the issue of product liability law in the context of Industrie 4.0 in its publication "Wie das Recht Schritt hält" [How the law is keeping up] of October 2016.

At the time, there was a basic consensus in the sub-working group that there were no gaps in the existing legal framework which would have required the legislator to intervene immediately. The group formed the view that product lia-

bility law on the one hand and general tort law on the other afford sufficient protection for the 'protagonists' (especially in the event of a cyber attack). At the time, the group spoke out against introducing strict liability rules akin to those of Section 7 Road Traffic Act (vehicle owner's liability) or of Section 833 German Civil Code (animal owner's liability), but agreed that it was necessary to keep an eye on the situation.

A (potential) need for other solutions were discussed only for damage occurring in the context of the manufacturing process itself and in the absence of a causal contribution to the act attributable to a person involved in this process: if the (German) employers' liability insurance system for employees does not cover such situations, it might be necessary to make some careful adjustments to the Liability Insurance Act.

The sub-group stands by this view, including in cases in which AI is used.

With regard to the first two questions, there is no difference between scenarios using AI and other 'Industrie 4.0' scenarios. Both scenarios feature a physical product that causes direct physical harm. The question about the causal contribution of AI and/or the coder's liability is no different from the question as to the extent to which a manufacturer of a part is to be held liable. Whether or not this is the case under product liability law depends on whether software is a product within the meaning of product liability law. There are no AI-specific aspects to this. The legislator ought to actively steer the deliberations and discussions at European level and thus contribute to the development of positive solutions.

The analogous question relating to cyber attacks is the following: who is to be held liable for accidents provoked by an external intervention of AI in an Industrie 4.0 manufacturing site, for instance if one or more steps in the manufacturing process have been disrupted? Again, the results are no different. Unless the AI intervention constitutes a form of foreseeable abuse and unless there is an obligation to provide for the best possible level of IT security at all times (cf. II above), there is no need to hold the operator liable in any other way as would be the case if the unlawful intervention had been caused by a (human) hacker. In individual cases, it may be true that the use of AI makes it more difficult to attribute the relevant action to a subject that is liable and to enforce a tortious claim than would be case with a "human delinquent", but this alone does not suggest that it would be wise to suddenly hold the operator of a plant liable in such cases. The decision to launch an attack was certainly not taken by the plant's operator.

In this last scenario, too, the mere fact that AI is involved does not result in a different outcome of the appraisal as the prima causa is obviously unknown. Any gaps in the liability regime could be filled by way of developing a strict liability regime if necessary (perhaps complete with insurance models).

D: Options and recommendations for action

In actual practice, the use of AI systems might make it more difficult to establish and attribute liability, compared to a conventional product. This is due to the fact that the establishment and use of an AI system has more of an impact on potential infringements committed by the system than would be the case with conventional products (especially where the system is populated with user-generated data). It is also more difficult to establish the exact cause of a damage.

At the present time, it is almost impossible to imagine how an AI system could cause infringements without human intervention ("software alone cannot cause bodily harm"). For this reason, the existing liability regime is sufficient and can be applied to AI systems. At this point, there is no need for the introduction of a separate provision establishing strict liability for AI.

AI-generated IP

A: Factsheet

What is the issue at stake?

German businesses regularly excel in global competition on account of their creativity, which, often and in many fields, translates into a competitive edge. But creative achievements tend to turn into marketable goods especially once they are protected in a way that guarantees the company an exclusive right to use them – be it through patents, design rights, copyrights or other IPRs.

The use of artificial intelligence (AI) might result in a situation whereby this type of creative achievement could be made increasingly by machines rather than humans. Notwithstanding the rather philosophical question as to whether machines have the capacity to be creative (lat. creare ~ "to create something new"), there has to be a legal debate about whether the outcomes of an AI working process can benefit from traditional IPRs or whether a company that uses AI for creative purposes might be at risk of being unable to adequately protect the outcomes of this production process.

What questions/challenges are there for Industrie 4.0?

- Under what circumstances should the outcome of a work process be considered as "generated by AI"; under what circumstances should AI be seen merely "as a tool" used by a creative human mind?
- Is artificial intelligence capable of "inventing something new" within the meaning of patent law?
- Is it possible for AI to develop a personal intellectual creation within the meaning of copyright law?
- What are the options for protecting AI-generated outcomes of a work process under the existing legal framework?

B: Legal assessment

The 'machine v. man' debate that was touched upon in the previous section is not only an all time favourite of science fiction, but has also long been the subject of legal debate. In 1964, Fromm, in the context of his writings about 'art robots, wrote the following': "The slave revolt staged by the automatic devices which we designed for our pleasure, to improve our lives and make life less cumbersome, is in full swing" (Fromm, GRUR 1964, p. 304, p. 306).

The difference between the art robots, drawing and composing machines of that time and AI lies not in the lack of foreseeability of the results (the products of 'random art' were also not foreseeable), but in the greater degree of autonomy that must be attributed to AI. Whilst the debate of 1960s centred on apparatus that conducted mere technical calculations, comparisons and classification, or randomly arrived at certain results without actually creating things of their own (cf. Fabiani, GRUR Ausl. 1965, 422, 423), it would appear that artificial intelligence in the 2010s has reached an altogether new level.

It is against this backdrop that the older views formed on the role of technology in the creative process must be challenged and revised.

