

Data logistics and AI in insurance risk management

A White Paper by

*Dr. Christian Spindler
&
Dr. Christian Hugo Hoffmann*

Contact:

Dr. Christian Spindler
cspindler@dataaheadanalytics.ch

DATA AHEAD ANALYTICS GmbH
Technopark
8005 Zürich
Switzerland

Dr. Christian Hugo Hoffmann
christian.hoffmann@syntherion.com

Syntherion KIG
ETH Risk Center, Scheuchzerstrasse 7
8092 Zürich
Switzerland

Zürich, August 2019

What others say

“An excellent and mature reflection on important aspects of an evolving topic for the industry.”

– Joachim Masur, former CEO of Zurich Insurance Switzerland

“As the (re)insurance industry struggles with how to transform the insurance value chain with machine intelligence, we need more of these kinds of papers where the authors expertly illuminate the opportunities and challenges related to using artificial intelligence and machine learning for intelligently automating the business of insurance.”

– Dr. Jeffrey R. Bohn, Head of Swiss Re Institute

“Big Data and AI are real game changer for insurance companies and their relationship to the customer. The insurer evolves from a claims manager to a real lifetime prevention partner. This deep-dive-paper sets an excellent framework how to face risks and use opportunities for the future of smart insurance.”

– Dr. Andrea Timmesfeld, Head of Public Affairs, Generali

“Insurance companies realize that playing alone is no longer an option to stay competitive. Instead, only trustfully exchanging data with externals will allow to generate AI solutions for challenges in risk management or fraud detection. This White Paper pinpoints in a convincing way the crucial steps insurance companies need to execute if they are to thrive, embedded in a multi-stakeholder ecosystem of connected data and service providers.”

– International Data Spaces Association, Germany

“This paper, presented by an energetic team, is an excellent appetizer of the manifold exciting future applications of AI and ML in (re)insurance processes.”

– Adriaan Ras, Senior Risk Expert at Santam, South Africa

“Insurance companies traditionally used data to make decisions more than any other. However, other industries showed that switching to new methods such as artificial intelligence and machine learning always comes with risks and uncertainties. Insurers additionally have moral responsibilities; it is therefore important to address these challenges in a white paper before moving forward.”

– Swiss Alliance for Data Intensive Services

Contents

Introduction	6
Our understanding of AI	8
Overview of AI in insurance	9
The future of risk assessment	11
The future insurance value chain	12
Case study: Opportunities in AI-supported underwriting in commercial insurance	15
Risk assessment	15
Algorithms today and in future.....	16
Example: Use-based machine insurance.....	17
AI-driven opportunities in the underwriting work process	18
Using AI to deal with new and complex risks.....	19
Example: Climate-change-related business interruption risk.....	20
Challenges in AI adoption and how to address them	21
Where to obtain the data to realize a value-creating AI?.....	21
The Internet of Things as a new source of risk information	22
Small Data	24
How to integrate the data into the corporate legacy?.....	25
From closed data systems to an interoperable world of APIs and micro-services.....	25
Heterogeneous data.....	27
Culture eats technology for breakfast.....	28
How to use data and AI in a fair and responsible manner?	31
Ethics in AI.....	31
Algo.Rule 1: Strengthen competency	32
Algo.Rule 2: Define responsibilities.....	32
Algo.Rule 3: Document goals and anticipated impact	33
Algo.Rule 4: Guarantee security	33
Algo.Rule 5: Provide labeling	33
Algo.Rule 6: Ensure intelligibility	34
Algo.Rule 7: Safeguard manageability	34
Algo.Rule 8: Monitor impact.....	35
Algo.Rule 9: Establish complaint mechanisms	35
Conclusion and take-aways	38
About the authors	42

Introduction

Data is quickly becoming one of the most – if not the most – valuable asset for any organization, given its tremendous volume in our digital era. But not only that – on the one hand, it will continue to grow much further in the future due to e.g. the avalanche of new data coming from a variety of sources, such as connected devices. On the other hand, digital tools such as Artificial Intelligence (AI) and Machine Learning (ML) that harness big data and complexity abound and are becoming more powerful. Many industries are heavily impacted by these trends and technologies, the insurance sector being no exception. The explosion of e.g. their transaction data in combination with AI and ML, methods that are able to map complex networks of interdependencies and that perform best when they can build on high computational power principally available to insurance firms, will generally allow insurance companies to understand their clients and risks more deeply. This is expected to result in improved processes, new product categories, more personalized pricing, and increasingly *real-time* service delivery and risk management.

However, in light of these manifold opportunities and leverage points not unheard of in the insurance business, it resembles a mystery why the digitization is not evolving more rapidly. In fact, according to a recent Gartner¹ study, 85% of data science projects (which together constitute digitization) fail, which provokes the question as to why this happens. More to the point, how can insurance companies make sure their projects are among the successful ones²? General industry literature provides a host of reasons inspired by common-sense: inadequate or unstructured data³, technology for its own sake in lieu of drivers of profits and key performance indicators (KPI)⁴, poor communication, insufficiently incentivized cross-functional teams⁵, missing executive support, over-complicating⁶, and, too narrow a problem focus, not to speak of a lack of transparency about the problem to solve and short feedback cycles⁷.

The question remains whether these are just symptoms or really their root causes. Based on our interview results with executives and risk professionals in mainly the Swiss and German insurance industry⁸, this white paper tackles the following key business problem:

¹ <https://gtnr.it/2oT1Mpn>

² <https://bit.ly/2l1ScuT>

³ <https://tek.io/2K2LIUa>

⁴ <https://blog.datarobot.com/why-most-ai-projects-fail>

⁵ <https://towardsdatascience.com/why-data-science-succeeds-or-fails-c24edd2d2f9>

⁶ <https://hbr.org/2016/12/why-youre-not-getting-value-from-your-data-science>

⁷ <https://towardsdatascience.com/why-data-science-projects-fail-revisited-85fe242c3931>

⁸ In alphabetic order: AXA, Generali, Mobiliar, Santam, SwissRe, Zurich Insurance Group

Legacy infrastructure, missing interoperability, and a lack of comprehensive knowledge on AI and its use cases hinder the adoption of advanced, AI-based, ethical insurance schemes, as appropriate methods or necessary data cannot be provided to the right spot, at the right time.

To address this challenge, we raise the three subsequent questions and provide answers below:

- What is state of the art of algorithm-based insurance / risk assessment methodology today and what is currently being discussed?
- To what extent can data logistics and AI technology enable schemes and practices for insurance companies?
- How can insurers encourage using AI systems in a morally responsible manner?

This white paper is structured as follows: In the next introductory sections, we present our definition of AI and provide a synthesis for the insurance sector. Then, delving directly into the heart of the matter, we bring the case of underwriting with its facets of risk assessment, AI-driven opportunities and complex risks into the spotlight. In the third part of the white paper, we elaborate on challenges of AI adoption in insurances and outline how to overcome them. The paper is closed by reflecting more broadly on shaping ethical AI systems in insurances and beyond.

Our understanding of AI

The precise meaning of AI is somewhat ambiguous; there is a perplexing cacophony of conflicting opinions. John McCarthy originally coined the term AI in 1956, defining it as "the science and engineering of making intelligent machines"⁹. When the field of AI today appears broad and vague, it is perhaps for this reason. Under this definition, all the software (or hardware or robots) we use could be considered as AIs, depending on how exactly we grasp "intelligent".

Interestingly, it seems that AI is actually a rolling target: Had Machine Translation been regarded as the main representation of AI only some years ago, we may regard it as a mere (yet impressively powerful) ML algorithm today. By analogy, almost everything we hear about AI today actually refers to Deep Learning and points to artificial neural networks used in mimicking human abilities to interpret sensory input, e.g. from seeing and hearing. The bottom line with regard to the larger picture of AI is that such sudden rises and falls of different techniques have characterized AI research for a long time.¹⁰ Every decade has seen a heated competition between different ideas, including Bayesian networks, support vector machines, and evolutionary algorithms, all of which take different approaches to finding patterns in data. Such frequent regime changes led to the emergence of the so-called "AI effect", our tendency to discount and underestimate the behavior and power of a (narrow) AI algorithm by arguing that it is not *true* intelligence, but *mere* computation.¹¹

Instead of searching for a rigid binary definition of AI, we could argue that there is a capability spectrum of artificial intelligences and every "agent" can be located on it. A light bulb has close to zero intelligence, whereas, on the other end of the spectrum, we envision a super intelligent system¹² outperforming humans in all relevant domains. For the sake of this White Paper, we regard AI as any data analysis technology that enables business decisions or reasoning about these decisions in an environment which is too complex (in terms of speed or heterogeneity) for humans to perform on the same level for the respective use case.

