Artificial intelligence in the insurance sector

An exploratory study

DeNederlandscheBank
EUROSYSTEEM

AFM

Summary

Developments in the field of artificial intelligence (AI) have gained tremendous momentum worldwide in recent years. The Dutch insurance sector has also shown an increasing eagerness to take advantage of the possibilities presented by AI. There are many opportunities, but they also give rise to uncertainties and risks for the sector. It is essential that insurers are fully aware of these uncertainties and risks. Only then will they be able to use AI responsibly and in accordance with the requirements regarding sound and ethical operational management, product development and duty of care.

Current AI applications

Dutch insurers are cautiously experimenting with both self-developed and externally acquired AI applications, often with the aim to automate or optimise existing processes and sub-processes. These AI applications can be classified as 'narrow' AI, which means that they focus on a specific task. Examples include applications for improving fraudulent claims detection, predicting customer questions, or providing better and faster damage estimates. In the years to come, a stronger focus is expected on applying AI in other areas, such as pricing (setting premiums) and customer acceptance.

The 'fuel' for AI models, the input data, mainly consist of data that insurers possess in-house. For certain processes, the internal data is enriched with external data sources, such as weather databases, data from the Dutch Central Information System Foundation (CIS) or data from the Chamber of Commerce.

Focus on technical aspects of AI

The application of AI techniques in the insurance sector is continuously developing. If insurers intend to use AI, it is essential that they are aware of the various technical aspects of AI models, especially when these techniques are not yet fully developed.

First and foremost, it is important that insurers, from the start, systematically define the restrictions in the use of AI and take its technical aspects into consideration. Knowledge of AI needs to be embedded within all levels of the organisation along with internal policies for its use. This must be anchored in clear governance structures. These are prerequisites for deploying AI responsibly and for triggering critical questions throughout the development and deployment stages.

More specifically, three technical components should be emphasised: the input data for the model, the model/technique as such and the model outcomes (and how to act upon them).

- The quality of a model (and this is especially true for AI applications) is largely influenced by the input data. It is crucial to focus on data quality and the suitability of the data for the intended AI application. This focus becomes even more vital in cases where insurers collaborate with external parties to enrich existing datasets or to develop AI applications. In this regard it should be noted that the insurer remains responsible for the data used and the AI application developed.

- Secondly, when developing an AI application, one can choose from multiple AI models and AI techniques. The question is whether and to what extent an insurer possesses the expertise to make such a choice. Such a decision process may include a comparison of the pros and cons of individual techniques and models, and how these pros and cons relate to the data available and the circumstances under which the AI application will be used. Another question includes the consistency in the use of AI techniques and models within the organisation. The same questions and considerations apply to externally developed and outsourced AI applications.

- Finally, models produce outcomes, which can result in action. This third component evokes new questions such as: to what extent can the outcomes of an AI application be traced back to individual input parameters? To what extent is such traceability required, taking into account the process in which the AI application is to be used? In addition, insurers need to consider how the process can be validated to ensure that an AI application continues to do what it was designed for. They also need to decide how frequently such validation procedures should take place.

Focus on consumer behaviour and social acceptance

The social context in which AI is deployed is equally as important as the emphasis on the technical aspects of AI. For example, it is undesirable for AI applications to encourage customers to act contrary to their own financial interests. This could be the case, for example, in online decision environments where AI has the potential to unreasonably influence expected heuristics and biases in consumer behaviour. Such insights should rather be used to encourage appropriate choices for and by customers. Moreover, the datasets (which are often very large) in combination with AI applications can reveal new patterns. Exploring whether these new patterns and outcomes withstand the test of social acceptance remains vital in key areas of insurance processes.

AI and solidarity: opportunities and risks

The adoption of AI by the insurance sector provides new opportunities for consumers who may be excluded from insurance options in the current system. However, the technology may also have a negative influence on the Dutch solidarity principle between groups of insured consumers. This dilemma is primarily a concern of the insurance sector, although both the AFM and DNB are willing to contribute to this discussion. It should be noted that not only the use of data and AI, but also other factors can potentially affect the solidarity principle.

Key considerations

The AFM and DNB have identified ten key considerations for the use of AI in the insurance sector. These considerations serve as a means to stimulate awareness among insurers and help to encourage a meaningful dialogue. Both the AFM and DNB consider these considerations as input for the further development of views and insights on this topic. These key considerations are based on discussions with insurers, stakeholders and experts.

Key considerations

Embedding AI in the organisation

1. How can insurers develop policies for the use of AI?
2. How can these policies best be embedded in the organisation?

Technical aspects of AI

3. What measures can guarantee the quality and completeness of the input data used for AI applications?
4. What is the best way to choose between specific AI techniques and models?
5. What degree of explainability is appropriate for AI applications?
6. What can be done to avoid illegal discrimination when using AI applications to identify causalities?
7. Which governance structures and criteria are appropriate for AI applications which are (partly) outsourced?
8. What is an appropriate method for validating AI models?

AI and the consumer

9. How can AI applications be prevented from taking unreasonable advantage of expected patterns or biases in consumer behaviour?
10. How can the outcomes of AI applications be assessed in terms of social acceptance?

1. Why this exploratory study?

In recent years, data has been the word on everyone's lips. Data is the 'new gold', and all signs are pointing to a modern-day gold rush. All kinds of organisations, including those in the financial sector, have more and more data at their disposal, and this data shows more and more variety. Thanks to the greater availability of data combined with massive gains in computer processing power, the development and application of artificial intelligence (AI) has gained considerable momentum in recent years.

The use of AI is expected to continue its growth in the years to come, certainly in the financial sector, and the Dutch insurance sector is no exception, where we also observe an increase in the use of AI applications. At the same time, there are greater societal concerns about consumer privacy and the way in which organisations and companies handle data. These concerns have also taken up a more prominent place on the legislative agenda.

A broad debate is also ongoing about the desired degree of solidarity between individuals and groups, and whether this solidarity will withstand more pervasive data analysis practices.