I. Under what circumstances should the outcome of a work process be considered as "generated by AI"; under what circumstances should AI be seen merely "as a tool" used by a creative human mind?

Prior to answering this, it has to be established whether today's artificial intelligence really is different from the creations generated by apparatus in the last century. Whilst this factsheet certainly cannot provide a full answer to this question (not least because this is a matter that requires biological, medical, neurological and IT expertise more than jurisprudence), it can at least be established that AI must not be misunderstood as denoting a silicon-based version of our own carbon-based intelligence, but must be classified as 'something other':

What is special about artificial intelligence – and also explains its rapid developments in recent years – is the formula of mathematical procedures combined with near infinite amounts of data made available through 'big data' and with algorithms that are ever more complex. The database to which AI can have access to 'take decisions' and instigate actions is many times larger than it was 20 years ago.

This is why AI – unlike human intelligence – still continues to rely on mathematical rules for its results. The method used by AI, its connection with big data, and its capability of using algorithms for self-learning have, however, resulted in a perception (at least by outsiders) that likens AI more to autonomous mental work than the simple 'if... then...else' strings used by the venerable apparatus of the twentieth century.

To illustrate this, it might be useful to draw a comparison with the 'ape selfies' phenomenon (cf. König/Beck, ZUM 2016, p.34 ff. for more detail): A photographer gave several apes access to a camera, which the primates used to take selfies. As apes do not have legal capacity (which de lege lata is the same for AI), they did not have rights to the photos. Whether or not the photographer has IPRs to the photos depends on whether the ape was to be regarded as the photographer's instrument, i.e. whether the outcome (the selfie taken by the ape) was to be more or less expected (then, yes), or whether the outcome was the result of an autonomous decision taken by the ape (then, no).

Despite the fact that the decision-making processes of AI are rooted in maths, their much higher complexity compared to traditional IT systems makes them more similar to those of autonomous beings than IT systems. The reactions of animals (and those of AI) to certain life situations are similarly impossible to predict for third parties as human reactions. This is another reason why there is a debate on whether the liability regime for AI should be analogous to Section 833 German Civil Code (liability of animal owners) (to learn more about the current state of the debate, cf. Borges, NJW 2018, p. 977; 980 f.).

Whether AI is considered to be performing 'actions of its own' or to be used as a tool by humans must therefore depend on its ability to act autonomously in a given environment. If the AI is operating in a fully controlled setting and therefore produces results that are ultimately foreseeable or if the differences between the potential outcomes are very limited, i.e. if the outcome depends on the creative potential of whoever 'deploys' the AI, the latter is used by the human originator/inventor as a tool that allows him to fulfil a certain objective defined by the human originator/inventor.

If, by contrast, the AI were to be able to act autonomously and if it were impossible to predict or at least plan the outcome on account of the setting in which the AI was being deployed, meaning that the AI operator's contribution would be ancillary, it would no longer be possible to attribute the 'creation' or 'invention' made by AI to the operator.

A regular objection to this bifurcated approach is the argument that a person presented with the outcome will usually not be in a position to know about the setting in which the AI was used. Whilst this is true, it is ultimately a mere 'point of uncertainty regarding the circumstances' – something that is not alien to IPR law. There will also be individual cases of simultaneous invention or creation, adding elements of uncertainty. Whilst the proposed approach might not always deliver clarity in terms of the legal situation ex ante, it will always be able to do so ex post. This does not mean that the approach is 'poor', it just goes to show that assessing real-life circumstances is always fraught with general risk, which ultimately, is an issue of fact.

II. Is artificial intelligence capable of "inventing something new" within the meaning of patent law?

First of all, the question as to whether AI is capable of "inventing something new" must be seen as separate from the question of whether the AI itself is patent-protected. This second question, which is not to be further explored here, is about whether patent-protection can be granted to those creating AI and, more specifically, to what extent the mathematical procedures and algorithms underpinning AI can be patent-protected (the answer to this question will depend on each individual case and is everything but trivial). Not least, it touches upon the issue of whether computer programs in general can benefit from patent-protection (cf. Section 1(3) No. 3 and 4 Patent Act and, more specifically BeckOK PatR/Hössle, 8th ed, 16.4.2018, PatG § 1 Rn. 189 ff.).

The question to be addressed here is separate from that as to whether AI itself can be patent-protected and is about whether outcomes produced by AI are eligible for patent protection and who is to be recorded as the inventor. Our starting point here must be a basic, unequivocal under-

standing, i.e. that IPRs can only exist for human achievements. There is therefore a general understanding in patent law that the process of invention is always the work of one or more human beings (Kraßer/Ann, Patentrecht, 7th ed. 2016, § 19 Rn. 7). This is on account of the ultimate purpose underpinning patent protection: A 'patent', i.e. an exclusive right to something is granted mostly by way of recognition for a special achievement in technology and is regarded as a service rendered in exchange (or as a reward) for the inventor's contribution to technical progress and to collective technical knowledge – and it also serves as an encouragement to make further such contributions (Federal Court of Justice GRUR 1996, 109, 114 – Klinische Versuche). Encouragement and reward, however, are genuinely human in quality (at least for now) and cannot be used as grounds to grant AI the status of an inventor. Whilst it is true that this would create an incentive for creating and investing in AI, this would bring us back to the question of whether AI in itself can benefit from patent protection and away from the question about patent protection for AI-generated outcomes.