⁹ <https://www.theguardian.com/technology/2011/oct/25/john-mccarthy>

¹⁰ <https://bit.ly/2Tlj3oM>

¹¹ http://www.pamelamc.com/html/machines_who_think.html

¹² <https://www.ft.com/content/021d3484-fd1d-11e3-8ca9-00144feab7de>

Overview of AI in insurance

Going one step back and acknowledging digital transformation first, we find, in line with our interview results, two basic value drivers contributing the most to that change in businesses: process automation/rationalization and faster, more robust or cheaper predictions. The former does not necessarily presuppose AI, but requires standardization and regularity. Nonetheless, it is often (but mistakenly and narrowly) regarded as the classical use case of AI.¹³ By contrast, AIs ought to be conceived of as prediction machines.¹⁴ In predictive modeling, for instance, AI paves the way for real-time automated decision-making by determining decision criteria from data. Since the insurance industry is largely a data-driven business, AI has the potential to heavily impact their business processes and decisions by monetizing the value of own and third-party data and coping with core business issues. Key AI-enabled strategies in insurance include:¹⁵

- Efficiency creation in underwriting and risk monitoring, but also in other more mundane business functions such as finance, policy admin, and call centers, gives insurers a competitive edge, especially in commoditized markets. A crucial ability in this field is to analyze unstructured datasets using AI to find new risk patterns or to automate contract understanding and processing. An area in which AI-supported risk analysis is viewed as particularly valuable already today, is property insurance, with predictive modeling of hazard exposure and sensitivity of assets. Natural language processing (NLP) plays an important role as underlying technology. Although there are vendors on the market who cater to this field, insurers have repeatedly reported that this problem is far from being solved. Global re-insurers such as Swiss Re and Munich Re started testing the IBM Watson software by exploring the impact of broad market trends on their portfolio to better price risk, but reports showed mixed results in the meantime.¹⁶
- Creating claims workflows that are more accurate and responsive to customer needs. For instance, as pointed out by one of our interviewees, optical recognition software assists with claims management via mobile phone apps. Such examples greatly improve customer experience, since customers may have to pay out of pocket and wait with uncertainty about timelines and

¹³ <https://schweizermonat.ch/kuenstliche-panik/>

¹⁴ <https://www.spektrum.de/rezension/buchkritik-zu-prediction-machines/1603000>

¹⁵ http://www3.weforum.org/docs/WEF_New_Physics_of_Financial_Services.pdf

¹⁶ <https://www.inside-it.ch/articles/48789>

outcomes today. For example, Zurich Insurance employs AI to review paperwork (e.g. medical reports).¹⁷ Shift Technology, as another example, incorporates AI to find patterns of fraud in deep claims datasets, which can then be applied to incoming claims in order to flag potential instances of fraud, and recently raised USD 60 million in Series C funding¹⁸. Moreover, one of our interview partners is exploring Deep Learning NLP to augment their Contact Center with information extraction from free text and speech.

- AI also provides agility to insurers to enable them to develop products against new types of risk in new ways. Alternative data sources, such as repair bills, sentiment analysis of news, remote sensing aggregation etc., can serve as damage proxies if AI is able to model the correlation (or rather the cause-effect relation) between proxies and damages. Apart from that, new contract models such as dynamic pricing and underwriting allow different components of a policy to be priced, bundled or sold separately, such that insurers can develop pay-per-use products. Trov¹⁹ is one example which allows users to purchase miniature insurance policies for specific electronic devices and offers them to turn that coverage on and off at will, equipping users with the opportunity to dynamically steer and control their risk exposure.

In the meanwhile, it is a common practice that insurers make use of their internal data to offer additional services that, on the one hand, complement the product offer and, on the other hand, often help customers to prevent damages. A prominent example are companies which measure driving behavior²⁰ and produce safety recommendations that prompt drivers to improve their behavior, if necessary.

In summary, AI plays a role in any process, where data-based predictions²¹ are made. Various areas of the insurance business model are subject to predictions, such as:

- *Contract granularity*: coupled with goal-based systems that capture clients' aims and lifecycle moments, AI may design insurance contracts in ever higher granularity, both in terms of scope and contract lifetime. Event-based insurance

¹⁷ <https://reut.rs/2Zisubr>

¹⁸ <https://www.shift-technology.com/de/aktuelles/>

¹⁹ <https://www.trov.com/>

²⁰ A great example of the use of telematics combined with behavioral economics can be found on the Discovery Insurer's Vitality Drive: <https://www.discovery.co.za/car-and-home-insurance/why-choose-discovery-insure>

²¹ <https://www.spektrum.de/rezension/buchkritik-zu-prediction-machines/1603000>

is possible if the risk of an event can be quantified hinging on the holistic context the event is embedded in. AI is able to take this context into account by approaches such as pattern identification, scenario analyses or probabilistic modeling in knowledge graphs.

- *Underwriting*: in the risk-taking moment of an insurance company, AI helps making better underwriting decisions by exploiting data in a less biased and more holistic way.
- *Claims management*: AI serves many purposes, such as triaging upon FNOL (first notice of loss), identification of fraud and claims handling preferences of the customer.

The future of risk assessment

In this paper, we exemplify the impact of AI and big data availability on insurances' risk management functions by examining the case of underwriting processes, in particular.

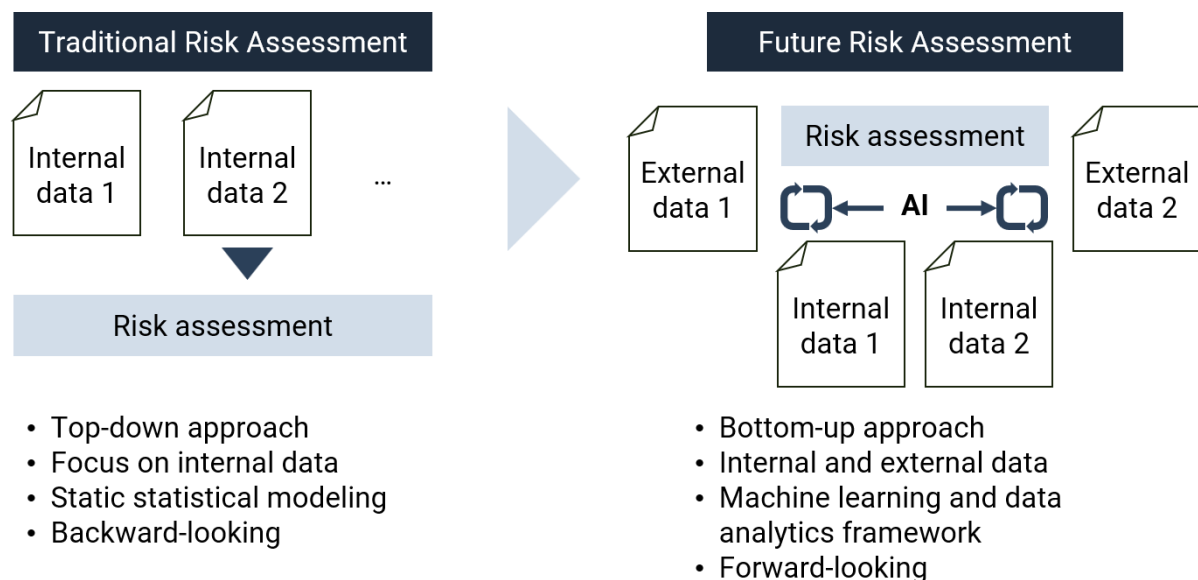


Figure 1. Change in insurance risk management from a top-down to a bottom-up approach

Historically, insurance companies have been evaluating risks for pricing and hedging purposes by building on statistical and actuarial modeling of technical data. The left side of Figure 1 shows the traditional way of risk management in insurance. Based on various data sources within the company, such as claims and contract statistics,

statistical modeling was applied to arrive at a certain risk statement. This process is very static in nature and focuses on historical data only.²²

With the availability of AI methods, increasing tractability and emerging access to external data sources, such as technical data of insured assets, this process can be turned upside down. In future risk management, not only internal data, but more and more external data sources contribute to risk modeling. The process is much more dynamic in nature, as data sources can be on- and off-boarded with little effort, while ML enables model-free risk estimation. Given the importance of systemic methods in complex systems²³ in general, the evolution of risk assessment will envisage a *compositional*, ML-based risk modeling approach of technical, behavioral etc. data, resulting in a more accurate and dynamic evaluation of risk.

The future insurance value chain

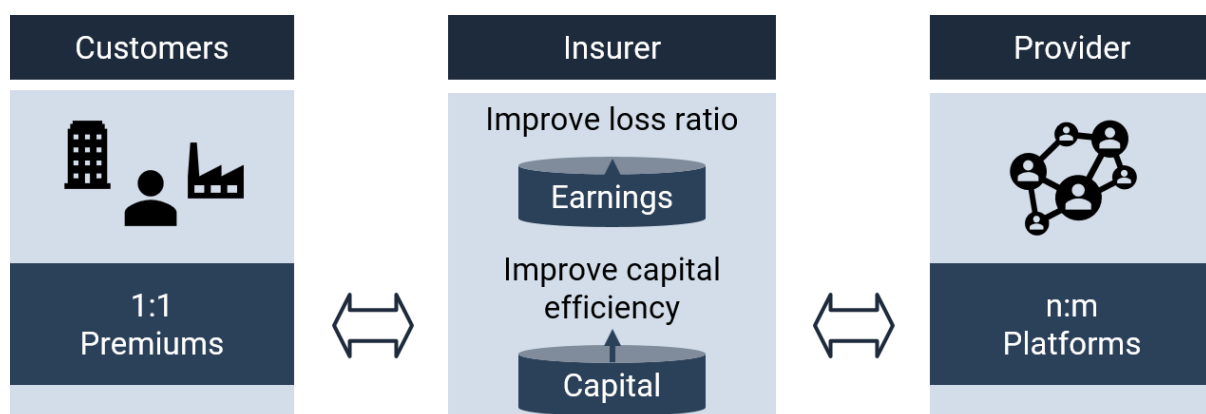


Figure 2. The (future) insurance value chain

Figure 2 indicates that the innovative use of AI has an impact on all elements in the insurance value chain.

For the insurance company in the center, the innovative use of AI/ML can reduce their loss ratio. The significant leverage can be demonstrated by considering that an insurance company with premiums of CHF 1 billion could increase their earnings by CHF 10 million just by improving the loss ratio by 0.1%.