The rise of AI applications touches directly on various aspects of the insurance business. It not only offers new opportunities, but also gives rise to risks for the insurance sector.

Definition of artificial intelligence in this exploratory study

The concept of artificial intelligence (AI) can be defined in various ways. Additionally, many other related concepts are used interchangeably when referring to applications designed to analyse increasingly varied data points faster and more efficiently. These include concepts such as machine learning, data science, data analytics, advanced analytics and big data & artificial intelligence (BDAI).

This exploratory study uses the terms AI, AI models or AI applications to refer to applications that are based on analysing varied and large amounts of data using techniques such as machine learning.

The techniques used in AI applications (such as machine learning) are 'intelligent' in the sense that they are able to optimise rationally: for a given task they are able to choose the best action to achieve a specified goal in accordance with predefined criteria.

Alongside intelligent techniques, datasets are also an integral part of AI applications. The development of AI is being fuelled by the ever-larger volumes and diversity of available data. This exploratory study uses the term 'big data' to refer to such massive volumes of usable data.

Objective of the exploratory study

In this exploratory study, the AFM and DNB examine the developments, opportunities and risks associated with AI for the insurance sector. The study was conducted from the perspective of product development, duty of care and sound and ethical operational management. The objective of the exploratory study is to initiate a dialogue with the insurance sector and other stakeholders. The 10 key considerations that emerged can serve as a basis for such a dialogue. Now is a good time to initiate this dialogue, as the use of AI in the insurance sector is still in full development, meaning risks can be identified and mitigated at an early stage. These key considerations are explained in further detail in this study. The considerations should be interpreted as areas for a meaningful exchange of ideas in order to arrive at a clearer and more uniform approach to AI for the insurance sector.

Scope of the exploratory study

Many different terms and concepts are used to refer to AI, and they often have overlapping definitions. This exploratory study examines the use of AI in the insurance sector, where AI refers to applications and models for analysing varied and large amounts of data using techniques such as machine learning.

This study only examines AI applications in processes that are insurance-specific in nature, such as selecting, appraising and pricing risks, handling claims and detecting potentially fraudulent claims. It does not discuss more general applications of AI, such as in marketing or back-office processes, nor does it examine the potential impact of AI on the value chain or market structure of the insurance sector.

Reader's guide

The exploratory study first examines AI developments in the Dutch insurance sector: how is AI currently used in insurance-specific processes? The study discusses the new types of data that are being used, and it also looks at the use of techniques such as machine learning. Additionally, it sketches the expected development of AI applications in the insurance sector in the coming years.

Next, the study focuses on key considerations for the use of AI, e.g. the explainability and traceability of outcomes of AI applications. It also explains how the use of AI impacts governance, model validation, outsourcing, the product development process and duty of care.

Finally, the study discusses how AI affects the broader debate on solidarity in the insurance sector.

2. Developments

There is much discussion about AI applications and how they will change the insurance sector. The use of AI applications by insurers is, however, still in full development. Therefore, the first questions to ask are: Which applications are already being used by Dutch insurers? And which developments can be expected in the next few years?

This section first discusses the types of AI techniques and data that insurers currently use, and their expectations about how these techniques and data will evolve in the years to come. Next, the discussion focuses on value chain components where AI is currently used, and how this will evolve in the near future.

2.1 Use of AI techniques

Which AI techniques are being used by Dutch insurers?

In recent years, an increasing number of new techniques have become available to analyse (large amounts of) varied data. One of the most important techniques for the insurance sector is machine learning, where a system is trained to perform a specific task over and over while optimising the process each time. This training is not the result of explicit instructions. Rather, it is based on optimisation algorithms. We also see the emergence of AI applications that are based on 'deep learning' techniques, i.e. more complex machine learning techniques that use multiple layers of analysis (see Figure 1).

Various Dutch insurers already use different types of machine learning applications in their processes. Some of the most commonly used techniques are clustering, random forests, gradient boosting and deep neural networks. The use of machine learning applications is often still on an ad hoc basis, usually with the objective to support or challenge more traditional models.

One specific application of machine learning, Natural Language Processing (NLP), involves the use of algorithms to understand and interpret textual data. Various Dutch insurers already use NLP techniques, though primarily for back-office tasks such as sorting and allocating e-mails or post, or in customer contact through virtual assistants.

In general, Dutch insurers do not yet use machine learning techniques on a large scale or in a structured way for their primary insurance processes. For these types of tasks and processes, the focus for now is primarily on expanding existing statistical (regression) models, for example by adding new data and parameters to these models (also see Section 2.2 Use of data).

Figure 1 Interrelation between AI concepts

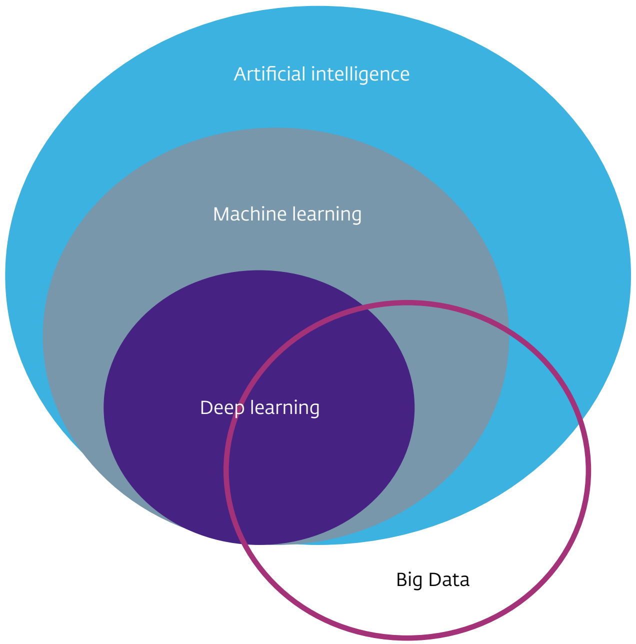

2.1.1 Internally built vs. externally acquired models

At present, Dutch insurers use a mix of self-developed and externally acquired (off-the-shelf) models and platforms for their AI applications. However, there is a trend towards building models in-house, especially as insurers develop more and more expertise in this field.

