Artificial Intelligence

What are the key features of artificial intelligence (AI)?
How will AI affect the insurance value chain
and which issues arise as a result?

8th AIDA Europe Conference,
„Landfall of the Tech Storm“,
Lisbon, October 3, 2019

AIDA EUROPE

CALOUSTE GULBENKIAN FOUNDATION

大成 DENTONS

Contents

Introduction and Moderation (Kai Goretzky)

The Silicone Valley mantra is that artificial intelligence will always be faster, more thorough and more accurate. Undoubtedly, artificial intelligence transfers Isaac Newton's formulas into today's digital space: work is force times distance, and performance is work multiplied by time. Insurers have always been excited about processes that are faster, cheaper and produce higher quality. 87 % of global insurance managers currently consider AI to be the most important topic for the future, but only 44 % have introduced AI so far. Is this because, once again, the law is lagging behind technology - and does artificial intelligence challenge previously established legal concepts such as legal capacity, declarations of intent, attribution of knowledge, liability concepts, data protection and regulation? These questions about artificial intelligence, and the role of both the private sector and the state, differ fundamentally between the USA, Asia and here in Europe.

While the vast majority of today's technology can only fulfil clearly defined tasks within a narrow scope (so-called weak AI), deep learning will soon enable AI to achieve intelligence that is at least very close to or even exceeds the depth and breadth of the intelligence of a human being (so-called strong AI). The defeat of the chess world champion Kasparov by the "Deep Blue" system is a pertinent example, and further evidence can be found in the shape of high-speed algorithmic trading on stock exchanges and the encroachment of Bitcoin on traditional payment systems.

Since 2017, the European legislator has been rolling out several initiatives in response to these developments, such as the "Directive on Digital Contracts", the "Deal for the Consumer" and, just this summer, guidelines on the use of artificial intelligence. These guidelines focus on the protection of people from machines and lay down basic principles on how to achieve this objective. However, they do not contain any reliable or granular concepts for the use and control of AI.

This panel is a forum to discuss selected AI issues affecting insurance from a claims, governance, liability and data protection perspective. The use case discussed during this panel will be the design of an accident and liability insurance policy compensating for losses caused by a highly or fully automated driving system. Let us assume for the purposes of this panel that this driving system is either completely autonomous or recognizes its limits in good time and passes control back to the human driver independently, for example when approaching the end of a traffic jam.

This use case shows that, in the future, AI will be both a subject matter of insurance policies as well as a tool for the design of coverage. Most likely, artificial intelligence will extend to the entire insurance value chain: from product development, to marketing, underwriting, claims and payment.

When designing future coverage, a traditional manual approach will compete with an AI-driven approach, where insurers will convert unstructured data into structured data more rapidly and use specific data for product design, risk underwriting tools and prompt exchanges with external service providers and customers via chatbots and robo advice. This is the area where InsurTechs play an important role today.

Discussion Points on Claims Handling (Ralph Echensperger)

- **Which technologies does Zurich use to foster its claims process?**

  We use different kinds of technologies for different steps along the claims handling process:

  To simplify and speed up the claims process we use robotics, analytics, machine learning and predictive modelling: for instance, incoming documents are automatically processed, key words are captured and documents are classified, indexed and triaged to the respective claim in our system. Furthermore, AI technology generally helps us automate or partially automate different process steps (automated payments, verification of coverage, etc.).

  Cognitive computing helps our employees derive liability in complex cases, evaluate possible case strategies and detect recovery. AI technologies also evaluate car claims from telematic data and derive the best settlement option such as reparation, replacement or cash settlement. To prevent or minimize fraud, we use machine learning to detect fraudulent patterns.

  In addition, we also provide our customers with enhanced customer service through machine learning: in a property claim, we offer our customers different settlement options such as cash, a replacement but also suggestions for other products based on their claim history. Furthermore, we use cognitive computing and machine learning to generate risk advice for our customers or direct them to the best possible provider in the case of a claim.