²² To avoid misunderstandings, please note that Big Data sources are "historical" too. However, they expand the scope of what one looks at in an attempt to be forward looking, enabling more circumspect predictions that rely less on limited and biased human judgement of what the future may hold.

²³ <https://onlinelibrary.wiley.com/doi/abs/10.1002/sres.2414>

For the insurance customer, AI-enabled calculations of insurance premiums might lead towards 1:1 schemes (fully personalized insurance), which challenges existing risk and business models of today's heterogeneous risk communities. Nothing less than a new regime will enter stage: If real-time data on individual risks becomes available, then classical risk assessment which is based on analyzing probability distributions becomes obsolete. The question opens up, how much (or few) capital would still be needed by an insurance to fulfill its tasks. The situation challenges the current mandate of insurance companies, which is to pool risks across insurance subjects. It is moreover an open question if regulators will accept too granular risk models or if they insist on pooling to a certain extent. Furthermore, it depends on the degree of diversification of risk, regarding a trade-off between the fairness / accuracy and *cost* of risk (which pooling reduces) which, therefore, has a direct implication for capital efficiency.

1:1 schemes require granular data from customers. This raises the question what an insurance company could offer to customers in exchange for their very personal data. One possibility could be financial reimbursement. In fact, insurance schemes are discussed on the market that would offer premium reductions in exchange for customer data²⁴. More sophisticated but also often more value-adding approaches trade data for information ("data-for-good"): in the case of a property insurance, this could be on-site burglary and flood statistics, a health insurance could distribute recommendations on sports training or nutrition. In turn, more and more customers expect additional features to their insurance, e.g. to support risk prevention. If designed in a smart way, such information can reduce actual claims substantially, for the benefit of the insurer and the insurance subject.

For infrastructure and data providers, the use of AI could enable n:m platforms for digital business ecosystems with which insurers and providers could share more fine-granular data and thus improve cost efficiency. Such ecosystems typically involve multi-sided business models that bear a higher complexity than traditional insurance-customer relationships.

One enabler of multi-sided business models are data ecosystems that allow easy access to broadband data from multiple, unrelated sources. There are various initiatives evolving that foster data sharing among organizations, while ensuring data sovereignty for data owners. In the GSA region (Germany, Switzerland, Austria), prominent examples are the International Data Spaces Association (IDSA)²⁵, Data Market Austria²⁶, and the Swiss Data Alliance²⁷. Multi-sided business models can be

²⁴ <https://www.reuters.com/article/us-insurers-bigdata-consumers-idUSKCN12N05R>

²⁵ www.internationaldataspaces.org

²⁶ <https://datamarket.at/>

²⁷ <https://www.swissdataalliance.ch/>

designed in various ways, but our hypothesis is that platform models with a central broker are most suitable. The International Data Spaces Reference Architecture²⁸ is an example, where such a platform architecture has been realized successfully, and, among other members, supported by the global insurer Allianz.

²⁸ <https://bit.ly/2KzAJvx>

Case study: Opportunities in AI-supported underwriting in commercial insurance

In this chapter, we put a special focus on the underwriting part of the lines of insurance business which provides insurance to companies. We investigate which features of underwriting will benefit the most from AI support. Moreover, new characteristics will emerge from incorporating AI technology, such as premium reduction through risk transparency or real-time, mission-based dynamic insurance.

Underwriting is typically a lengthy process. In commercial insurance, it most often involves the physical presence of an underwriter and inspector at the customer's place, and it presupposes subject matter expertise for assessing risks. If those time and resource-intensive tasks can be, even just partially, supported or automated by AI, the underwriting process as a whole would be more cost efficient and may even transform into a profit center or value-generating activity.

The changes in underwriting are large for applicants as well as insurances. AI connects the applicants with the insurers and also streamlines the process of insurance. Risk assessment seems to be predominantly affected, while also, for example, fraud detection is improved in the wake of AI (the latter, however, is not addressed by this white paper).

Risk assessment

AI technologies are able to process relevant information and assess the risks around an insurance applicant often more precisely and certainly much faster than humans. Previously, underwriters relied mostly on information from the applicant for risk assessment, with little means of verifying the validity and completeness of information.

However, scenarios that insurers and applicants face today are becoming ever more complex. This is articulated in

- more diverse risk scenarios customers wish to insure for
- more and more diverse information underwriters have to process
- shorter time to process individual applications, leading to higher risk of making errors (both operational and judgmental).

AI provides solutions to all three aspects. ML, given sufficient data, finds patterns in diverse scenarios and helps facilitate/automate the underwriting process. Relevant

data will stem from additional new sources, including social media, where human underwriters alone would find it difficult to exploit such data, to integrate and process it in a cost effective and efficient way.²⁹ Concrete use cases in risk assessment for underwriters employ a spectrum of AI technologies, such as:

- Computer vision for remote sensing, e.g. flood risk
- Sentiment analysis with NLP approaches for real-time risk scoring
- ML in terms of providing price suggestions for different customers based on their individual risk factors

In practice, insurers experience difficulties to realize the opportunities in real business processes. Multiple academic examples have been successfully demonstrated. Yet, no productive implementation of these methods came to our attention in the investigated markets.

Algorithms today and in future

A broad range of algorithmic methods are applied in insurance companies today. Customer intelligence seems quite mature, with typical classification and recommender approaches used. Pricing and risk modeling is often done with General Linear Models (GLMs) and credibility models, while more advanced models have not been observed in the investigated market.

However, many insurance companies assess more advanced algorithms that are useful as soon as the underlying data logistics have been established. For future applications, particularly policy optimization using reinforcement learning and Bayesian methods, that can e.g. better cope with small data³⁰, are named by our interview partners. Recurrent neural networks and ex-post gradient analysis are further regarded as useful for risk modeling. Finally, causal inference³¹ should be the preferred way for risk analysis, above correlation-based approaches of today. Deep Learning, despite its presence in public discussion, is not regarded as very useful, apart from the special application of image processing (e.g. in the claims process or in remote sensing). For managing insurance portfolios, Deep Learning is not seen as a viable option, mainly due to the missing explainability of models as well as the minor added value above simpler methods.

²⁹ By that, we do not wish to underestimate the costs of successfully implementing AI, especially in the short-medium term.

³⁰ <https://towardsdatascience.com/what-is-bayesian-statistics-used-for-37b91c2c257c>

³¹ <https://www.ibm.com/blogs/research/2018/12/ai-year-review/>

Example: Use-based machine insurance

Today, production machines in industrial manufacturing are typically insured on an annual basis. Some machines carry individual insurance contracts, otherwise a frame contract for an entire production plant applies. The insured risk consists of two elements: Known and unknown risks. The former are inferred from the type and purpose of the machine and from the typical production environment the machine is embedded in. The latter may arise from peak-load periods in case of high machine utilization and from different tool handling by different machine operators. An insurance estimates the known risks based on experience but has typically little means to assess the unknown risks.

Industrial IoT has the potential to measure the machine state at all times and to deliver a *real-time* picture of machine operation. With such condition monitoring, unknown risks are getting transparent. If the machine's data is fed into a Digital Twin, a model can infer the risk of failure based on the current data and the entire historic mission profile of the machine. This goes even further. If the Digital Twin is interoperably connected with a production planning system and a customer order book, the model can take this future machine payload for a scenario analysis to create a *forward-looking risk estimation*.

For insurances, interesting business models arise from this scenario. According to the value chain model in Figure 2 above, insurers may use such a model to help their customers in taking preventive measures against machine failure. To realize advanced and robust models, insurers can draw data from specialist providers, such as aggregators of production fingerprints of machines, or commodity providers for data such as weather and alike. The customers benefit from lower business interruption risk, i.e. higher overall equipment effectiveness, a standard KPI in assessing manufacturing processes. Customers may also profit from lower and use-specific premiums if the insurance company offers such a scheme.

One insurer raised the question whether insurance companies will be the first point of contact in this scenario. Rather than that, predictive asset management could go via asset manufacturers and respective service level agreements between the manufacturer and the owner of the asset, up to the point of asset-as-a-service business models, where ownership stays with the manufacturer and only machine uptime is purchased by the customer. The latter model provides strong incentives for the machine manufacturer to care about the health state of the machine.

Another insurer sees medium-term opportunities in deploying more parametric insurance models in the supply chain or in agricultural value chains. New data plays a crucial role. For instance, Concirrus Quest Marine provides live marine weather, fleet

and port data to insurers (e.g. warning when a vessel enters an exclusion zone)³². Furthermore, technologies such as the 5G mobile network standard will be helpful in reducing the latency of data transmission, while the distributed ledger technology (Blockchain) provides immutability and traceability of the data. Irrespective of the concrete application, the great potential of AI in commercial insurance is seen particularly in catering to SMEs, as scale effects from technology can close their underinsurance gap. Thereby, larger companies are typically more proactive in adopting new technologies.