At most, it might be possible to argue that the outcomes generated by AI are somewhat inherent in the AI itself, which would mean that any potential patent protection for the AI would then extend to its derivatives, i.e. all the outcomes that are eligible for patent protection and generated by the AI at a later point in time. This would mean first of all that the inventor is not the AI but the code-writer (read on for more about the potential legal status of the code-writer as an inventor). Second, it is important to note that the principle of sufficiency of disclosure, an important concept in patent law, only applies if both the process used to arrive at the derivative invention and the derivative invention itself are disclosed in the documentation submitted as part of the patent application. It is therefore impossible for abstract IPRs to be granted to derivative inventions made by AI when these inventions are not known in sufficient detail (Hetmank/Lauber-Rönsberg, GRUR 2018, p. 574, p. 577). We can therefore say that the current legal situation and the basic principles underpinning today's patent law do not allow for AI to 'invent'. Inventions are made by natural persons. This brings us to the next problem, namely that of who the inventor of an AI-(co-)generated outcome is: is it the user of the AI system, or the code-writer, or perhaps the owner?

It has been pointed out already that the use of software and computers for generating inventions is not a development of the 21st century. Computers have long been used as tools to assist with inventions. There is a general consensus that the 'inventor' of a computer-aided invention is the person who has arrived at and recognised the solution of the technical problem through coding and analysing the computer's output – as opposed to the person(s) who have constructed or own or possess the computer or who otherwise operate it (Benkard PatG/Melullis, 11th edition, 2015, PatG § 6 Rn. 32).

This has the following immediate implications for the use of AI: the smaller the extent to which the AI acts autonomously in the environment in which it has been deployed, and the stronger the influence of the human (and obviously technically skilled) 'user' of the AI on the creative process and the process of analysing the results, the more logical it would seem to abide by the established principles of computer use and regard the human 'user' as the inventor.

By contrast, the position of the 'user' as inventor is to be challenged in those cases in which AI is acting (mostly) autonomously, with the human 'user' (potentially someone with little technical understanding) perhaps merely switching on the AI system. This is on account of the fact that switching on an AI system does not qualify as a 'special feat in the field of technology' that would warrant a reward. By the same token, it would not seem appropriate to award the 'user' the privilege of a patent, i.e. a monopoly.

It has to be noted, however, that the number of cases in which this kind of autonomous AI has been deployed to date seems, at best, very limited. This also means that it will be possible for the vast majority of scenarios that can be expected to happen in the foreseeable future to be resolved according to the principle set out above, i.e. by regarding the AI as a 'tool' in the hands of a user who is also the inventor. It may well be the case, however, that this changes in the near future and that inventions made by autonomous AI systems will become the new normal. If we should find that it is impossible, for the reasons already given, to grant the 'user' of the AI system the status of an inventor, we would have to see whether the existing law provides for a viable, alternative way of identifying the inventor.

In these cases involving an autonomous AI system, the code-writer who created that system could potentially be regarded as the inventor. The code-writer of the AI system is the only natural person who has accomplished a special feat in the field of technology – a feat that is causally linked to the invention itself. Even if one were to disregard the fact that this solution would see the code-writer be

rewarded – perhaps for the second time – for writing the code of the AI (for which they might already have been granted a patent) rather than the actual invention itself, there would also be some significant practical issues: in many cases, the code-writer will not know about the outcomes generated by the autonomous AI system (unless the AI remains in their sphere of influence). Moreover, there would be little incentive for using AI in the first place if the outcomes generated by it do not accrue to the user but the code-writer. For these reasons, there is not much of a case for rewarding the code-writer of the AI by way of elevating them to the status of 'inventor' of the outcomes generated by the AI system.

The only other solution that could potentially make sense might be to grant the status of 'inventor' to the owner or proprietor of the autonomous AI system. The main problem here, however, is that the owner or proprietor of the AI may not necessarily be a natural person and that this is incompatible with the requirement that an inventor within the meaning of patent law must be a natural person. Legal persons can also act as owners and proprietors. Furthermore, unless the owner or proprietor (assuming they are a natural person) has personally 'switched on' the AI, their specific contribution to the invention will, in most cases, be even smaller than that of the 'user' of the AI. Lastly, whoever owns or possesses an AI system may depend on chance. For instance, it would seem arbitrary and potentially economically unsound to attribute an invention to a person simply on account of the fact that the machine that physically embodies an AI system has recently been transferred or handed over to them as collateral. It therefore follows from this that declaring the owner or proprietor of the AI the 'inventor' within the meaning of patent law is not a convincing approach.

We can summarise the results of this analysis as follows: under today's legal framework, it would be impossible for AI to 'invent' or be classified as an inventor within the meaning of patent law. It follows from the established principles that, if the AI has been used as a tool by a human user (who would necessarily be well-versed in technology), and if the user has made a significant contribution to the invention, and if the AI has not acted autonomously, the user of the AI will be regarded as the inventor. At present, this approach works for at least the vast majority of potential cases. After all, there are very few, if any, cases of autonomous AI today.