2.1.2 Ways to train machine learning models

Machine learning models are often trained through *supervised or unsupervised learning*. A technique like *reinforcement learning* – where the model adjusts itself based on a result or an outcome – is far less used. The same applies to models that are continuously updated based on newly available data. Insurers do retrain their models periodically (on a weekly or monthly basis) with newly available data.

2.2 Use of data

Which types of data do Dutch insurers use?

The significance and impact of 'big data' are often stressed in relation to AI applications. Insurers have large volumes and a great variety of data at their disposal, but even these volumes do not come close to the amount of big data generated by processes at Big Techs such as Tencent, Google, Alibaba or Amazon.

In recent years, insurers have therefore focused primarily on improving their internally available data (Figure 2), of both a structured (databases) and unstructured (texts and scans) nature. Insurers consider their own internal data to be of the greatest value, especially when the data is directly related to customers or can be put to immediate use in insurance processes, e.g. data on claim behaviour.

In addition to internal data, insurers are also trying to incorporate external databases. These databases may contain information on payment habits or creditworthiness, but they may also include data from the Chamber of Commerce, the police, the Employee Insurance Administration Agency (UWV) or the Dutch Central Information System (CIS). In addition, and especially in traditional insurance segments (vehicle, home, agriculture), data from weather databases or satellites (geocoding, soil data) are used for improved risk appraisal and mitigation. Especially a combination of internal and external data is considered to be very valuable.

The use of data derived from the Internet-of-Things (IoT) is increasing, but in most cases this is still in an experimental phase. This data mainly involves vehicle telematics, with a focus on driving behaviour and driving habits (routes, times, etc.). Large-scale commercial exploitation of IoT data is not yet taking place in the Netherlands or most other countries. The United Kingdom is an exception, where IoT data is being used, for example in vehicle insurance products for high risk groups (e.g. young drivers with little experience on the road).

Figure 2 Use of new data sources
(% of participating Dutch insurers)

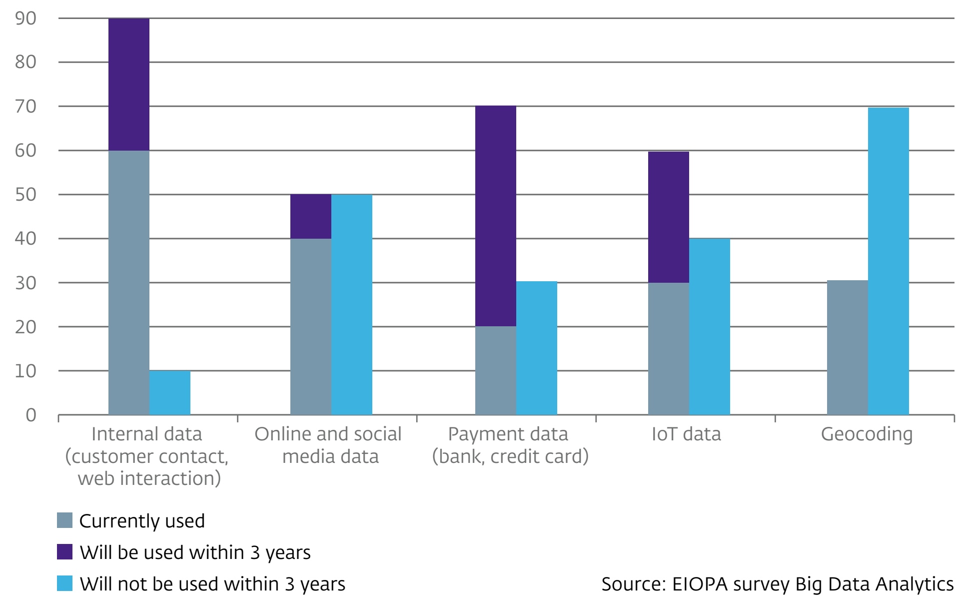

Source: EIOPA survey Big Data Analytics

There are several reasons for the limited enthusiasm to use IoT data in the Dutch insurance sector. Consumers are not keen to share such data due to privacy concerns. Additionally, it is questionable whether IoT data can be sufficiently contextualised to provide reliable input for analysing behaviour and related risks. The insurance sector itself is also unsure whether the available IoT data reliably reflects causal relationships.

Finally, the use of data from social media platforms in insurance processes is still very limited. Some insurers have conducted trials, but the results show that the added value of the data is still insignificant.

Furthermore, the privacy risks associated with the use of this type of data are considered significant. A number of insurers are nevertheless looking into the possibility of using online data, such as data on consumers' activities on the insurer's website.

2.3 Use of AI in the value chain

Where is AI being used in the value chain?

AI applications can be used in various insurance processes, from product development, risk selection, pricing and customer acceptance, to damage estimation, claims management and fraud detection. The European supervisory authority for the insurance industry, EIOPA, recently conducted an extensive study (including among Dutch insurers) into the use of AI. From this study – and from discussions with insurers – one can conclude that in the non-life and income segment Dutch insurers expect the greatest opportunities and impact from AI in the coming years to be in the areas of premiums, customer acceptance, claims management and fraud prevention (see Figure 3). For Dutch health insurers, the impact of AI in the area of premiums and customer acceptance will likely be less significant than for other segments, because health insurers are not allowed to refuse customers (at least for basic health insurance), and premium differentiation is not permitted.