In the medium and long run, enhanced risk assessment altogether could become the key competency and business of insurances as new fin- and insur-techs are entering the market, offering more user-friendly, cost-effective, and appealing products at the front office. In that case, intensified by a low profit margin environment, steering and optimizing risk-return ratios will turn out to be the essential KPI of insurers and financial intermediaries altogether.³³

AI-driven opportunities in the underwriting work process

Today's office processes are often still characterized by a great variety of different software applications that constitute a diverse value chain, such as customer, claims, and product databases. Underwriters use, if not waste, much of their time with manual data transfers from one application to another, while spending only little time with higher-value tasks such as reasoning from information, selling, or engaging with brokers. Specific AI-powered applications, delivered at the right spot and time, may significantly boost the performance of underwriters³⁴:

- In an automated interaction with agents and customers, cognitive applications may sort service requests and respond to queries by automatically providing targeted information. As around 80 % of data in today's insurances is text, as reported by one of our interview partners, NLP is regarded as one of the most important AI technologies for now.
- Underwriting tasks are per se very document-prone and often paper-based. Any structured information extraction from scanned documents (AI-enabled optical character recognition, named entity recognition, etc.) saves time and manual labor. The structured data can then be used for creating recommendations to the underwriter (by enabling the user to address queries such as which similar cases do exist, etc.).

³² <https://www.concirrus.com/quest-marine>

³³ <http://www.syntherion.com>

³⁴ <https://insuranceblog.accenture.com/ai-offers-insurers-a-big-step-up-in-underwriting-performance>

- In the process of closing new contracts or renewing existing ones, AI-powered calculators may infer the win probability and current risk patterns to help underwriters focus on prioritized tasks.

Using AI to deal with new and complex risks

Within the new edition of the Risk Barometer³⁵ of global insurer Allianz, business interruption, cyber incidents, and natural catastrophes are named as top business risks in 2019.

Despite cyber incidents being still seen as a new risk, global premium revenues for cyber-insurance are supposed to rise up to USD 7.5 billion by 2020, according to the researcher Statista³⁶. However, professional insurers such as Warren Buffett's Berkshire Hathaway AG³⁷ already warn of the largely unknown and potentially too big risks behind a cyber-cover.

Apart from new risks, the complexity of risks continuously increases. Corporate ecosystems and supply chains have spread globally across many countries and have increased the complexity of commercial risk assessment. The rise in trade and business regulations³⁸ imposed around the world has led to a similar development.

Looking at the Risk Barometer ranking, many risks are strongly interconnected with each other. For example, rank 8 risk, "climate change", is interwoven with rank 1 "business interruption" (e.g. in terms of supply chain interruptions after climate-attributed extreme weather events), rank 3, "natural catastrophes" (e.g. growing frequency of hurricanes as a result of warmer mid-Atlantic ocean water), rank 4, "regulatory changes" (e.g. emission taxes), and so on.

In line with the report, catastrophe cover is a strongly rising need due to climate change. In 2017, the net underwriting deficit among US property and casualty insurers leaped to USD 23.2 billion, from USD 4.7 billion in 2016.³⁹

Applied to business risks, AI, given interpretability, will have a major impact by drawing on data describing processes or properties that are unrelated at first glance, but relate them to each other, by extracting patterns and finally making complex risks manageable.

³⁵ <https://bit.ly/2WztBqp>

³⁶ <https://www.statista.com/statistics/533314/estimated-cyber-insurance-premiums/>

³⁷ <https://bloom.bg/31jlxHV>

³⁸ <https://insuranceblog.accenture.com/rising-trade-barriers-snap-insurers-digital-transformation>

³⁹ <https://vrsk.co/2uYIt0R>

Example: Climate-change-related business interruption risk

Business interruption from supply chain vulnerability is one of the most complex risks in global business environments. Climate change is a major root cause of supply chain interruption, mostly resulting from acute risks of extreme weather. However, also company-internal processes are subject to business interruption because of climate impacts. In the extended heat period in Central Europe in 2018 that caused high temperatures and low water levels e.g. in the Rhine, the global chemical manufacturer BASF had to adjust production and logistics at its Ludwigshafen headquarters.⁴⁰ The quantitative assessment of such risks is crucial to adequately price insurance products, and also to design parametric products around business interruption.

Climate change also has material regulatory and liability implications. Since conventional emission trading schemes have been in place for several years, emissions regulation is tightening, and new targets or instruments are shaping. Germany's recent discussions about emission taxation impacts on business risks of companies and requires strong improvement in both fact-based carbon accounting of business processes and industry-wide climate risk reporting and disclosure standards. The latter will increase exposures for directors and executives who oversee disclosing the risks.

A comprehensive framework for risk disclosure has been developed by the TCFD⁴¹. For insurers and climate risk analysis experts, who work with these risks on a daily basis, it is clear that the complexity of climate risks basically eliminates the suitability of generic, industry-focused approaches. Processes differ substantially among companies within the same industrial segment. An AI fed by an underlying modern data logistics layer is best suited to account for such evolving risks in a robust and reliable manner.

⁴⁰ <https://www.basf.com/global/en/media/news-releases/2018/08/p-18-287.html>

⁴¹ <https://www.fsb-tcfid.org/>

Challenges in AI adoption and how to address them

The previous chapter described the opportunity of AI-supported underwriting in general and gave two examples. We will now turn to the question why the observed progress of AI adoption in insurance is rather slow, given the articulated business potential.

Digital transformation and the substitution of manual processes by AI-driven automation is an ongoing activity in companies. Initially, low-hanging opportunities were realized during the raising branch of Gartner's Hype Cycle⁴². Regarding the next stage of digital transformation, it seems that certain factors prevent a more rapid evolution. In the following, we raise and elaborate on three questions along the value chain of AI in the insurance business and identify ways to overcome real or virtual obstacles for pushing digitization to the next level:

- Where to obtain the data to realize a value-creating AI?
- How to integrate the data into the corporate legacy?
- How to use AI in a fair and responsible manner, ensuring long-term business success and its *raison d'être* tout court?

The following sections are addressing these questions. First, we introduce the (industrial) Internet of Things (IoT) as a fast-growing source of risk-relevant data. Next, we investigate how companies can make efficient use of new data and thereby, step-by-step, transform their infrastructure. Finally, we demonstrate how companies can draw business value from employing algorithms while adhering to ethical and compliance standards.

Where to obtain the data to realize a value-creating AI?

Insurance-internal data are an important source of information also for new business cases. Yet, this data is limited in size and scope and only grows with a very moderate rate. According to one of our interview partners, insurance data is notoriously poor, in particular on the non-life insurance end of things. He identifies three contributing factors to this:

1. The power of the broker and outsourced underwriting functions (in the case of specialist lines of business) in terms of control over client and risk data.

⁴² <https://gtnr.it/2vTTphv>

2. The sheer inherent and rapidly evolving complexity of the risks non-life insurers underwrite, as one key reason why most contracts are limited to one year at a time.
3. Competition amongst insurers – the "irritation factor" and "anti-selection motive" presented by more stringent underwriting requirements contributes to brokers and clients preferring insurers that ask less questions.

External data sources greatly enlarge the scope of accessible information (weather forecasts as a standard example) and grows with much higher rate. Two main sources seem particularly relevant to trigger new applications or efficiencies in insurance: Behavioral data and the (industrial) IoT. With regard to fraud prevention for instance, exploiting behavioral data such as on the tone of voice and facial expressions at the moment of underwriting is a classical case for ML. One of our interviewees estimates that in life and health insurance, more than 40 % of risk information can be obtained from behavior monitoring alone. In non-life insurance, behavioral data is also important (e.g. does a certain pattern on how an employee operates a machine indicates process issues that could lead to insurance claims?), but it will most likely be a much less low-hanging fruit. In the following, we particularly emphasize sensors as a risk-relevant data source, which are the main data providers of the IoT.

The Internet of Things as a new source of risk information

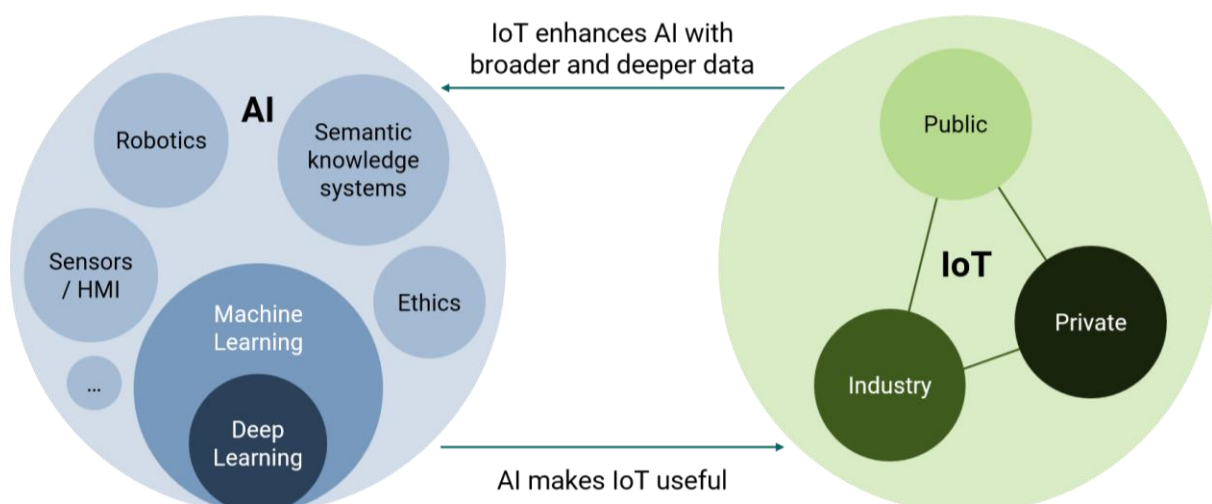


Figure 3. The relationship between AI and IoT

All AI-based processes in commercial insurance presented in the previous section benefit from a growing data source: the (industrial) Internet-of-Things (IIoT). Condition monitoring provides accurate information on an asset's historic mission profile and current state. Asset tracking by means of RFID, LoraWAN, GSM network or Bluetooth technology allows for an evaluation of the current risk exposure of assets or trade goods. And the combination of data from a company's ERP system with external, open

data sources amounts to an enrichment of the risk assessment with supply chain information, facilitating the quantification of complex risks such as business interruption.