- **How do you assess customer behavior and AI acceptance?**

  In view of the results of our customer surveys (before and after a claim), we know that, generally speaking, our customers want either a fast and digital claims process or personal and customized customer service - depending on the nature of the claim.

  Regarding AI acceptance, it has been our experience that customers appreciate the AI-based service they receive as they benefit in various areas: they receive suggestions for replacements after a property claim based on their claim history; they receive advice regarding risk exposure, and so on. Therefore, we see a high level of acceptance among our customers. However, we must take into consideration that AI might be disadvantageous for customers if it detects fraud or takes a decision that does not match customer expectations. Situations like that may reduce acceptance even though the risk community and therefore, our customers, benefit(s) overall.

- **Are there differences depending on countries/customers/cultures?**

  Yes, there are many differences for us to consider when creating a digital customer journey. For instance, in Switzerland we possess a high level of knowledge regarding new technology thanks to good education, and we export

digital services all over the world. Paradoxically, a Swiss customer is rather reserved when it comes to accepting AI technology – because of our mentality. Therefore, Swiss companies aren't as digitized as our knowledge would allow. However, if a Swiss customer has a good experience with AI technology, they are willing to try new things. Younger customers obviously approve more strongly of AI technology as they grew up in a digital age. For us as an insurance company, those differences mean that we have to offer a wide range of services in order to meet all customer requirements.

- **What is the relationship with InsurTechs like today (disruption vs. cooperation)?**

  We have to take advantage of the inventions and ideas of InsurTechs from around the world. We cooperate with them if they can offer a service our customers can benefit from. However, as innovation is also a competitive factor, we want to create and implement as much as possible on our own. This keeps us independent and closer to our customers.

- **In what areas might AI be beneficial and deployed in claims over the next 5 years and are there limits from an operational perspective?**

  At the moment, the sky is the limit: we see a lot of potential in areas such as risk management, fraud detection, chatbots, personalized customer service based on preferences and past interactions and fast-track claims. However, we have to evaluate which processes or services are beneficial for the customer as well as the industry and in which areas acceptance is insufficient - for instance services mainly based of personal interaction and trust. In conclusion, nobody knows where we are heading: AI may be obsolete in a few years, but equally, everybody may be using it in their everyday life.

Discussion Points on Governance and Risk Management (Johannes Kinold)

- **From a governance perspective: what will change if AI becomes part of the insurance value chain?**

  Regarding the principles laid out by SII and national regulation – not that much, it just adds a new perspective/aspect that has to be observed.

  The first principle would be that the responsibility for all processes remains with the management/board of executives, incl. those processes operated by AI and machines and incl. those machines that are hosted by external service providers via outsourcing agreements. Management has to prepare and install adequate precautions, such as oversight rules and authority directives.

  The second principle is that all processes and activities as well as the respective roles and operative responsibilities have to be defined and documented comprehensively. Where activities and operative responsibilities, e.g. for the verification of claims, are delegated to a machine/AI, the underlying rules and algorithms must also meet the requirements as defined and documented. Black-boxes will not be an option here. So far, it can only be assumed, but I am very confident the regulators will not accept any Blackbox-excuses, considering the high level of diligence they are already postulating for comparably simple individual data processing applications such as Excel sheets. Additionally - and intrinsically -, the organization needs a profound understanding of the models for both analytic purposes and the continuous controlling of machine-based learning.

  A third principle would be the separation of functions, meaning that the creation of risks on the one side and the monitoring and evaluation of these risks on the other, are to be delegated to separate units/responsibilities. Leaving the responsibility of monitoring and evaluation to the algorithm as well, e.g. to further increase efficiency gains, does not comply with the principle of separated functions.

- **What are the options when setting up an effective and efficient control system for AI?**

  That is tricky, because it is a greenfield area. When we were only talking about process automation, the answer was relatively easy: test management. You would implement control mechanisms such as plausibility and integrity checks into the automated routines and prior to launching it, you would run and document a sufficient number of test cases to ensure the operating effectiveness of the control mechanisms. And then, on every new release you would just run the same tests again, regardless of whether the specific controls to be tested had been changed during the release or not.