However, we may well see an increase in the number of cases involving autonomous AI generating inventions. In these cases, there is no viable justification for rewarding the 'user' of the AI by granting him the status of 'inventor'. To merely 'switch on' an AI system is not a 'special feat in the field of technology', as would be required under patent law for both granting inventor status and for a patent to be awarded. Nor does the existing law provide us with a suitable alternative solution that would allow for an 'inventor' to be identified in cases of autonomous AI. Most importantly, it would not make sense to attribute this status to either the code-writer of the AI or its owner or proprietor. This means that there is a need for new legislation/regulation for cases involving inventions generated by autonomous AI systems (for more details, go to C.).

III. Is it possible for AI to develop a personal intellectual creation within the meaning of copyright law?

The question as to whether AI can produce a personal intellectual creation arises where AI generates a result that could also be the result of a human process of creation – but only if the role of the AI system is not reduced to that of a tool used by the [human] originator (cf. I above).

German copyright law (similar to that of most other countries in continental Europe) is heavily influenced by the concept of the right to protection of personality. It is predi-

cated on the notion that copyright-protected works are created by a person who, akin to Spitzweg's 'The Poor Poet', lays open his innermost self in order to give others access to his intellectual world. Copyright does not distinguish between different types of work. Music, poetry, drawings, architecture and software are all protected for the same reason: to protect the creative human intellect.

For this reason, the originator is always the human 'creator'. They might decide to part with their copyrights and use rights, but will never cease to be originators. This is different from the Anglo-American legal systems, especially US law, which uses a 'work for hire' concept in certain situations (such as in employer-employee relationships) and for certain types of work (i.e. contributions to works created in cooperation with others). In these cases, the copyright lies not with the originator of the work, but with whoever commissioned the work (cf. 17 U.S. Code Section 101).

If we adopt a German legal mindset, whether or not copyright protection is available depends not on whether a creative act has been undertaken autonomously but whether or not whoever produced it is capable of originating a personal intellectual creation. AI does not fulfil that criterion. Despite its ability to act autonomously and despite the lack of foreseeability of the outcomes, it is still 'only' a machine that generates results on the basis of complex mathematical operations. This means that de lege lata, there are no analogies with copyright protection.

We might discuss, however, whether AI could potentially be able to generate outcomes that are protected by ancillary copyrights, e.g. photographs. It is important to note that – unlike with copyright-protected works – ancillary copyrights are granted not to protect the creative act, but the end result (e.g. the photograph). The holder of the ancillary copyright, however, is whoever delivered this result – in this case the AI, unless it is used as a mere tool by a human being. As the current state of legal debate is that artificial intelligence does not have legal capacity, it follows from this that it cannot actually hold ancillary copyrights.

Finally, we will discuss if the outcomes produced by artificial intelligence might at least benefit from the right to protection for databases pursuant to Section 87a Copyright Act (cf. Hetmank/Lauber-Rönsberg, GRUR 2018, p. 574, pp. 578 f.). There may be a few arguments in favour of this, but they do not bring us any further along. First of all, this is another case where there is no person subject to the law to whom the right would accrue (unless the AI is not used as a mere tool that does not do its compilations of its own 'accord'); second, protection for databases exists specifically for a given combination of individual elements, not the individual 'works' themselves.

Against this background, we can assume de lege lata that artificial intelligence – no matter how autonomous it may be – is incapable of creating works protected by copyright.

UK legislation has already evolved and arrived at a different situation. Sec. 9(3) UK Copyright Designs and Patent Act 1988, for instance, stipulates:

> "In the case of a literary, dramatic, musical or artistic work **which is computer-generated**, the author shall be taken to be the **person** by whom the **arrangements necessary** for the creation of the work are undertaken" (emphasis added).

This provision is remarkable on account of two aspects that allow it to resolve the question we are debating here (at least for the United Kingdom): first, a computer-generated work will benefit from copyright protection regardless of whether it is the product of human creativity – all that is required is that a human being made the arrangements necessary for its creation by a computer. Second, the provision also clarifies who the originator is: it is the person who has made the necessary arrangements/preparations for the creation of the work by a computer/AI system.

At this point, we will refrain from entering into a debate on whether this approach would be compatible with the German approach to copyright law, which is heavily driven by the concept of the right to protection of personality, or whether an entirely new type of ancillary copyright law would have to be introduced. We merely note that the UK legislation makes interesting food for thought.

IV. What are the options for protecting AI-generated outcomes of a work process under the existing legal framework?

Despite the discussion outlined above, there can be no doubt that the outcomes of AI have commercial value and that, therefore, companies that integrate AI in their process have an interest in protecting these outcomes so that they can make commercial use of them.

There are two legal instruments available for doing so under the existing legal framework, namely protection of

know-how and also ancillary protection under competition law, which, however, can only provide for a rudimentary level of protection restricted to some fields of application.

Following the end of the transposition period for the Know-How Directive (EU 2016/943), the criteria for know-how protection are now harmonised across the EU (at least by way of a mandatory interpretation of national law in conformity with the Directive). Information is protected under the Directive if it is secret ("in the sense that it is not, as a body or in the precise configuration and assembly of its components, generally known [...]") and "has commercial value" (on account of being secret). However, there is an additional requirement in that the information must also be subject to action suitable to ensure that it is kept secret. The case law to be handed down by the courts in the coming years will certainly provide for greater clarity as to what such suitable action may look like.