Consider a concrete example of bottom-up digital transformation: The digital power hub for construction sites.⁴³ A power hub is literally the first device to place on a construction site, as it delivers electric power to all machines and tools that require it. It acts as an integrator for the diverse landscape of contractors and sub-contractors working on large construction sites. Anybody who needs power may access the power hub, plug in one or several devices and use it as long as required. So far so good. This typical scenario, albeit common practice on millions of construction sites around the globe, lacks any transparency of how the power is used. An insurer may ask:

- Is the employee getting power from the plug eligible to use the powered tool? If not, how is she covered in case of an accident?
- Is the process to be powered executed in the prescribed way? If the power is e.g. used to drive a screed drying fan, does the fan run at least long enough to guarantee that the floor has reached its contracted payload in the finalized building?
- Are groundwater pumps running stable or are there signs for pump failure that would cause insured damage to the construction site?

These and more such questions can be answered if the insurer obtains data from a digital and connected power hub. The insurer may use the data directly, but also employ AI to identify devices from power usage patterns, spot upcoming failures in machines and predict the real-time risk of the construction site based on the number and type of contractors present.

It is evident that IoT and AI form a virtuoso interplay and we expect data from one side and algorithms from the other side to create vast opportunities in the insurance sector. Insurers see the high potential in IoT, but currently lack ways of finding information in the breadth and depth of sensor data. Historically, the dimensionality of IoT data was typically broad, i.e. a large number of data sources produced only short coherent time-series samples. With growing digital connections of industrial IoT, this is currently changing and opens up possibilities for AI above basic regression approaches to be useful for underwriting. Successful IoT cases that are repeatedly named today are mainly two:

- Primarily automobile insurance, where the customer is awarded a premium reduction based on the data from a drive recorder or crash recorder.

⁴³ <https://www.twinstruction.de>

- Insurance of short-living goods, e.g. food during transport, where insurance companies develop IoT-controlled, parameterized products to insure cargo ship freight and logistics delivery.

Small Data

Many applications of ML rely on a massive quantity of data to train models. This is particularly true for Deep Learning models in applications that involve image, speech or text processing. In reality, large, clean data sets are hardly available in a business context. Instead, the heterogeneity of data sources, including IoT, leads to little data silos that contain actually comparable data, however many of them. In particular, IoT shifts AI closer to the edge, i.e. to the sensors, controllers or gateways, where data is generated and processed in the first place. On edge devices however, computation and memory capacity are typically a scarce resource. Complex, high-parametric and data-hungry algorithms can often not be executed due to hardware constraints. There are three approaches to mitigate this problem: (1) Equipping edge devices with better hardware. Indeed, we see a rapid development of GPU-powered edge devices, e.g. NVIDIA's Jetson product line⁴⁴. However, improving hardware is usually the last resort, given the comparably high costs. (2) Employing down-scaled, edge-adapted algorithms. Google's TensorFlow Lite⁴⁵, for example, enables ML models to be deployed on mobile and IoT devices. (3) Even if powerful hardware and edge-adapted algorithms are available, the business may still lack the right data in the required quality. In that case, future AI will necessarily have to develop strategies to use less data.⁴⁶

⁴⁴ <https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-tx2/>

⁴⁵ <https://www.tensorflow.org/lite>

⁴⁶ <https://hbr.org/2019/01/the-future-of-ai-will-be-about-less-data-not-more>

How to integrate the data into the corporate legacy?

From closed data systems to an interoperable world of APIs and micro-services

As in other business segments, IT and data systems in insurances are heterogeneous and hardly interoperable. This problem emerges from either a diverse landscape of applications and vendors serving different needs, or from monolithic Enterprise Resource Planning (ERP) systems that are complex in setup, cumbersome in operation and costly in maintenance and expansion. Use cases that employ AI to work with data across this system landscape are hindered by the lacking interoperability of systems, meaning it is almost impossible to consolidate and analyze data. Concretely, insurers are blocked in rapid AI adoption, because they need to:

- carefully evaluate which process is suitable for ML and which data is required to deliver a working AI solution.
- be clear on the desired outcome and the best-suited algorithm and ML technique.

The major lesson to bear in mind is that these aspects can frequently not be specified in a top-down fashion, but a more granular, bottom-up analysis of various hypotheses is required as singled out and suggested in Figure 1 above. A prerequisite, however, is interoperable data access for any authorized stakeholder at any time. Interoperable data access creates transparency about which data is available in the company, in what quality, and where data gaps exist that prevent business questions to be answered.

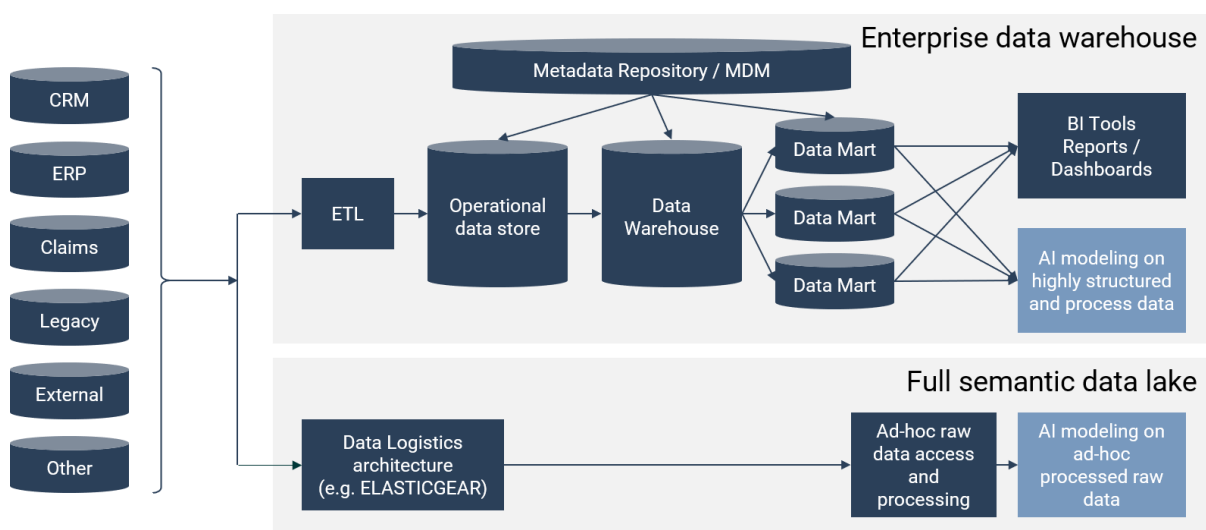


Figure 4. Enterprise data warehouse vs. full semantic data lake architecture

Today, typical IT architectures use data warehouses as the workhorse for data management. Data warehouses are highly organized and structured IT systems for data storage (such as indicated in the upper panel of Figure 4). Data is transformed and loaded in a predefined manner into the warehouse. Changes on the structure of the data warehouse are associated with cost and often hardly possible due to mutual dependencies of the internal sub-systems. Data warehouses are highly tuned to solve a specific set of problems but catering any new use case outside the original scope is typically hard.

The predefined structure and aggregation level of data, the complex system architecture of many interrelated databases, batch-oriented processes, and the missing interoperability to external data sources poses obstacles to the fast development of AI-supported applications. Interoperable data access would allow an AI to learn from a much richer set of data, while direct access to raw data sources without intermittent ETL processes would enable the AI to identify more fine-grained patterns in data if they exist, and to learn aggregating data by itself.

How can data of companies with a legacy data warehouse infrastructure be made accessible for using effective AI applications? Robotic Process Automation (RPA) is a technology that raised commercial interest in the last years. RPA is the concept of using software programs to mimic user behavior on a computer system. An RPA program can observe a user navigating through a data warehouse, collecting data from various systems and consolidating the data as input feed to an AI algorithm. Given that the systems in the data warehouse do not necessarily support single-sign-on access, do not necessarily output data as text files (instead, users have to copy data out from a BI view), and may require user input before displaying data (e.g., a time range or scope in which the data is to be supplied), automating this journey for daily execution saves much time. However, when AI-driven, the output is again focused on a certain use case and training an AI for different cases typically requires different data, which is not within the scope of the designed RPA application.

Our recommendation to resolve this dilemma is to modularize data systems. Instead of designing complex processes and systems to integrate data, companies should focus on making disparate data sources accessible, making any data point searchable and instantly retrievable, and using micro-services for dedicated applications on the data. This approach is visualized in the lower panel of Figure 4. One technology that allows for full interoperability of corporate data builds upon an orchestration and industrialization of open software.⁴⁷ The transition to micro-services can be performed step by step: Processes that ought to be transformed with priority are those that are of high value (either internally or towards customers or partners) or create

⁴⁷ <https://elasticgear.com/en/>

major waste. Based on the applications used in the processes and the way to use them, targeted micro-services can be designed with APIs (interfaces) as publicly accessible as possible. Insurers in Switzerland started to follow these recommendations and to make databases of legacy systems API-ready.

Heterogeneous data

On top of missing interoperability, the existing and growing heterogeneity of data is an obstacle to AI adoption. Heterogeneity bears two meanings in this context: First, in a complex world, many different processes need to be described and quantified, e.g. to estimate risk in insurance. Think about new common sources of real-time data across the industry, such as

- Sensor data from policyholders' cars, homes and workplaces
- Data from drones and satellite imaging
- External data sources like government databases and social media.