But now, artificial intelligence enabled for machine-based learning will just launch new releases on its own, again and again, potentially without human supervisors noticing right away or even at all.

The only way forward, in my opinion, would be to make the algorithm subject to the same control environment any human resource would have to endure.

Preventive controls here could be based on the financial volume of the transaction. Regarding underwriting, the algorithm could only accept and issue policies until an annual premium of 500€, e.g., as soon as the required premium is above this limit, implying that the underlying risk is higher, the transaction would be forwarded to another instance/authority for additional checking and approval. The same method would of course work in claims processing, where a financial threshold or a specific type of claim could lead to the requested approval by a higher authority. These kinds of controls would of course reduce efficiency gains, since human resources would still play a part in the operative processes.

Additionally, one could enact detective controls, e.g. quality assurance frameworks where, according to a defined schedule, random transactions would be picked and checked for adherence to given rules and targets, e.g. underwriting guidelines in the case of underwriting transactions.

Currently, another future option would be the delegation of the aforementioned control mechanisms to AI. In this scenario, it is important to use a different algorithm/machine for the control mechanism than for the operative conduct of the processes in order to comply with the principle of the separation of functions.

- **Could outsourcing help with providing separate functions in that matter? What are the opportunities and constraints when outsourcing to Insurtechs or BigTechs?**

  Outsourcing either the operative or the control instance alone would not be sufficient to comply with the separation of functions principle, since there always needs to be a dedicated in-house role responsible for the outsourced activity. Compliance can only be attained where this role is separated from the in-house activity, so that again, the principle needs to be adhered to within the organization. The same is true for management responsibility and comprehensive documentation of processes, roles and responsibilities: all principles need to be met no matter whether the activity is carried out in-house or via outsourcing.

  There are, of course, opportunities that make outsourcing with regard to technical innovation matters very appealing to insurance companies. The first of these is the possibility of accelerated development and implementation. Traditionally, insurance companies are rather slow at developing and continuously optimizing their IT infrastructure and prefer to spend time administrating multiple inventory management systems. Now, by cooperating with InsurTechs, they have the opportunity to gain rapid access to new technology and innovative customer-interfaces.

Secondly, there are advantages for both sides, i.e. insurance companies and InsurTechs, which is why we generally see cooperation in this field. The two are not actually competitors. InsurTechs do not generally have the trusted brand, financial capital or regulatory experience required to take on a role as risk carrier. Insurance companies are, traditionally, not skilled at developing technological solutions and applications in a fast and agile manner, even if many of them are increasingly eager to appear that way. In my experience, InsurTechs often confine themselves to the role of a licensed broker in order to enter the insurance value chain. In my opinion, this positioning is a very smart and distinct move towards cooperation, and thankfully, insurance companies accept it.

On the other end of the value chain, at least in our line of business (legal expenses insurance), we are increasingly dealing with what we call LegalTechs, and the approach here is not so cooperative. There are case examples where we cooperated with LegalTechs by bringing them new business while simultaneously enhancing our own claims processes and improving our results. One of these examples is “flight right”. However, there are also case examples where we were in competition with or even in direction opposition to LegalTechs, either because their offer was a substitute for our insurance product or because they sought out and approached our customers in order to handle their legal conflicts for them, and thereby increased our claims numbers.

There is also one constraint or downside with regard to outsourcing. Traditionally, one of the key risks of outsourcing activities to external service providers was the loss of know-how and expertise. This risk is not that relevant while outsourcing to InsurTechs, since insurance companies never had the innovative and agile technology know-how in the first place. The risk here, however, is that outsourcing acts as a barrier to building up the required expertise within the insurance organization in the first place.

Luckily, effective mitigation measures for this risk are self-evident. If management intrinsically accepts its overarching responsibility and effectively monitors the extensive creation and documentation of all processes, roles and responsibilities both in-house and externally, the information and know-how is at least available for someone to absorb it.