A scenario in which the output of AI systems are kept secret – perhaps even by way of 'confidentiality by design', i.e. from the very beginning on, is also possible. Any protection will, however, only be effective if said outcomes can be kept secret throughout their commercial use, for instance where they are used in internal manufacturing or business processes and do not show in the product that is placed on the market. By contrast, know-how protection is not available for outcomes that are to be placed on the market directly (e.g. a sculpture 'created' by AI).

These are cases where Section 4 No. 3 German Unfair Competition Act may apply – a legal instrument that is known for being fraught with uncertainty. The criteria that need to be met include not only that a third party must have created an avoidable false impression regarding the company that has produced products or services similar to those generated by the relevant AI, taken advantage of or diminished the level of appreciation for the products or services generated by the AI, or come into possession of expertise or documentation required for imitating the products or services by dishonest means, but also that the output of the artificial intelligence must be of a nature that is relevant for competition (cf. Köhler/Bornkamm/Feddersen, UWG, 36th ed. 2018, § 4 Rdnr. 3.24). This is the case if the specific form or some of the characteristics of the AI output are of a nature that allows for conclusions to be drawn about the company where it was generated, or if there are other specific characteristics.

Given that ancillary copyright protection is granted not for the way in which an outcome was produced but for the outcome itself and its characteristics, the question of whether a product or service can be granted this type of protection is no different in this case compared to that of outcomes produced in a 'conventional' way. The fact that 'something' has been generated by AI fades into the background here. What matters instead is that 'something's' ability to indicate which company it was produced by. This means that ancillary copyright protection will still be reduced to the role of a catch-all clause, even in the case of AI products that have been disclosed or placed on the market.

C: Options and recommendations for action

The current toolbox of commercial law allows for the protection of outcomes of AI processes in which the AI was not used as a mere tool as part of a inventive or creative process conducted by a human only in exceptional cases where these outcomes can benefit from know-how protection or ancillary copyright laws.

This leaves law-makers with the following options:

- **Maintain the status quo:** the argument in favour of this can be found in the 'reward' character of IPR law for inventions and creative output, which are predicated on the human intellect. It follows from the very nature of the system that output generated by AI machines is not protected. IPR law must not be amended in any way that would result in further restraints of competition.

- **Establish 'ancillary copyright protection for AI output' de lege ferenda:** having a separate legal instrument to protect AI output would allow for the specifics of 'AI creations' to be properly taken into account and for issues linked to copyright etc. to be dealt with in a centralised manner. It would also do justice to the fact that artificial intelligence is not human intelligence in an electronic form, but aliud. The new provision could be inspired by the UK provision outlined above.

- **Adjust the existing IPR regime de lege ferenda:** this approach would offer the advantage of not involving a separate new 'right' (namely one that would accrue to the 'generator' of the outcome) to be introduced alongside the existing ones that are specific to the type of outcome (e.g. patent law for technical inventions, copyright law

for intellectual works). Second, it would also make it easier to decide whether the AI system has 'generated an invention' or been used as tool. Again, the UK provision outlined above can serve as inspiration.

The first step, however, would have to be to see whether there is actually a need to protect AI output at this point in time.

One argument against changing the existing legal order is that the scenario above in which there is a lack of legal clarity on account of the fact that an AI system generates creative or inventive output autonomously and without considerable human contributions is not something that happens in real-life practice, or at least not other than in exceptional situations. In a standard scenario today, machine-generated instructions and other actions can be attributed to their operator, who uses the machines as a tool. The mere fact that some of the output generated by AI (e.g. deep-learning technology) cannot be predicted (at least not in detail) does not necessarily mean that this output cannot be attributed to the natural or legal person behind the AI system. At present, machines do not perform actions on their own, independent accord and do not create their own creative output.

However, it is worth exploring the introduction of legal provisions that would support the use of AI. For instance, a legal provision that would clarify that the rights to AI output accrue to the operator of the AI system would work as a kind of investment protection clause for AI systems.

Our recommendation for action to policy-makers would be that they should start preparing for the changes that might soon be happening at technological level and should begin discussions about the best possible system to allocate rights to AI output. This will allow them to be able to respond to technical progress. At present, it would seem compelling to attribute AI output to the system's operator, i.e. whoever has invested in the creation of the system (or their legal successor who might have purchased the system).

An alternative option – albeit one that should be carefully explored in advance – could be the creation of an 'ePerson', which would not only exist in liability law, but could also act as an originator or hold patents. At the present state of the art, however, this would seem to be an idea for the distant future, especially as the concept would only make sense economically and for society at large once AI systems have become capable and interested in making use of their protected rights and inventions. That said, there are many researchers who believe that AI technology is about to develop at an extremely fast pace over the coming years and decades, which is why legislators are well advised to prepare for even such scenarios as these.

Labour law

A: Factsheet

What is the issue at stake?

Artificial intelligence will transform the world of work and this will pose various challenges for labour law. The Federal Government studied this topic in detail and went on to present its AI Strategy on 15 November 2018. One of the objectives set out in the Strategy is for the development and use of AI in the world of work to be human-centred and designed around the labour force. Developing it in this way will allow skills and talents to be developed, enable self-determination, provide security and protect health. A further aim is for the IT systems that use and apply AI to be equipped with a high level of IT security. This also applies to operating systems. In addition, the principles developed at EU level which are to underpin the development of ethical guidelines for the use of AI will become particularly important as the discussion on AI continues.