This kind of heterogeneity is unavoidable or even desired if the insurance cover is to be offered in a granular and process-specific manner. Secondly, data capturing the same process is often not formatted in the same way, as can be seen for example in the different ways to write today's date. This sort of data hygiene issue is hindering fast adoption of AI because it artificially shrinks the volume of useful data to train algorithms.

There are two ways to improve the second aspect of data heterogeneity: First, by establishing standards for data formats and exchange, software guaranteeing unique representation of data anywhere in the system can be developed. For instance, XBRL⁴⁸ suggests a standard exchange language in financial reporting. The drawback of this approach is that it takes time to train people and develop systems that allow the seamless handling of a new standard, as well as the risk of missing out real-life scenarios in an artificially created restriction. The second way is to design an own AI to cure the data hygiene heterogeneity on the fly, i.e. at the very time the data is needed in an application, e.g. to train the actual business process. AI companies such as Tamr⁴⁹ are exploring this direction further. A drawback of this approach is the lack of data integrity after automated processing – a majority of insurance executives agree that automated systems contribute to new risks including fake data, data manipulation, and inherent bias.⁵⁰

⁴⁸ <https://www.xbrl.org/>

⁴⁹ <https://www.tamr.com/>

⁵⁰ https://www.accenture.com/_acnmedia/PDF-79/Accenture-Technology-Vision-Insurance-2018.pdf

Culture eats technology for breakfast

We trust that the reader is familiar with the phenomenon that technological challenges, as described above, are only one part in the equation of the digital and AI transformation. In fact, technology is usually the part which is solvable by a sequence of concrete steps.

Another, substantial, but vaguer part in the equation is the adoption of a digital mindset in the company's management and workforce. Insurance insiders observe that many companies talk about being "big-data-driven", whereas the number of firms actually acting on their culture and organization is not as large. Instead, one is often confronted with rigid structures and path dependency in insurance firms (especially, in larger ones characterized by hierarchies and long decision times). Yet, our interviews revealed a broad spectrum of maturity levels in insurance companies, from leaders that have a sophisticated understanding of whether and where to employ current AI technology in their processes, to companies that are just starting to build up a data analytics group.

Data-driven processes, technology and business models require different mindsets in quite some respects:

- Becoming more agile, forward-looking and experimental, but if you fail, fail fast
- Transition from deterministic, rule-based calculation systems to algorithms yielding probabilistic results ("The commercial insurance risk score for this client is A because it belongs to sector X" vs. "I am 90 % confident that the risk score is A because the client is similar to group Y of clients")
- Data quality is match-winning in AI, whereas poor data creation through staff leads to poor performance of AI (or the need for expensive data cleaning)
- Change of business models from product sales to providing solution ecosystems to clients (e.g. "underwriting a commercial asset insurance" vs. "offering a risk management solution to clients that helps preventing losses") necessitates thinking in long-lasting, service-focused client relationships.
- The insurance business is regulated and actuaries work in a quite standardized environment. Changing actuaries and the supervising regulator to rely equally well on probabilistic outcomes of ML in the future (in contrast to actuarial models of today), poses another cultural challenge.

Even though it is a task for change management and leadership to build a corporate framework for the evolution of a digital mindset, technology can help here. We suggest that the lower-skilled employees have instant access to all raw data of a company, constrained only by a minimum of legal access boundaries, and that the higher-skilled

ones learn to process and analyze data on their own behalf, for the benefit of their own work. This democratization of data access fosters data thinking and creates a digital culture in a bottom-up approach, without managerial initiatives. Another aspect to consider is today's frequently poor communication and visualization of a data-product-output. Many products focus on the establishment of management dashboards, but merely create an overview of aggregated data and fail to articulate clear information to enable decision-making. A third improvement option is constituted by enhancing IT in companies to approach the work of business-sided data scientists and data engineers, for instance by opening and facilitating IT to work with Python.

Paving the way for and cultivating data thinking as well as establishing a bottom-up digital culture across functions and hierarchies are key to prevent data science projects from failing because a poor fit between a data problem and its solution or method, as well as a selection bias towards certain hyped data science tools, have been observed. In other words, in opposition to a flourishing digital culture, many current data science projects have a tendency to fail at the end of the day because they often start in search for the problem to a stipulated solution or method and not with the business problem that one intends to solve. Data managers then make the mistake to start working without clear managerial guidance. The result is disappointment on both ends. Hence, it is also critical to maintain a *common language* and a *clear protocol* on how actions are triggered in order to not walk into Maslow's trap: if all you (choose to) have is a hyped data science hammer, every data-related problem seems to be a nail. At the same time, managers and business leaders should make sure that the problem they pick is overly relevant for their business.

Business model innovation is a further aspect required for (or facilitated by) AI adoption where cultural aspects come into play. One insurer regards first-mover risks or disadvantages as an obstacle. Why should a company engage in adopting AI, particularly from external vendors? If the business case was successful, then the technology would be available for other insurers as well, preventing any long-term competitive advantage. We do not find this concern persuasive as the argument would also hold for all the other cases where insurers rely on third-party providers to tackle a business problem (successfully!) - e.g. in terms of application-specific risk assessment software solutions: even though those tools are available to their competitors as well, the performance of the different institutions in the respective regard is far from being uniform; at the end of the day, it matters what decisions are made by the people in the organization and not the output from a software tool or AI technology that are, on top of that, not used in isolation, but in a system made up of many different software packages, tools, humans, teams, processes, external effects, etc.

In practice, we observe such complex interwoven systems and, therefore, acknowledge particularly the cultural inertia limiting the speed of technology penetration, hence competitive advantages may last longer than anticipated.

How to use data and AI in a fair and responsible manner?

Ethics in AI

Regulators (such as Germany's financial regulator BaFin⁵¹) and large insurance companies (such as Allianz⁵²) or specialist insurance markets (Lloyds⁵³) recognized the risk behind and ethical relevance of AI systems which in more and more cases are being used to make decisions that have a profound impact on our lives. For example, what happens if an AI-powered claims adjuster rejects a life insurance claim based on an ill-founded suspicion of fraud at a time that a grieving parent needs the money to pay a mortgage after the death of the family's breadwinner? Or would it be fair for an algorithm to decide that someone with a low credit score would pay a higher premium, not because of being a less safe driver, but because of being more likely to file a claim for a small accident than a wealthier driver who can pay out of pocket? And, profoundly, what would be the impact on the insurance systems as well as insurers' business models if AI-based risk measurement procedures became more and more granular and individual risk estimates became more and more precise? By raising such questions, we enter the realm of ethics which, of course, does not lend itself to closed-form solutions. (Business) Culture, for example, plays an important role in shaping what people conceive of as morally good and bad.

However, with the potential non-universality of ethical principles in mind, it would not be desirable or feasible to teach organizations and their employees which actions and practices are right or wrong when it comes to employing AI responsibly. A meta-ethical endeavor seeking to understand the nature of ethical properties and evaluations, however, can be more promising by addressing questions such as "What are algorithmic systems designed for the benefit of society?" and "How can we tell that individual and collective freedoms and rights (comprising human rights) are strengthened, not undermined, by AI systems?" Regulations and businesses' codes of conduct designed to protect these norms must remain enforceable. Against this background, the Bertelsmann Foundation suggested the following 9 meta values, i.e. values that hold irrespective of concrete moral systems or corporate responsibility (CR) programs company stakeholders adhere to. These so-called Algo.Rules⁵⁴, released in March 2019, add value to the insurance context and, therefore, we restate and embed them briefly in contrast to other proposals that lack universality, reference

⁵¹ <https://bit.ly/2Xuat9j>

⁵² <https://bit.ly/2Xv67Pz>

⁵³ <https://www.lloyds.com/news-and-risk-insight/risk-reports/library/technology/taking-control>

⁵⁴ <https://algorules.org/startseite/>

to AI and comprehensiveness (individual codes of conduct or CR programs) or applicability to the insurance industry (e.g. the project of creating values-aligned AI by the well-known Future of Life Institute⁵⁵ or the OECD⁵⁶) or both; e.g. also recall in this connection the set of principles for the healthy development of AI regarding R&D, use and governance proposed by the Beijing Academy of Artificial Intelligence in May 2019.⁵⁷

Algo.Rule 1: Strengthen competency

This postulate emphasizes that the function and potential effects of an AI system must be understood by those who develop it (software engineers) and make decisions based on it and its predictions (managers). Sharing individual and institutional knowledge between departments and risk silos such as operational risk and underwriting⁵⁸ as well as promoting interdisciplinary exchange across task areas are just as crucial as ensuring adequate skills development which includes workshops and trainings for developers and new employees to complement their technical expertise. However, (at least) for the insurance sector, emphasis should be put on the second part of this principle (“the function and potential effects must be understood by decision-makers”) because there is evidence⁵⁹ resulting from peculiarities of large organizations such as Allianz that less (or insufficient) technical knowledge and understanding exists the further up we go in the hierarchy:⁶⁰ To managers and executives actuarial mathematics has been a black box for many decades, and therefore it is not surprising if the situation will be the same once data analytics and ML algorithms are in place. Therefore, even organizational and structural changes might be indicated.