- **Do regulators actually provide useful requirements and tools to deal with new developments? Is further AI regulation to be expected?**

  I am afraid so. And I regret that this is the case because I actually think that the existing principles and regulatory requirements are sufficient even when AI is added to the insurance value chain.

  For example, the German regulator did the following: they created an extensive report on Big Data and Artificial Intelligence with the support of external consultants and then published it together with a questionnaire on expected future developments and potential requirements for new or additional regulation. We participated in said study, but have not received any new information on the matter since then, so I only know our answers and the answers of the German actuary

association, and we both agree that existing instruments are designed effectively enough to embrace AI. The key issue in this regard is whether any regulatory initiative would actually result in new regulation – which would be fine in my view, even though I do not know what that would encompass – or in additional regulation, which – like I said – is not necessary.

The problem with additional regulation is that it would simply be a result of the original regulation not being enforced effectively. Naturally, I see the challenges of enforcing a principle-based regulatory regime rigorously throughout the industry. It should not lead to additional requirements, however, but to measures that are meant to clarify and explain the diligent implementation of the underlying principles across the entire value chain and in all other relevant aspects of insurers' business activities and infrastructures.

To sum up, I believe that new AI-related regulation is not necessary but nonetheless to be expected.

Discussion Points on Product Liability and Insurance Cover (Jens Gal)

- **When considering Level 3 and Level 4 driving, does AI require changes to the concepts as set out in the Vienna Convention on Road Traffic?**

  While it is claimed that the Vienna Convention is based on the assumption that the driver is always fully in control and responsible for the actions of the vehicle in traffic, I don't think that such understanding could not be harmonised with third and fourth level driving. The steps taken by the Economic and Social Council of the UN in 2014 show that the convention only assumes there to be a driver, who, however, may be aided by AI. Therefore, when we are speaking of conditional automation and high automation, this does not run counter to the concept of a necessary driver. It is only when we enter the stage of full automation, level 5 driving, that the concept would appear to be outdated.

- **Do the rules of the Products Liability Directive suffice to address the risks of autonomous cars, i.e. unbundled software, b2b claims?**

  Well, firstly, we should ascertain that there is a universal understanding that AI is software and that software is a product in the sense of the Product Liability Directive. While this is now common understanding in Germany, this is not necessarily so in other jurisdictions. This should probably be addressed by the Commission. Several questions still remain that may not be optimally addressed by the Product Liability Directive. If we take the example of unbundled software, it will become increasingly difficult to say what the complete product actually is. If a user combines two programs with each other and an error occurs, this could be the fault of one or both programs, or be due to a faulty interaction. When we are talking about car accidents, this would not initially create a problem, since the owner of the car would be strictly liable towards third parties. Since they are presumably at fault in relation to the rearrangement of the software, it will not be unduly burdensome for them to prove what software was defective or created an interaction problem that should have been foreseen. In other instances, this seems problematic. If producers could always seek exemption from liability due to unbundling, this would further a system in which competition is severely hindered. Concerning b2b claims, I am still on the fence. The directive is an instrument for consumer protection and its implementation in the b2b environment does not seem entirely appropriate.

- **Is there need for a strict liability concept to cover AI and (mandatory) insurance coverage? What may the impact be on innovation and the product design of insurers?**

  Emphatically, no there is not. To impose a system of strict liability with regard to AI would be the beginning of the end for fault-based tort law. Considering the number of products and services that already involve some sort of AI, such a system of strict liability would turn lawsuits on their head, since the injured party would supposedly initially seek damages from the producer, who would then

seek redress from the user. This would turn producers into professional respondents and claimants and would require an enormous amount of resources. It would also tend to hinder innovation, since the extent of insurance coverage needed even during the testing phase but also during the initial rollout would be incredibly cost intensive. This would particularly deter start-ups from producing AI and would further concentrate innovation in a small number of affluent multinational concerns. I would therefore say that we, as societies, should stop seeing strict liability and mandatory insurance as the passe-partout answer to every challenge. AI should simply stand for artificial intelligence, rather than being all-inclusive in the sense that every risk a citizen is subject to is shared between all citizens and that mandatory insurance covers every risk of every citizen.