What questions/challenges are there for Industrie 4.0?

Based on the working paper entitled Industrie 4.0 – How well the law is keeping pace, the following points emerge: What impact does AI have on

- vocational training and the safeguarding of jobs?
- health and safety at work?
- efforts to render working hours more flexible?
- the protection of employees' data?
- structures for giving instructions?
- HR decisions?
- the protection of personal rights?

B: Legal assessment

The primary purpose of labour law is to protect employees. In areas where companies are required to provide effective and appropriate protection for their employees, their scope for action is therefore limited. Co-determination and the work of elected bodies help ensure that companies respect these boundaries in practice. It is not yet clear whether the legal framework will have to be adapted to cater for AI. In any case, the use of AI-based systems must be built upon the principles of transparency, the traceability of AI-based recommendations and fairness, as well as the important ethical requirements mentioned below. It will be necessary to evaluate whether AI applications will be properly covered by the existing legal protection mechanisms.

C: Options and recommendations for action

The employment forecasts and scenarios made to date need to be scrutinised and strategies for designing the way in which we work, including an ongoing focus on the human factor, will have to be readjusted. The Federal Government has underlined the need for a human-centred approach to be taken as the requirements relating to skills, jobs, the organisation of work, and labour relations are changing. Against this background, it is important for Germany to have a National Further Training Strategy that also takes artificial intelligence into account, and for a joint approach to be taken by the social partners.

When it comes to professional development and the safeguarding of employment, AI will unleash a considerable need for further training to be provided. There is already some discussion on whether all employees should have a right to further training and whether they can be obligated to undergo such training. It now needs to be discussed whether co-determination rights and ways of safeguarding employment should be further developed and whether there ought to be a modern (digital) further training regime that goes beyond the Opportunities for Qualifications Act.

With regard to occupational health and safety, the question arises as to whether and in what ways the use of AI systems in particular is causing greater mental stress, and to where AI systems are helping to relieve the burden on the

employee by assuming hazardous or monotonous tasks. In this context, the existent right of the employee not to be contacted outside working hours needs greater clarification.

The flexibilisation of working hours in a digitised industry is not specific to AI. In this context, discussion is needed on whether existing leeway should be "exploited" within the framework of EU law. If staff are given greater autonomy over how they organise their working hours, employers will have to take greater responsibility for making sure that the designated number of hours is actually worked. There would need to be discussion about updating statutory safeguards or/and modifying the co-determination rights of works councils to protect workers against excessive demands.

A further approach could be to assess and, if necessary, further develop opportunities for staff to co-determine the introduction of AI applications in their work. It will be necessary to check whether opportunities for co-determination need to be adapted to the digital age. Concerning the question as to whether co-determination rights should be extended or modified and whether a specific Workers' Data Protection Act is needed, opinions between employees and employer representatives differ.

The use of AI applications also has implications for the protection of employees' data. In this context, introducing co-determination rights in data protection is the subject of controversial debate. It would be conceivable for minimum technical requirements to be set for personal data processing systems. For systems that autonomously decide which employee data are collected and processed for which purposes, requirements such as automatic deletion after the retention period or defined access rights could become particularly important for AI applications. The basis for this can also be found in Article 25 of the General Data Protection Regulation.

AI can alter structures for giving instructions. The question that needs to be addressed here is to what extent managerial functions may be performed by the computer, i.e. to what extent the computer can give instructions for work to humans. This is likely to primarily involve assigning instructions to each task of work to be fulfilled.

It would also be necessary to discuss the use of AI in personnel decisions. Sections 95 and 99 of the Worker Participation Act are relevant in this regard. It should be noted that personnel recruitment processes could, in future, be handled entirely using AI applications. The way in which AI systems are designed must be readily comprehensible and fair. AI systems can help make selection process less biased and render decision-making processes quicker and more transparent. On the other hand, however, they can also carry discrimination and bias in a way that seems neutral. Ultimately, processes like these which are supported by AI must in most cases end with a human making the decision – a decision for which he or she is then responsible. (This is a principle that applies to labour law in general.) Since the decision can can thus be attributed to (a) human being(s), there is no acute need to introduce new regulations. Where the protection of employees' data is concerned, Art. 9 GDPR, sections 22 and 26(3) Federal Data Protection Act (BDSG) and Art. 22 GDPR provide a level of protection. In individual cases, it will be necessary to examine whether AI systems are compatible with the personal rights of employees and their right to control their own data.

Finally, as already mentioned above, AI must be developed and applied in a broadly transparent, readily comprehensible and fair manner. AI must meet ethical requirements, such as the principles of human dignity, personal rights, non-discrimination and co-determination, as the EU expert group on this issue has rightly called for. This means that red lines must also be drawn in labour law – lines that must not be breached. This also places special demands on those who develop AI.

IT security and AI

(A review of current recommendations for action)

A: Factsheet

What is the issue at stake?

In order to guarantee IT security in B2B, it is fundamentally important that the ICT systems involved function reliably. This is also vital in order to create trust in the use of AI systems – not just in specific AI applications, but also in AI development and the underlying infrastructure used. As AI is more widely used and the amount of human-machine interaction increases, the demand is being made that the development and use of AI be governed by the highest security standards that are appropriate in each case. In addition to the danger that vulnerabilities can be attacked and exploited by hackers, the risk of damage can be increased due to the ability of corrupted AI systems to self-learn (cf. paper entitled Künstliche Intelligenz in Sicherheitsaspekten der Industrie 4.0 drawn up by the Plattform Industrie 4.0 Working Group on Security of Networked Systems).