Figure 3 Impact of AI on the insurer's operations
(0 = low impact, 5 = high impact)

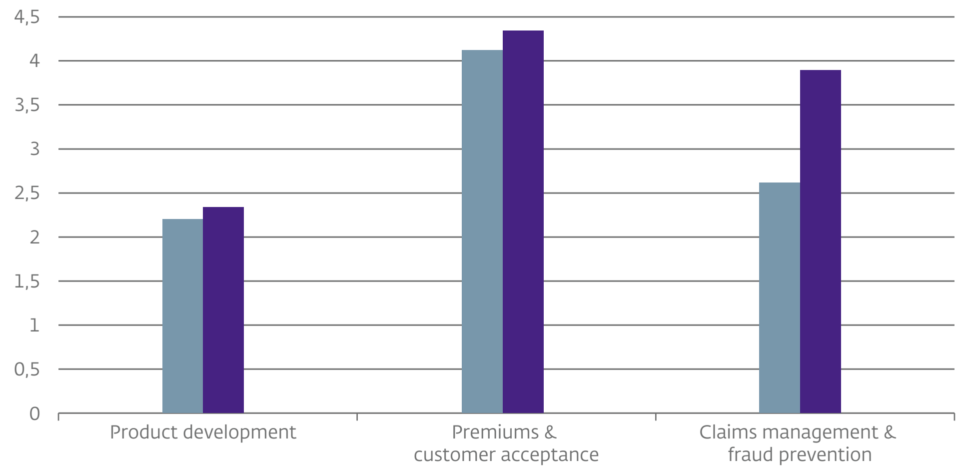

Source: EIOPA survey Big Data Analytics

2.3.1 Claims management & fraud prevention

Fraud

In combating fraud, new data sources (e.g. photos or prices of car parts) and machine learning techniques, especially anomaly detection (see Figure 4), are already being deployed on a larger scale.

Claim scoring, in which a claim is assessed based on its characteristics, is also frequently used. If the score exceeds a certain threshold value, it will be flagged as anomalous or potentially fraudulent. Such claims are then usually assessed by a claims expert. Fully automated processes, where models make autonomous decisions about claims, are currently not being used. However, insurers are endeavouring to increase the number of claims that are scored automatically. Insurers take various aspects into account: the costs of allowing false negatives (approval of claims that should have been rejected) due to automated handling, the costs of manually checking potentially fraudulent claims to avoid false positives (rejection of claims that should have been approved), as well as reputation risks associated with false positives.

Insurers also consider social network analysis as a promising AI application. This involves analysing an individual's social network for indications of potential fraud when a claim is submitted. Behavioural modelling (analysing an individual's conduct and behaviour) is another area of interest for insurers. It is expected that such analyses will be used with greater frequency in the near future.

Figure 4 Use of AI applications in combatting fraud
(% of participating Dutch insurers)

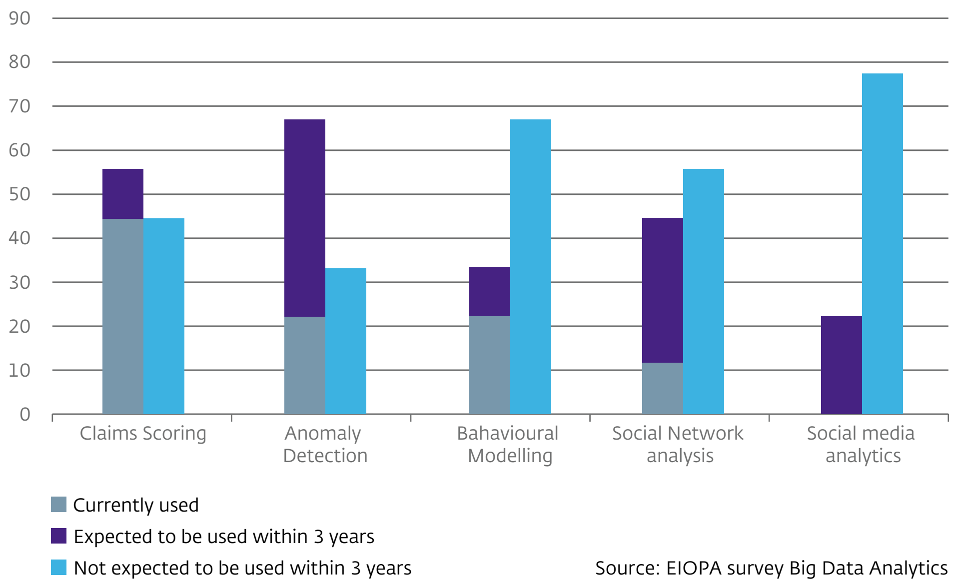

Source: EIOPA survey Big Data Analytics

Conversely, automated analysis of social media is not yet seen as a very useful or applicable anti-fraud instrument, because insurers consider the added value of social media analytics to be relatively limited and the associated privacy risks to be relatively high.

Claim management

Apart from fraud detection, AI is being used sporadically to handle, manage and classify claims more accurately or to predict claim characteristics. External databases are often used for this purpose. These databases contain data on, for example, prices of car repairs or car parts price lists. A few insurers are also using machine learning techniques (e.g. for photo recognition or for analysing IoT/sensor data) in order to estimate insured losses faster and with more accuracy. Nevertheless, AI applications are currently not widely used in the field of claim management.

2.3.2 Pricing & customer acceptance

Risk pricing

AI applications have already been used in recent years, albeit to a limited extent, for assessing and pricing risks. This has been the case, for example, in the non-life segment (vehicle insurance), and to a lesser extent in the income protection segment.

As indicated earlier, these developments are expected to continue in the years to come. Various insurers are already using additional data to refine their risk assessments. This additional data consists of internally available data (often on existing customers), publicly available data and externally sourced data. Sometimes, the additional data is used to reinforce existing risk factors, and sometimes it is used to introduce new risk factors in models or to create more sophisticated risk segments (see Figure 5). These new parameters are usually of a fairly traditional nature, meaning that they have an intuitive causal link with the claim.

The number of risk factors and risk pools applied when pricing have increased both for Dutch vehicle insurance and for occupational disability insurance. Expectations are that this trend will continue in the coming years; insurers expect that the number of risk factors used in models will grow, because of the increase in the availability of data and data analysis techniques.