Algo.Rule 2: Define responsibilities

According to this guideline, a natural (a manager, a coder or a group) or legal person (the insurance company) must always be held responsible for the effects involved with the use of an e.g. AI underwriting system⁶¹ whose performance is easily undermined if its predictions hinge on inaccurate, manipulated and biased data. Despite the fuzziness and complexity of processes, activities and projects at insurance firms -

⁵⁵ <https://futureoflife.org/valuealignmentmap/>

⁵⁶ <https://bit.ly/2Y11pOk>

⁵⁷ <https://technode.com/2019/06/13/insights-beijing-ai-principles-a-step-in-the-right-direction-but-still-not-enough/>

⁵⁸ <https://riskandinsurance.com/managing-claims-across-silos/>

⁵⁹ <https://hbr.org/2009/10/the-six-mistakes-executives-make-in-risk-management>

⁶⁰ Hoffmann, C.H. 2019. Thinking about Thinking and Practices – What it Means to Reach Effective Risk Management Decisions in Banking. Forthcoming: Journal of Applied Corporate Finance.

⁶¹ <https://emerj.com/ai-sector-overviews/ai-for-claims-processing-and-underwriting-in-insurance/>

many parties contribute, where does one process start and end - accountability must be clearly assigned⁶². The respective accountable person ought to be aware of the responsibilities associated with their tasks, which also applies to responsibilities shared⁶³ by several persons. The allocation of responsibility must be fully documented and transparent for internal and external parties. Given the character of AI systems as tools rather than persons⁶⁴, responsibility may not be transferred to the algorithmic system itself, nor merely its users or people affected by it.

Algo.Rule 3: Document goals and anticipated impact

The objectives and expected impact of the use of an AI must be documented and reviewed by decision-makers within the firm prior to implementation. The documentation includes revealing the underlying data and calculation models used in risk management⁶⁵ teams to management boards and interested/affected parties within the firm. Particularly in the case of Deep Learning systems and in dynamic areas of application that are subject to frequent changes, an impact assessment ought to be conducted frequently. The risk of discrimination and other adverse consequences should be accounted for in risk management functions.

Algo.Rule 4: Guarantee security

The security, reliability and robustness of an AI system including its underlying data should be tested before (security by design) and during its implementation to prevent accidents, misuse and manipulation⁶⁶. At the same time, it must be noted (which the Algo.Rule team neglected) that AI itself too is key to increase cyber-security.⁶⁷ The effects should be investigated⁶⁸ in detail.

Algo.Rule 5: Provide labeling

The use of AI must be identified as such. The firm's stakeholders (e.g. customers) interacting with algorithmic systems, e.g. with AI-powered chatbots on which insurance companies are betting⁶⁹ must be able to identify that they are communicating with a program or that a decision or prediction is based on an

⁶² <https://www.signavio.com/post/process-thinking-insurance/>

⁶³ <https://bit.ly/31hTbP2>

⁶⁴ <https://bit.ly/211tRW5>

⁶⁵ Hoffmann, C.H. 2019. Thinking about Thinking and Practices – What it Means to Reach Effective Risk Management Decisions in Banking. Forthcoming: Journal of Applied Corporate Finance.

⁶⁶ <https://bit.ly/2WQWvRP>

⁶⁷ <https://bit.ly/2K9zDYb>

⁶⁸ <https://medium.com/swlh/how-artificial-intelligence-is-changing-cyber-security-a243294ccdfc>

⁶⁹ <https://bit.ly/2MHXfx>

algorithm. This is particularly important in cases (of chatbots) where the AI imitates a human being in how it interacts (e.g. through language or appearance). It is less important in cases where data analytics (one black box) replaces actuarial mathematics⁷⁰ or other very technical approaches (another black box) because one way or the other those systems are not accessible by non-experts and labeling would not change much about that, which also informs the specification of the next principle.

Algo.Rule 6: Ensure intelligibility

The decision-making processes within an AI system ought to be comprehensible and transparent⁷¹ to enable stakeholders to question and review decisions resulting from an AI. Without releasing its secret sauce, the insurance company should publish information about the data and models on which the system is based, its architecture and potential effects in “easily” understood terms. To not overwhelm insurers that deal with sophisticated risk measurement approaches, it ought to be taken into account that this call for transparency and intelligibility should not primarily empower those external stakeholders that are complete lay people or end customers (since they are just too far from the engine room of risk assessment), but the already more competent mainly internal ones. On the other hand, it is important for the modeling teams and developers to check if an objective can be achieved without a significant loss in quality by the use of a less complex program that involves an easier to understand mode of operation.

Algo.Rule 7: Safeguard manageability

An AI must be manageable throughout the lifetime of its use, which goes hand in hand with postulating rule #2 because to bear responsibility presupposes to maintain sufficient joint control over the AI system. This involves ensuring broad oversight of the entire system, even when tasks are distributed across various teams, departments within the firm as well as avoiding that the system’s complexity ever exceeds the capacity of human oversight and interference. This rule is of particular concern for Deep Learning approaches that are often regarded as black boxes⁷². If sufficient manageability cannot be guaranteed, the AI system in question should not be used.

⁷⁰ <https://bit.ly/31d5n3G>

⁷¹ <https://on.wsj.com/2WP7GKR>

⁷² <https://algorithmenethik.de/2019/04/17/ki-und-moral-eine-grundlagendebatte/>

Algo.Rule 8: Monitor impact

This rule says that the effects of an AI must be actively monitored and reviewed on a regular basis since it is embedded and employed in a dynamically complex world⁷³. The purpose is to determine to what extent targeted objectives are actually achieved, and to examine the potential violation of existing legislation. Adding to that, a specific external audit by a third party without compromising legitimate concerns regarding business confidentiality is possible and indicated. Should a negative impact be determined, the cause must be identified, and the program be adapted accordingly.

Algo.Rule 9: Establish complaint mechanisms

If an AI leads to questionable decisions or decisions that affect stakeholders' rights, it ought to be possible - to the extent that this is not already captured by existing regulations - to ask for and receive an explanation and to file a complaint by having easy access to feedback functions that e.g. allow for requesting appropriate and detailed information regarding a specific decision and the considerations that have fed into it. Complaints and actions taken should also be documented.

The open, participatory and interdisciplinary approach by the Bertelsmann Foundation to develop those 9 Algo.Rules involved more than 400 participants from

- science and research
- companies and professional associations
- civil society
- NGOs, politics and administrative bodies

who contributed their knowledge working on models for the practical application of each rule. To further increase their practical value, the next steps consist of refining the Algo.Rules for specific focus groups in (insurance and other) firms and the context of their application: AI developers, company executives, etc. (which we touched upon above already). Even though we would wish to conduct an empirical comparative analysis for the insurance sector to determine the firms' achievements so far with regard to each of the nine rules, i.e. to a transparent, holistic and responsible use of AI systems for their business activities, we have to admit upon critical reflection that such an approach is not feasible, at least currently. In our interviews with managers and executives of Swiss and German insurances, it turned out that the Algo.Rules are too broad, vague or generic to be linked to concrete processes or activities by the respective companies. For lack of specification (where in some cases it is not even clear how a specification can be attempted), interview partners would, for example,

⁷³ <https://onlinelibrary.wiley.com/doi/abs/10.1002/sres.2414>

simply agree to “overcoming silos & providing trainings” (#1). All in all, we made some observations however, which did lead to the following four hypotheses that ought to be tested by future studies.

- Insurance businesses have not yet sufficiently ensured the long-term successful and responsible use of AI technology. For some, ethics is not even a relevant topic.
- At first glance, the Algo.Rules attract much attention and agreement by insurance businesses. However, in terms of implementation within the organizations, they do not play a guiding role, at least in their current form. Some principles, e.g. complaint mechanisms (#9), do not seem to be necessary or relevant as additional guidance, since they are already covered by existing legislation. Others, e.g. crude labeling and intelligibility (#5 and 6), do not make much sense in an insurance sector to the extent that black boxes have been around before.
- Insurers understate the scope, power and status of AI in opposition to classical statistical approaches once it comes to its ethical implications.
- Focus on shaping ethical algorithmic systems ought to be complemented by a data perspective, and the latter is more a matter of individual responsibility than of the regulator’s business.

Furthermore, it is important to note that those justified claims in the form of the 9 Algo.Rules must not leave our interview partners, let alone their smaller competitors, behind and overwhelmed. A strong cultural aversion to take risks by investing into AI in insurance and risk management which might not meet the Algo.Rules requirements right away should not be indicated. On the one hand, as seen above, some Algo.Rules need to be restricted for insurances. On the other hand and generally, lethargy is not acceptable, but entrepreneurial courage is required, not just on the side of the startups. Despite all the snares on the way of establishing ethically acceptable AI systems which foster sustainable business success, the risks which are reflected by regulatory requirements for transparency, auditability, and completeness⁷⁴ go hand in hand with manifold opportunities, particularly once we focus on risk management in insurance.

To put it in a nutshell, AI in insurance risk management is able to improve the accuracy and consistency of models by offering deeper insights into data, by identifying complex patterns (such as sets of transactions indicative of invoice fraud) and making more accurate, faster predictions of default, and other risk events. In the wake of a combination of big data and smarter AI, insurers are empowered to better calculate risk on a more granular or even an individualized basis, rather than pooling risk across large groups of people or organizations. But since the latter strategy (conflicting with

⁷⁴ <https://mck.co/2gq3AWD>

the former) constitutes insurances' classical business model, a philosophical reflection about systemic effects of new technologies with their possibilities and limits becomes important.

Proposals such as the Algo.Rules thereby set reasonable boundaries for business opportunities, since society and regulators are likely to reject an insurance business model that employs algorithms and data in a radical, unbalanced, and short-sighted manner, e.g. to price insurance beyond the reach of people in need. Businesses that hesitate to consider their AIs as something that must be "raised" to maturity will remain struggling to catch up with new regulations and public demands – or worse, cause strict regulatory controls to be placed upon the entire industry.