- **Which other liability concepts are conceivable and what might an ongoing duty to test and monitor AI look like?**

  One concept would, of course, be the duty of integrated monitoring, testing and updating. In view of the increased level of connectedness such monitoring and updating is already possible and taking place. One could even ask if such a duty of integrated monitoring is not already inherent in current tort law. At least German tort law imposes so-called *Verkehrssicherungspflichten*, i.e. a duty to implement safety precautions, on every person creating a source of danger. The type of precaution that must be taken depends on the amount of danger and the level of safety expected by the public as set against the costs and efforts incurred by the originator of the source of danger. Depending on the type of AI, courts might come to the conclusion that such ongoing integrated monitoring is already owed. Such monitoring will, in most cases, be performed by another AI, so such a monitoring duty would not be the philosopher's stone. It does, however, highlight that the current fault-based liability system, in conjunction with product liability provisions, is able to evolve organically to meet new challenges posed by the advent of AI.

- **In an increasingly connected world, parties designing or offering AI are more likely to become subject to joint liability at first instance: how might this impact insurers' partition agreements?**

  Since I am an academic and not an industry insider, I cannot speak as to what changes have already occurred in the framing of partition agreements. I suppose that they are still, for the most part, based on the assumption that first instance proceedings only involve one insurer who will lead the proceedings and if necessary, settle the claim. Where the insurers of several AI-producers, who are jointly sued for damages, are now each responsible for leading the court proceedings concerning their policyholder, the situation would alter significantly. In order to manage court proceedings more efficiently and reduce their costs, it would be advantageous to have them led by one insurer.

- **When does AI produce a "wrong" result and how can this be proven?**

  If you want the short answer: when the judge says so. Since the categories of right and wrong are highly subjective, even metaphysical, I do not believe that we will come up with a right or wrong answer to the question. I suppose we will

have to make do with our time-honored categories of negligence. We will have to ask if the AI took an action that would not have been taken by someone who exercised the appropriate or ethical standard of care expected in the given situation. The problem will be who the person of reference is, the long-standing *bonus pater familias,* or whatever the person is called, or the *bonus intelligentia artificialis*. No matter the person of reference or the level of diligence required, one of the problems becomes that the AI will often not be able to explain why a certain action was taken or not taken. A car driver may provide testimony as to the factors which moved them to change lanes. In order to submit such proof of correct action (or at least reduced negligence) for AI, one would have to inspect the protocol leading up to the accident and assess it in light of the algorithm. From this, one would have to distil whether the AI implemented it correctly (if it did not, this would still not mean that it produced a wrong result) and what factors it took into consideration and weighed up when deciding its action. Which brings us back to the ethical question: is it right for AI to change lanes if it would otherwise hit a pedestrian, even though this might endanger drivers on the other lane and the driver of the car itself? If such a lane change works out without an accident, we would consider it the right decision; it is only when an accident occurs that we would ponder whether it was wrong. AI and its programmers do not have the luxury of making purely spontaneous decisions. While certain decision parameters will be set by legislators, others will only develop over time through case law.

Discussion Points on Data Protection (Helena Tapp Barroso)

- **How does AI affect the concept of data ownership (controllers, joint-controllers and processors)? Is co-ownership an appropriate concept? Is it clear who will respond to customers' data protection rights?**

  The concept of a data controller and its interaction with that of a data processor play a crucial role in the application of the GDPR. The legal obligations that apply to controllers are different from those that apply to processors. 'Data controller' means the person/entity who/which, alone or jointly with others, determines the purposes and means of processing of personal data. 'Processor' means the person/entity who/which processes personal data on behalf of the controller. That is to say, that data processing is carried out by both controllers and processors and the definition of processor is dependent upon that of controller; you need a controller before you can have a processor.