In the following, Plattform Industrie 4.0 primarily sets out its views on the need to protect the AI applications themselves – this being of major importance as AI is used ever more frequently and widely in the industrial sector. Irrespective of what kind of final legal assessment is made, the prevailing political will to shape the future and the current discourse must be taken into account at this point.

What questions/challenges are there for Industrie 4.0?

- Does AI create a fundamentally new basis for evaluating the security of IT systems used as part of Industrie 4.0?
- Does the legislator need to take action to assign responsibilities for ensuring the integrity and confidentiality of AI systems?
- Are liability rules for IT manufacturers and providers of IT services/systems that are designed to eliminate deficiencies in data protection/IT security sufficiently regulated for AI?

B: Legal assessment

I. Does AI create a fundamentally new basis for evaluating the security of IT systems used as part of Industrie 4.0?

With respect to the underlying definition of AI (see chapter 1), the question of responsibility and who this is borne by is also important when it comes to security. The relevant comments on questions of liability law (cf. Chapter 4 Liability) can therefore be referred to again here. One of these comments states that "...there is no difference between scenarios using AI and other 'Industrie 4.0' scenarios". This means that the assessment of IT security for AI that was made in 2016 also continues to be valid in principle.

In its first recommendation for action in 2016, the Plattform Industrie 4.0 Working Group "Legal Framework" primarily focused on the need for practical measures to be taken based on the development of industry standards and certification to strengthen IT security. This approach is set down in the recently adopted Cybersecurity Act, which provides for the development of certification schemes. The European legislator has, quite rightly, not demanded that, at this level of abstraction, there be mandatory certification for categories that are still too abstract in terms of how these categories might be applied in practice.[10] Rather, there will need to be discussion along each of the different value chains as to whether providing certification for particular measures would help to increase the level of IT security in each case, and what measures it would be useful to provide this certification for. This discussion needs to be based around the particularities of each value chain and of the specific sector, as well as take account of the particular way in which AI is being applied. In the context of each of these value chains, it is then possible to look at whether or not a binding requirement for certification should be introduced in each specific case, as is provided for in principle in the Cybersecurity Act. However, it must always be borne in mind that regulatory adjustments also always encroach on the business and contractual autonomy of companies.

10 The Regulation provides for the development of three different assurance levels (basic, substantial, high), each of which will be linked to different protection requirements that have to be fulfilled in order to be issued with the particular certification.

When it comes to security protection requirements for Industrie 4.0 and AI, the greatest risk to operational safety does not stem from the AI system itself but from adding further components (such as production robots) that depend on an increasing degree of automation or autonomy. The rules designed to protect the security of plant and products (in order to protect people and surroundings) which are formulated at legislative level must also be complied with in the digitalisation of processes and their networking. The risk of persons gaining authorised access to such plant and products from outside must therefore be given the attention it is due. There is no clear opinion or position that emerges from present discourse on this matter. However, both IT suppliers and their industrial clients are endeavouring to conduct a targeted discussion with a view to strengthening IT security. This has, for example, led to the development of the Tech Accord and the Charter of Trust on Cybersecurity.

Companies that have signed the Tech Accord focus primarily on developing cooperation with each other and with institutions, designed to serve their own interests: "We will work with each other and will establish formal and informal partnerships with industry, civil society, and security researchers, across proprietary and open source technologies to improve technical collaboration, coordinated vulnerability disclosure, and threat sharing, as well as to minimize the levels of malicious code being introduced into cyberspace."[11]

A similar approach is also pursued by the members of the Charter of Trust. This, however, also refers to a need for regulatory support: "Companies and – if necessary – governments must establish mandatory and independent third-party certifications (based on future-proof definitions and especially where life and limb are in danger) for critical infrastructures and IoT solutions."[12]

From the point of view of the Charter of Trust on Cybersecurity, the possibility of adapting the European Machinery Directive 2006/42/EC should at least be discussed as the rules on critical infrastructures also continue to be developed (see remarks below on orientation toward the public interest).

The comments on the evaluation category "Orientation toward the public interest" presented in the working paper of 2016 also do not suggest that the use of AI creates any clear need for a differing recommendation where AI is used.

When it comes to the orientation toward the public interest, the regulatory focus is on ensuring that the internet and IT systems are able to function properly in areas that affect public interests and are therefore classified as critical infrastructures.[13] Since the use of AI systems is no means restricted to critical contexts only, and since the AI system itself does not necessarily make the context in which it is used critical as criticality is usually only produced through the use of different systems[14], "orientation toward the public interest" can also be ruled out as a criterion why the use of AI systems in Industrie 4.0 scenarios should per se be assigned to the regulatory area of critical infrastructures.

II. Does the legislator need to take action to assign responsibilities for ensuring the integrity and confidentiality of AI systems?

Since "orientation towards the public interest" is not used as a regulatory basis, for the time being, the following remains true for AI as well: "...it is ultimately up to user demand – or in interconnected Industrie 4.0 structures, to the interconnected companies acting as a consortium – to decide whether they individually will implement a higher security level than the minimum standards."[15] This also means that measures to increase IT security should currently also be laid out in contractual arrangements.