Conclusion and take-aways

AI will transform the whole future value chain of insurances. Use cases and opportunities abound at all corners. To name just a few leverage points we reviewed, AI improves and expands distribution strategies or provides agility to insurers to enable them to develop products against new types of risk in new ways. In particular, AI fosters more powerful risk assessment systems gaining their capabilities from

- a more *granular* approach,
- *forward-looking* analytics,
- making use of both *internal and external data*.

New-generation risk management systems can indeed then become so powerful and disruptive that they are able to turn insurers' business model upside down: from pooling to personalizing risks.

We highlighted the currently long, resource-intensive, error-prone process of underwriting as a predestined case of applying AI in risk management. This process will massively benefit from AI-triggered automation, partially because AI technologies such as NLP are able to process relevant data which will increase tremendously in size over the next years. Secondly, those new technologies will empower their users to assess the increasingly complex risks around an insurance applicant often more precisely, but certainly much faster than humans.

This process will profit massively from AI-triggered automation, partially because relevant data will increase tremendously in size over the next years and risks around an insurance applicant are becoming increasingly complex. AI technologies like NLP are able to process these data and assess these risks, both often more precisely and certainly much faster than humans.

Advanced, sophisticated AI-powered risk models, however, have not been largely observed yet in the insurance market. AI adoption is low and slow due to methodological, data-centered and cultural issues.

- 1) In terms of data, several problems stand out. 1.1) While many applications of ML rely on masses of data, large, clean data sets are hardly available in insurances. 1.2) AI will benefit tremendously from a growing data source: (industrial) IoT, but insurers currently lack ways of identifying information in the breadth and depth of sensor data. 1.3) IT and data systems in insurances are heterogeneous and hardly interoperable. 1.4) The existing and growing heterogeneity of data e.g. stemming from new common sources of real-time data across the industry, turns out to be a veritable hurdle to AI adoption. 1.5) IoT enhances AI as an alternative data source with complementing

characteristics to existing data: IoT delivers data with high granularity in time, making it possible to monitor the state of insured assets or processes in real time. If integrated well with contractual and claims data, AI can draw from a much richer context around the insured subject compared to today. IoT data plays a role in risk management of existing insurance contracts (e.g. enabling parametric insurance), in the risk estimation for underwriting new contracts, and also in claims management (e.g. the last seconds of IoT data in a car's drive recorder before an accident).

- 2) With regard to methods, the devil is in the details. Whereas popular and hyped methods such as Deep Learning lack real impact in most insurance settings, 2.1) Bayesian modeling (to tackle small data problems) and 2.2) causal modeling (to efficiently select interventions to test putative causal relationships, and to make better decisions by leveraging knowledge of causal structure) have the power to pave the way forward. 2.3) IoT data are often time-series or event-series of sensor measurements. As such, time-series methods, among them recurrent networks, are suited when it comes to forecasting the dynamics of a risk-indicator based on a set of IoT data sources. Moreover, state-space methods may apply if the time-series are aggregated to a state of a monitored asset or process.
- 3) Culture-wise, 3.1) many companies talk about being "big-data-driven", yet the actual number of firms acting on their culture and organization is by far much smaller. 3.2) Digital cultures in organizations, if existing at all, are too hierarchical, lack cross-functionality and are organized in a bottom-up rather than top-down way. As a consequence, many data science projects fail because they often start by searching for the problem to be solved with a stipulated (AI-) method and not with the business problem that deserves attention the most. Data managers then make the mistake of starting working without being aligned to the overall business needs.

Insurers ought to resolve methodological, data-related and cultural issues, as this will boost the entire digital transformation journey of a company, not only the particular ability to integrate AI.

- Resolving methodological issues first necessitates the development of a clear understanding of the desired outcomes of specific AI applications. This development must be complemented by a careful assessment of which process is appropriate for a given ML technique, and of which data is truly required to deliver a working AI solution. Not least, it must be checked if the data points are continuously available, as well as in sufficient quality, at the insurance or the external sources it can access. The expected volume of data

determines if parameter-intense methods such as Deep Learning are applicable or if algorithms with lower need for data are favorable. A further ingredient for deciding which algorithm class to use in development, or which solution to buy on the market, is the necessary level of interpretability of the AI model. Compliance factors such as GDPR and ethical values are to be accounted for. Finally, the algorithm class to select evolves from these assessment steps and data scientists' work can continue by elaborating on the details of the respective AI solution.

- Resolving data issues goes hand in hand with a shift of mindset away from complex integrated solutions towards an agile orchestration of micro-services. Accordingly, insurers have to undertake small initial steps to create initial return-on-investment and subsequently fragment monolithic data systems. This will result in modular, interoperable data systems that make data sources accessible for any application, even if the concrete use case and application is still unknown. Instead of designing complex data processes, companies ought to focus on defining few and simple access rules for the company-internal data space and a secure but accessible interface to external data providers and aggregators.
- Resolving cultural issues emerges from sorting out data issues. Helpful initiatives aim at bottom-up enablement, democratizing data access to employees rather than top-down communication of digital values. Insurance companies should empower their employees to interact with data and might experience growing engagement with digital transformation, when exposure to data translates into successful projects. In particular, cultural change is key to change the business models from product sales to providing solution ecosystems to customers (e.g. “underwriting a commercial asset insurance” vs. “offering a risk management solution to clients that helps preventing losses”). The latter presupposes to spell out service-focused customer relationships that bring about a competitive advantage by seamlessly integrating into every situation where customers have insurance need.
- Ethical development of and responsibility for AI by insurances is crucial for long-term business success and thought leadership has recently progressed massively. Insurance companies ought to consider taking up elements from frameworks such as the Algo.Rules not just in their Code of Conduct in terms of acknowledging reasonable principles on a written declaration. But, corporations need to effectively find and implement answers to ethical questions on and challenges of AI raised by society and its institutions to keep their social license to operate. On top of that, a broader, more philosophical or systemic reflection on AI by businesses is indicated because, as argued above, its likely impact is of fundamental magnitude (e.g. in terms of individualizing in

lieu of pooling risks). For example, on the one hand, we could envision (and prepare for) a scenario according to which risk management of the future is on assessing the risks of machine error as machines take over decision-making along insurers' value chains. We could then also ask ourselves if machines are of crucial importance on that meta level too (i.e. not only for risk evaluation in underwriting, but also for assessing AI employed in underwriting processes). Or do we wish to preserve a human element⁷⁵ in the game? On the other hand, would we like to live in a (brave new) world where insurance companies turn out to be sufficiently powerful to dictate behavioral norms;⁷⁶ where their customers wish to influence their scoring in a positive manner to prevent that a possibly constantly monitoring insurance provider cancels the insurance policy since "reckless" behavior was recorded by some of the many sensors?⁷⁷

This White Paper set out to elaborate on questions about state-of-the-art of algorithm-based insurance and risk assessment methodology today and how this state may transform in the not too distant future. We found that insurance companies are quite heterogeneously positioned with regards to AI maturity, but observed as a common denominator that the entire industry is actively exploring the opportunities in their respective focus areas. Algorithm-based insurance builds on data, which is often not in place today. Productive examples are parametric insurances, e.g. in the area of logistics or agriculture, where conditions are triggered by sensor signals. Forward-looking analytics in such contracts or for underwriting is hardly used today.

Data logistics is an essential prerequisite for AI technology to be employed in underwriting and algorithm-based insurance schemes. In-house complex data systems must be converted into interoperable services, catering to data users (be they human users or software) in a flexible and adapting way. We emphasized that a changing value chain will position insurance companies in the center between customers and external providers of data and services, with both the need and the opportunity for insurers to add valuable differentiating services to customers.

Finally, we explored and evaluated a catalogue of rules to ensure that AI systems are activated in a morally responsible manner. The Algo.Rules are one example out of many frameworks that are currently developed to guide developers and users of AI

⁷⁵ <https://schweizermonat.ch/kuenstliche-panik/>

⁷⁶ Such a dystopia (without involving insurance companies) is also portrayed in Netflix's Black Mirror episode Nosedive.

⁷⁷ For instance, think of a car full of sensors which is reporting to the insurance firm that the driver is driving recklessly, having passed the speed limit by 50 km/h for the last 20 minutes. Wouldn't it then be almost realistic that the driver receives a warning on his display, informing him that his insurance coverage expires if he doesn't slow down by 50 km/h within the next 10 minutes?

with reasonable principles. Insurers now have the opportunity to position themselves vis-à-vis a range of options from a purely compliance-driven operating model to a reflective, inclusive model that actively shapes answers to the societal challenges ahead raised by AI.

About the authors



Dr. Christian Spindler is founder of DATA AHEAD ANALYTICS, a FinTech and IoT company, enabling Financial Services and Audit firms with industrial data. Christian served as Senior Manager for IoT & machine learning at PricewaterhouseCoopers and received extensive training in industrial IoT and processes as Senior Researcher at the global electro-technology company ABB. Christian graduated in physics (PhD) and earned an MBA from the University of St. Gallen.



Dr. Christian Hugo Hoffmann is an AI and FinTech entrepreneur and researcher who serves as co-founder of Syntherion, a B2B risk management software company, as well as senior scientist at the chair of Entrepreneurial Risks at ETH Zurich. Apart from that, he teaches FinTech in the executive program at the University of Liechtenstein and regularly writes pieces on AI in leading magazines and newspapers such as NZZ, Schweizer Monat and Towards Data Science.