  Knowing who the controller is also plays a crucial role in terms of how data subjects can exercise their rights in practice. When it comes to joint controllers, the GDPR requires that by means of an arrangement between them, joint controllers determine, in a transparent manner, their respective responsibilities with regard to GDPR compliance, in particular the exercising of rights by data subjects and controllers' duty to provide information to such data subjects. In this arrangement, a contact point for data subjects may be designated and arrangements shall reflect the role of each of the joint controllers vis-à-vis the data subjects. The joint controllers shall provide for the essence of the arrangement to be made available to the data subject. Irrespective of the terms of the arrangement, the data subject may exercise the rights provided to data subjects (e.g. the right to request from the controller access to and rectification or erasure of personal data) in respect of and against each of the controllers.

  The controller shall only use processors providing sufficient guarantees to implement appropriate technical and organizational measures in such a manner that processing will meet the requirements of the GDPR and ensure the protection of the data subject's rights, and the processor shall not engage subsequent processors without prior specific or general written authorization of the controller. Processing by a processor shall be governed by a contract, binding the processor with regard to the controller and setting out the subject matter and duration of the processing, the nature and purpose of the processing, the types of personal data and categories of data subjects and the obligations and rights of the controller. The GDPR additionally sets out a number of stipulations that must be contained within the contract, including that the processor processes personal data only after receiving documented instructions from the controller.

- **Which preparatory actions must insurers and their service providers take when mining data to ensure that they properly use personal data for AI purposes? How will the insurer match its AI usage with the GDPR principles of purpose limitation, data minimization and accountability?**

  Clear identification of the legitimate grounds for processing is a "must". This means that a clear delimitation of processing purposes is also required since the assessment of the processing grounds is always made in view of the processing purposes.

  The purpose limitation principle means that the reason for processing personal data must be clearly established and indicated when the data is collected. This is essential if the data subject is to exercise control over the use of his/her information. The purpose of the processing also needs to be fully explained to the data subject if he or she is to be able to make an informed choice about whether or not to consent. Yet the development and application of artificial intelligence often requires many different types of personal data, which was, in some cases, collected for other purposes. This recycling of information may be useful and may provide a more accurate analysis than was previously technically feasible, but can also be in contravention of the purpose limitation principle. In cases where previously-retrieved personal data is to be reused, the controller must consider whether the new purpose is compatible with the original one. If this is not the case, new consent is required or the basis for processing must be changed.

  Processing that is necessary for legitimate interests (pursued by the controller or by a third party) will provide legitimate grounds for processing except "where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, in particular where the data subject is a child". A balancing test is, therefore, required. It is critical in the context of AI, which must strike the difficult balance between the rights of data subjects and the legitimate interests of the data controller; there being no clear guidelines on how to strike such a balance. Setting up the legal basis and purposes of personal data processing remains one of the most important features to take into account when dealing with AI systems and related machine learning features.

  Carrying out Data Protection Impact Assessment (DPIA) will also be a preparatory measure. Under Article 35(3) of the GDPR, the controller is required to carry out an assessment of the impact of the envisaged processing operations on the protection of personal data, inter alia, in case of "a systematic and extensive evaluation of personal aspects relating to natural persons which is based on automated processing, including profiling, and on which decisions are based that produce legal effects concerning the natural person or similarly significantly affect the natural person". Accordingly, most AI systems would require a DPIA before carrying out any personal data processing. This will require a detailed assessment of AI systems from a data protection perspective,

also with regard to the relevant security measures which are applied. Supervisory authorities have issued lists regarding processing operations subject to the requirements of a DPIA. In Portugal, this includes the processing of specific categories of data, highly personal data or data on criminal convictions and offences or related security measures, the use of new or innovative technology, as well as processing that involves or consists of a substantial amount of profiling.