Whether and to what extent technical standards and certification can be demanded without an additional contractual basis remains to be seen, as an EU-wide certification framework for the cybersecurity of products, processes and services still has to be developed under what will be the European

11 Tech Accord, www.cybertechaccord.org

12 Charter of Trust on Cybersecurity, www.charter-of-trust.com

13 Cf. Federal Ministry for Economic Affairs and Energy/Plattform Industrie 4.0 (2016): Industrie 4.0 – How well the law is keeping pace p. 9

14 An example might be the interplay between an AI system and autonomous production environments, the combination of which could create a security risk. Therefore, the interaction between the different systems would have to be considered in sum and not just the effect of a single component be taken into account.

15 Federal Ministry for Economic Affairs and Energy/Plattform Industrie 4.0 (2016): Industrie 4.0 – How well the law is keeping pace, p. 10

Cybersecurity Act. However, the industry could think about developing specific verification and certification schemes for IT and AI components itself in order to adequately address the increasingly complex questions of who should bear responsibility.

The use of classic IT systems does, however, differ from that of AI systems in one specific area. Because of the specific way in which AI systems work, these depend to a especially high degree on the availability of high quality data. This means that ensuring data integrity is of crucial importance in order to protect AI systems. To make sure that the data is of a high quality and that AI makes integral decisions, it must be ensured that the basic data used is not manipulated in any way. This creates a need for new security requirements to be established which must be taken into account when developing and applying the relevant standards. These requirements should also guide what protection rules look like, including 'security by design'. Since it is questionable whether the legislator can stipulate a priori how collaboration between different entities in an AI context is to take place – for example through the provision of data – in order to assign responsibility clearly, the flexibility that can be created through companies concluding relevant contracts should be prioritised in this context, too. Market developments in this direction should be valued and given the necessary support.

III. Are liability rules for IT manufacturers and providers of IT services/systems that are designed to eliminate deficiencies in data protection/IT security sufficiently regulated?

In addition to the existing links between IT security and liability (cf. Chapter 4 "Liability"), the way in which AI systems are actually developing in the Industrie 4.0 environment should also be considered.

In contrast to many traditional IT components, it can be assumed that many AI applications or parts of these are provided not just as individual products, but also via digital platforms. Again, this creates a need for contractual arrangements to be established for companies that work together on the market. These arrangements ought to be based on cross-sectoral standards that first need to be developed and then implemented. However, if the way in which the market is designed does not go on to provide sufficient stimulus for the demand side to cover its IT security needs in a proper manner, for example based on contracts, it may become necessary to adjust the way in which responsibility for ensuring the security of IT systems that use AI is assigned.

In order to encourage practical action to be taken, as advocated in principle by WG "Legal Framework", extensive attempts should however be made to ensure that each industry can achieve an appropriate balance of responsibility by following examples of practice. After all, as already explained at the beginning, we are able to rely on the fact that both providers and the users of AI systems alike will endeavour to protect themselves. As such considerations are made – a process which WG 4 "Legal Framework" will also participate in intensively – it is also important that newer models are discussed, including models that cover legal and economic aspects together.

C: Recommendations for action

In summary, it can be said that i) the recommendations for action published in 2016 also apply in principle with respect to the growing use of AI systems in Industrie 4.0 environments, ii) measures to ensure IT security for the use of AI must first and foremost be motivated by a company's desire to protect itself, and iii) the provision of such security should be enforced via contacts.

In addition, recent debate – for example on a possible revision of the Machinery Directive – has highlighted the need for there to be more intense discussion on how the IT systems of facilities and products can be protected against unauthorised access.

At least until such time that the Machinery Directive might be revised and new rules transposed into national law, it would, however, seem more useful to focus in the B2B field on strengthening market mechanisms to ensure the level of security provided continues to rise, especially in view of the increasing complexity of technical IT systems in general and AI systems in particular.

The primary focus of discourse therefore needs to be on how the relevant standards should be adapted to the requirements of AI systems and IT systems incorporating AI components or, where such standards are lacking, on how these can be developed. Any emerging discussion about legislation that would impact value chains – such as a revision of the Machinery Directive – must in no way be self-referential. Rather, it must be guided by the needs of individual market players for protection and the actual possibilities for them to implement such legislation in practice, and should particularly be prioritised if it starts to become apparent that market mechanisms are failing.

AUTHORS

Dr Dennis Amschewitz, lawyer (Robert Bosch GmbH) | Prof. Dr Gesmann-Nuissl (Technische Universität Chemnitz) | Dr Philipp Haas, lawyer (Robert Bosch GmbH) | Nils Hullen, lawyer (IBM Deutschland GmbH) | Dr Ulrich Keil, lawyer (Schaeffler AG) | Dr Thomas Klebe, lawyer (Hugo Sinzheimer Institut für Arbeitsrecht) | Thomas Kriesel, lawyer (BITKOM e.V.) | Thomas Schauf (Deutsche Telekom AG) | Dr Johannes Schipp, lawyer (T S C Fachanwälte für Arbeitsrecht) | Marc Wirwas, lawyer (HARTING Stiftung & Co. KG)

This publication is a result of the work of Working Group “Legal Framework” of Plattform Industrie 4.0.

www.plattform-i40.de