Figure 5 Number of risk factors and risk pools are increasing

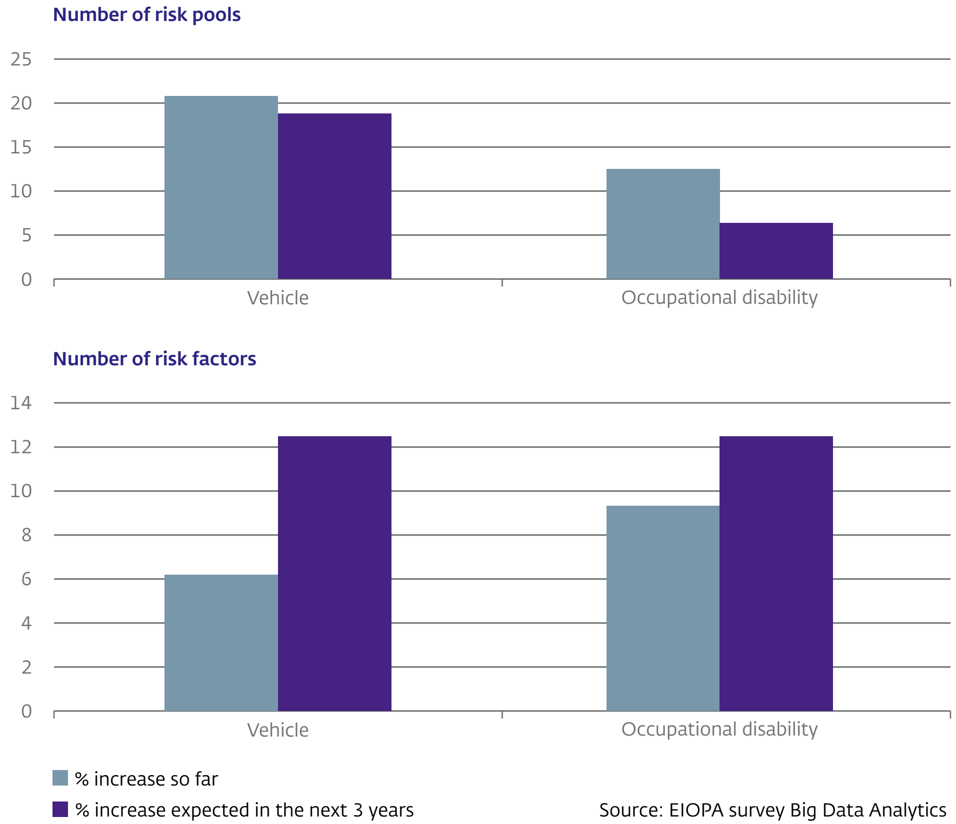

Source: EIOPA survey Big Data Analytics

Behavioural pricing (dynamic pricing)
AI applications are being used for behavioural pricing; more specifically the premium component that is determined based on market conditions and behavioural aspects. This is also referred to as dynamic pricing or pricing optimisation. Behavioural pricing is often based on the loyalty (elasticity) of customers, their lifetime value or the premiums charged by competitors.

Such pricing strategies are not a new phenomenon. However, insurers do make more use of new data – often a combination of 'new' internal data (e.g. for customer value or loyalty, or by analysing click and conversion information from the website) and external databases – in order to be able to focus more on dynamic pricing. Pricing based on very specific and individual characteristics of potential and current customers is not (yet) being applied.

2.3.3 Technical provisions
Little or no use is being made of new techniques such as machine learning for estimating expected insured losses (technical provisions); more traditional statistical models are generally used for this purpose. A few insurers have indicated an interest in new techniques, in particular to challenge current models.

2.3.4 Capital models
Currently, the Dutch insurance sector does not use AI in models for determining capital requirements or allocation of capital. It is debatable to what extent machine learning techniques can have any significant added value for such models. It is of great importance for these models to approximate a 1:200 shock scenario (the probability used for setting the legal capital requirement). This requires calculating a probability distribution, whereas machine learning techniques currently focus more on reaching a best estimate.

2.3.5 Customer experience & risk selection

Many insurers have put significant effort into digitising and improving the customer experience. Examples include anticipating customer questions by analysing website interaction, text mining or NLP (mails, call centres), improved website navigation, and even personalised website content. Such applications can also be part of the risk selection process.

3. Key considerations

Apart from opportunities, AI also presents challenges in a number of areas. Insurers themselves believe that the primary risks lie in dealing with consumers: maintaining consumer trust, using AI ethically and identifying potential risks that jeopardise the reputation of an insurer and the sector. In addition, insurers also see regulation and the supervisory authority as potential sources of risk: it is unclear for insurers how AI can be used without overstepping the legal and regulatory requirements (see Figure 6).

Figure 6 Challenges for a responsible implementation of AI applications

(Assessment of the challenge by insurers, 0 = no challenge; 4 = very big challenge)

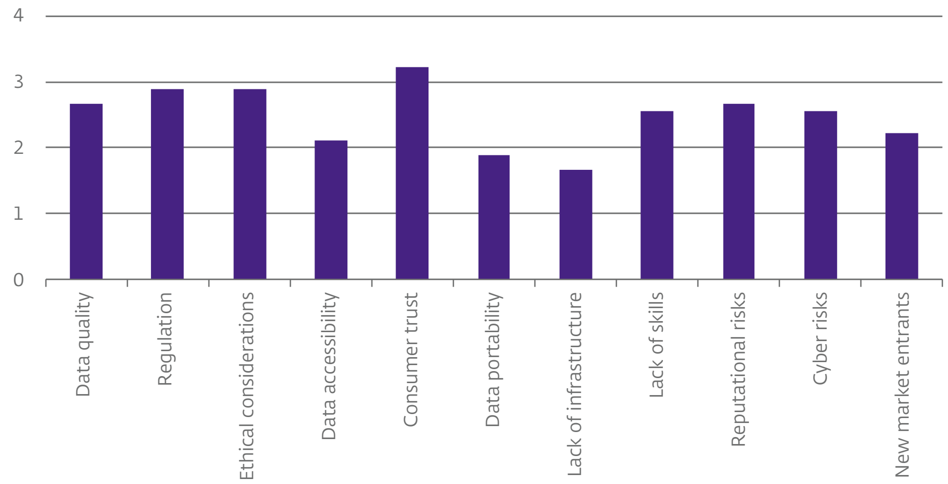

Source: EIOPA survey Big Data Analytics

This section discusses 10 key considerations regarding the application of AI by insurers that merit further scrutiny and dialogue. These key considerations are subdivided into 3 categories: embedding of AI in the organisation, technical aspects of AI, and AI and the consumer.