With respect to data minimization, developers should start carrying out research on possible solutions that use less training data. They should also conduct further research on anonymization techniques and solutions that explain how systems process data and how they reach their conclusions.

- **How do the GDPR principles "privacy by design" and "privacy by default" affect the design and control of AI?**

  One key change introduced by the GDPR is the obligation to integrate privacy into systems and operations when processing personal data.

  The concept of 'Privacy by Design' was developed by Ann Cavoukian and indicates a new "philosophy" and approach to embedding privacy into the design of information technology, network infrastructure, and business practices rather than adding it onto code after it is written. It is about integrating privacy in the full lifecycle of systems, operations and products.

  The 'Data Protection by Design and by Default' concepts in the GDPR text are explained as follows: "…the controller shall, both at the time of the determination of the means for processing and at the time of the processing itself, implement appropriate technical and organizational measures…". It is therefore about companies integrating privacy into their systems and operations as well as the end products and services they deliver.

- **Are there cases where AI may turn anonymous or pseudonymous customers into identifiable ones?**

  One must be aware of quasi-identifiers, which are not unique identifiers (e.g. gender, postcode, profession, languages spoken), as they can re-identify people when combined. An interesting report published by Latanya Sweeney in 2000 showed that combining just three quasi-identifying pieces of information enabled 87% of the U.S. population to be identified.

  Quasi-identifiers have been the basis of several attacks on released data (in Netflix and AOL logs cases, the de-identified data could be re-identified by merging the quasi-identifiers with the available information from other sources). In such reports, the practice of de-identifying data and of ad hoc generalization are shown to be insufficient to render data anonymous *because combinations of attributes often combine uniquely to re-identify individuals*.

- **Do GDPR requirements on automated decisions provide for a balanced approach between innovation, efficiency and data protection?**

  The GDPR applies to artificial intelligence whilst it is being developed using personal data as well as when it is used to examine or reach decisions about individuals.

  GDPR provisions that are squarely aimed at machine learning state that "the data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her".

  Article 22 is a general restriction on automated decision making and profiling. It only applies when a decision is based solely on automated processing – including profiling – which produces legal effects or similarly significantly affects the data subject.

  The restriction shall not apply if the decision: (a) is necessary for entering into, or the performance of, a contract between the data subject and a data controller; (b) is authorized by law (which also lays down suitable measures to safeguard the data subject's rights and freedoms and legitimate interests); or (c) is based on the data subject's explicit consent .

  In the (admitted) cases of contract and consent, the data controller shall implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, at least the right to human intervention on the part of the controller, to express their point of view and to contest the decision.

  Under this "human-based approach", even if the processing is carried out by automated means, a data subject is granted the right to object to the possible decisions derived from it. It may, however, be difficult to apply such a right from a practical perspective, since AI systems are designed not to allow human interventions which may replace automated decision making.

  Also noteworthy are Articles 13 and 15 which repeatedly state that data subjects have a right to "meaningful information about the logic involved" and to "the significance and the envisaged consequences" of automated decision-making. One of the most frequently debated topics in the context of GDPR and AI discussions relates to such "right to explanation." Despite common misinterpretations, the GDPR does not actually refer to or establish a right to explanation that extends to the "how" and the "why" of an automated individual decision.

  "Meaningful information about the logic involved" in relation to Article 22 of the GDPR should be understood as information around the algorithmic method used rather than an explanation about the rationale of an automated decision. For example, if an insurance application is refused, Article 22 may require the controller to provide information about the input data related to the individual and the general parameters set in the algorithm that enabled the automated

decision. However, Article 22 would not require an explanation around the source code, or how and why that specific decision was made.

Conclusions (Kai Goretzky)

There is nothing to be gained from demonizing artificial intelligence by putting forward fear-inducing scenarios. Instead, we should appropriate and proportional steps to ensure that the insurance industry and its customers reap the benefits of artificial intelligence whilst simultaneously ensuring that an appropriate legal framework for its control is in place, especially from a European value-perspective vis-à-vis the USA and the Asian markets.