3.1 Embedding AI in the organisation
How can AI be embedded in the governance structure and policy of insurers?

3.2 Technical aspects of AI
What are the key considerations regarding the development and application of AI?

3.3 AI and the consumer
What are the key considerations regarding the duty of care when applying AI?

3.1 Embedding AI in the organisation

How can AI be embedded in the governance structure and policy of insurers?

There is a risk that insurers may not have a fundamental vision on how to apply AI. Artificial intelligence is all over the news and on everyone's lips, meaning insurers may regard AI as a hype and fail to clearly define whether and how AI may be relevant to them. This is a risk in and of itself, but there is also an opposing risk: insurers may start using AI driven by a fear of missing out, but without a clear underlying strategy. If data analysts are not provided guidance, they may design AI applications that are not in line with the insurer's risk appetite or strategy. Underwriting risks may be the result: like bad choices when it comes to customer acceptance, pricing or other aspects. The insurer also runs the risk of reputational damage. This may occur for instance if customers are subjected to a data analysis that deviates from the insurer's own standards.

With these risks in mind, a clear policy must be developed at board level regarding the deployment of AI applications (see key consideration 1). The Solvency II Directive [1] requires that insurers assign responsibilities to competent persons and that (risk) strategy, goals, risk appetite and risk limits, processes and roles are defined. A clear risk policy on the use of AI is also part of this requirement.

Key consideration 1 - Determining policy for AI applications

An insurer's board should carefully consider whether and how AI should be used. Questions such as the following should be asked in this regard:

- What criteria serve as a basis for deciding whether to use AI (machine learning, big data)?
- For which processes and components in the chain does the insurer intend to use AI?
- What rules does the insurer apply for training and retraining models? How often does the insurer want to retrain its models? How does the insurer structure the processes related to training and retraining?
- In the field of ethics and social accountability:
  - How much differentiation does the insurer consider justified, both with regard to risk assessments and price optimisation (dynamic pricing, behavioural pricing)?
  - What type of input data does the insurer intend to use for differentiation?
  - To what extent does the insurer intend to use AI to enable customers to improve their risk profile (risk prevention in healthcare or in the home, or behind the wheel)?
  - Is the insurer considering offering a premium discount in exchange for submitting data?
- How does the insurer assign decision-making responsibilities, processes and roles within the defined policy frameworks?
- How does the insurer guarantee that the responsibility for AI applications has been clearly assigned within the board, and that the responsible board member is sufficiently knowledgeable and experienced to estimate, test and manage the risks of AI applications?

1 Articles 258 and 259 of the Solvency II Delegated Regulation

Insufficient knowledge sharing and testing can jeopardise the proper implementation of AI within the organisation. Formulating a policy that specifies how the insurer intends to use AI is a start. If policy awareness is lacking in the organisation and relevant departments (e.g. data science, actuaries, risk management, IT) fail to communicate and challenge one another, then AI deployment may not be in line with the insurer's risk appetite. A potential consequence could be that customers and insurers will be exposed to unforeseen underwriting and operational risks. Such risks are increased when AI expertise is fragmented throughout the organisation. This may be the case when individual departments deploy their own AI solutions without centralised internal supervision by a data science team, for example. Clear governance structures must be in place to safeguard AI expertise in the organisation. This will ensure that models (and their outcomes) can be carefully validated while also guaranteeing continuity. This can be accomplished in various ways. For example by centralising data science expertise in a specific team, or by taking a structured approach in sharing such expertise throughout the organisation.

Under Solvency II, the concept of challenging each other has been incorporated in the roles of key functions, in particular the risk management function (model validation) and the actuarial function (methods used, quality of data). The internal audit function in turn is responsible for independently auditing the other key functions.[2] These functions can also assume their respective roles when it comes to AI applications. In view of the technical complexity inherent to AI, the relatively high frequency with which AI models are adjusted and the fact that the use of AI by insurers is still a relatively new phenomenon, it is advisable that challenge discussions start at the level in the organisation where AI applications are designed and deployed.

Key consideration 2 - Communication about AI applications and embedding them

Clear internal communication about the vision/policy of the board with regard to AI applications provides guidance for the use of AI. In addition, the organisation can build its AI expertise by enabling structured sharing of knowledge and experiences. The importance of mitigating the risks of AI applications can be emphasised if the responsible board members and relevant key functionaries engage and challenge each other.

2 Articles 269, 271, 272 of the Solvency II Delegated Regulation

3.2 Technical aspects of AI

What are the key considerations regarding the development and application of AI?

3.2.1 Prior to training: data quality and model selection

The availability of high-quality and varied input data is a precondition for applying AI. Calibrating AI applications, especially machine learning models, is entirely based on the optimal analysis of patterns inherent in data. Factually correct and representative data are therefore crucial for calibrating machine learning applications. Additionally, when models become more complicated and less transparent, it will be more difficult to trace and resolve errors or biases in the data through sanity checks.

Practical experience suggests that, despite the work that insurers have done in recent years to improve the quality and accessibility of data, the volume, completeness and quality of data sometimes still falls considerably short of expectations and is too sub-standard to allow deployment of machine learning on a broad scale. In this regard, data on individual customers deserves extra scrutiny: when for instance intermediaries are used, such information is now often not readily available, or incorrect.

Lack of sufficient, correct, complete or varied data may give rise to underwriting risks. Such 'dirty' data may lead to incorrect patterns in the model. Furthermore, and particularly in the case of machine learning models - where the calibration and structure of the model are determined by the input data - patterns (risks) not included in the data will not be recognised by the model either. This risk is augmented as the data becomes less diverse. In addition, erroneous data may also pose risks for insured persons if such data leads to incorrect decisions, e.g. when it comes to setting premiums, accepting customers or selecting an appropriate insurance product.

Key consideration 3 - Being in control of input data

Before using AI, it is important that the insurer can demonstrate that it is in control of the relevant input data. The following questions should be emphasised:

- Does the insurer have an up-to-date overview of the data elements to be used?
- Have data quality standards been drawn up for the input data to be used?
- Are controls in place to monitor the quality of the input data on an ongoing basis?
- Has a risk assessment been conducted of the quality and completeness of the input data?
- Are any shortcomings in the data remediated appropriately?
- Does the input data satisfy the data quality standards set by the insurer?

In addition to the absence of material errors in the input data to be used, it is important that the datasets used are cleared of unwanted biases and assumptions to the greatest extent possible.

Effective risk management requires being in control of the 'sanitary' quality of the data: in other words, the datasets must be free of material errors or missing data points. To ensure that this is in fact the case, insurers could implement the processes described in DNB's Guidance on data quality for Solvency II reporting processes[3] in their AI applications too. The key functions (especially the Actuarial function) also have an important role to play when it concerns challenging the quality of data used in AI applications.

When assessing the suitability of input data for use in machine learning models, however, more is required than just checking whether the data are factually correct and complete. It is important that the data is sufficiently diverse and representative for the purpose for which it will be used. This includes questions such as: is the incidence of outliers in the dataset acceptable? Which input parameters are included in the dataset and for what reason? And which parameters have been left out? Representativeness can be enhanced by making use of scenario analyses and data simulations. Biases and assumptions in the data should also be scrutinised (see key consideration 6).

When deciding to use specific AI applications, insurers should take into account the statistical risks and the possible complexity of the applications, especially when it concerns machine learning models. Statistical risks arise when, for example, an insurer decides to use a machine learning model for which it possesses insufficiently diverse data. Even if the data is correct and complete, the risk of an overfit may increase, causing the model to produce erroneous results for insured persons with characteristics that differ from those in the training set applied. Another risk concerns the complexity of a model: if an insurer uses models that make it difficult to analyse patterns intuitively, then the insurer may run the risk of failing to grasp the outcomes and patterns, and may also lose control of the models. In the worst case scenario, this may result in wrong decisions, and could also lead to discrimination. Especially in processes that have a direct impact on insured persons, it is essential to take complexity into account when selecting a model. It may even be necessary to make an explicit trade-off between the effectiveness and complexity of models.

3 See DNB (2017) "Guidance Solvency II data quality management by insurers", www.toezicht.dnb.nl/binaries/50-236703.pdf

Key consideration 4 - Model selection

When selecting an AI/machine learning model, the following questions should be considered:

- Can the insurer systematically substantiate why a certain model and technology has been chosen?
- Is the decision for a certain model based on the quantity, quality and diversity of the available input data?
- When a certain model or technology was chosen, were factors such as explainability, complexity, and reliability taken into account alongside 'best fit' considerations?
- Is there a certain degree of consistency between the models and technologies used for determining premiums and those used for determining technical provisions?
- Were experts from the relevant business areas, e.g. the IT, Actuarial and Risk management (model validation) functions, involved in the selection process?
- Can the insurer give insight how the chosen technology works in a more general sense, and for which types of processes or types of datasets one specific technology is more suitable than others?
- Can the insurer describe circumstances under which the use of the chosen technology would no longer be appropriate? Is this checked periodically and, if so, how?

3.2.2 Dealing with outcomes

The explainability of the outcomes of AI applications is of great importance, especially when used for sensitive processes. Explainability first of all means that an insurer is able to indicate how the input data leads to a certain outcome. This means indicating which parameters have contributed significantly to the outcome and quantifying this contribution. Explainability also means being able to indicate which changes in individual input values are necessary to enforce a change in the outcome of the model.

For some machine learning technologies, it is very difficult or even impossible to achieve such a degree of explainability. Models built by means of such technologies are generally referred to as black box models. Examples include deep neural networks and, to a lesser extent and depending on how the model is constructed, random forests. Although there are various options to reveal patterns in black box models (for example, through explainer models or partial dependence analysis), such observations often only apply to one particular combination of input and outcome. This is an important difference compared to white box models, where, in principle, the patterns between input values and outcomes – as well as the statistical reliability of these patterns – can be observed for all input patterns.

The dilemma of explainability and black box models became more pressing after the introduction of the European General Data Protection Regulation (GDPR), which gives consumers greater rights when they are subjected to automated decision-making. In this case, insurers must be able to explain to consumers the underlying logic of the models used and the consequences of automated decision-making for the consumer. The consumer is also entitled to object to the decision taken.[4]

4 Articles 15(1)(h) and 21 of the GDPR

Key consideration 5 - Explainability of outcomes

When determining to what extent a model or machine learning technology is appropriate for a process with automated decision-making, the following questions should be considered:

- To what extent are models used in processes that have a direct and large impact on customers - and thus possibly involve risks for the insurer with regard to product development, duty of care, legal and reputation risks - or in processes that have a direct and large impact on the insurer's stability?
- To what extent can the model or machine learning technology be explained? In other words: to what extent is it possible to trace back patterns between input parameters and model outcomes?
- What degree of explainability is appropriate for that specific process?

Explainability can best be considered on a more case-by-case basis, because the degree of explainability is not the same for every machine learning technology. Moreover, not every process requires the same degree of explainability. Black box models do not necessarily have to be banned, but the degree of explainability that is required for a process must be carefully assessed, along with how the desired explainability can be achieved. A few examples are discussed in more detail below:

Back-office processes: such processes may include post or email sorting, or optimising the use of call centres. The impact on the customer – also with regard to discriminatory patterns – is often limited in such processes. In this case, using less explainable technologies and models has a small impact on the customer and on the insurer's stability. Any defects in a model's 'fitness for purpose' can be traced and addressed based on output checks. That is why the use of black box models in such processes is more acceptable than in other processes.

Models for pricing, customer acceptance and fraud detection: these types of models have a direct impact on the customer, especially when they are used for automated decision-making. There are, in principle, significant concerns with regard to product development, duty of care, explainability and discrimination, and the associated reputation and legal risks for the insurer. It is therefore important to be able to indicate whether and to what extent individual input parameters contribute to the outcome of the model, and which changes to the input parameters are required to result in a change in the model's outcome. A black box model, which does not allow for such an explanation (or only with great difficulty), will therefore probably not be suitable for a process designed for automated decision-making that directly affects the customer.

Black box technologies can, however, be used to support decision-making of other models or, for example, as support for fraud investigators or first-line actuaries. In this regard, an insurer should set an acceptable threshold for false positives and false negatives. Moreover, the insurer should see to it that human intervention does not become a mere formality.

Models for calculating technical provisions or capital requirements: in this case, concerns about discrimination or directly and unjustifiably placing consumers at a disadvantage are not as big as with pricing or customer acceptance. In principle, a smaller degree of model explainability would be acceptable compared to, for example, processes for pricing or customer acceptance. From a prudential point of view it is essential, however, that the insurer is well aware of the statistical reliability and the uncertainty margin of the outcomes. The use of a full 'black box' for automated decision-making in such processes would therefore be undesirable in most cases. It is of the utmost importance to be able to demonstrate that the model outcomes provide correct estimates of the required provisions. This can in part be achieved in the validation process, for example through back-testing and testing model predictions against actual insured losses. However, a certain degree of explainability must be maintained, also to allow effective challenges by the actuarial function. This is of greater importance for product segments with longer timelines, e.g. life insurance products, because in these segments it often takes considerably longer before the actual insured loss is known.

If discriminatory biases in AI applications cannot effectively be avoided, the insurer should consider not deploying these applications. Discrimination means the unlawful treatment, subordination or exclusion of people on the basis of personal or other characteristics. When insurers perform analyses, for example in the customer acceptance process, pricing or possibly as part of fraud detection activities, there is a risk that the patterns and parameters that are part of AI applications may result in discriminatory decisions.

The law sets the grounds on which insurers are not allowed to discriminate. Examples of such grounds are gender, ethnicity, religion, sexual orientation and disability.[5] Direct use of such parameters is relatively easy to avoid, but indirect use of discriminatory variables (or their proxies) must also be avoided.

This risk is not new. After all, less extensive or complicated models can also contain discriminatory proxies. However, in AI applications – and especially in less explainable machine learning technologies – there is a higher risk that patterns arise which are potentially discriminatory and more difficult to trace.

Key consideration 6 - Avoiding the use of patterns that lead to illegal discrimination

An insurer must have systems and processes in place that prevent AI applications from generating discriminatory outcomes.

The following questions need to be addressed when designing such systems and processes:

- How are input variables challenged to detect possible discriminatory bias?
- How are outcomes checked for discriminatory bias? This may include the use of adversarial modelling, as well as sample testing with identical test groups, where a discriminating (proxy) variable is the only difference between groups. If significant differences emerge from the model, then discriminatory bias may be present. Virtual cases may also be used for this purpose.
- How can checks for discriminatory bias be refined and made more robust? One possibility would be specific checks for biases in false positive outcomes (rather than restricting tests to overall model outcomes).

If unlawful discriminatory biases cannot be ruled out, then the insurer should consider whether it is prudent to use the model in processes that directly affect customers (e.g. pricing and customer acceptance, fraud detection).

5 Section 5 of the Dutch Equal Treatment Act (AWGB) and Article 9(1) of the GDPR

It is important that insurers put a process in place that allows them to continually challenge the input parameters and any discriminatory patterns that may be present in models. In view of the potential consequences that discriminatory biases may have for customers – and the associated legal and reputation risks for an insurer – an insurer should only use a model if it can be established with sufficient certainty that the model will not generate any prohibited discriminatory outcomes.

If an insurer nevertheless wishes to use a proxy for suspicious or unlawful discrimination, then it must substantiate such use in accordance with the usual legal tests: for example, an objective goal must justify the discrimination and the insurer must be able to demonstrate that the discrimination is proportional to the goal it hopes to achieve. The variable/proxy would have to be sufficiently delimited as well. Machine learning models with less explainability make it especially difficult to trace the precise patterns that are used in the model, presenting yet another legal obstacle.

Finally, it is important that insurers focus their bias analysis on the goal for which the model is deployed. For example, when designing technical provision processes, certain distinctions (e.g. between male and female) may be less problematic, whereas such a distinction may be highly undesirable and potentially unlawful in the case of pricing processes.

3.2.3 Use of external data and models

When AI applications are outsourced, it is essential to monitor, test and challenge these outsourced applications. Insurers make extensive use of external expertise in the form of data or algorithms/models. This usually involves acquiring external databases, working with external consultants to train machine learning models, or acquiring pre-trained models.

Without proper outsourcing processes, insurers may run the risk of acquiring datasets of dubious quality, or failing to understand how external models have been trained or function. Outsourcing key processes is especially prone to increased underwriting and operational risks.

With the increasing use of external data and models, it is of importance that insurers look closely at how they collaborate with external parties with regard to data and data analysis. This should be taken into account in the outsourcing policy: what kind of activities does the insurer want to outsource, and what activities should be kept in-house? Which external data sources and applications do insurers want to use and for what purpose?[6]

6 See also DNB (2018), "Good practice document for outsourcing by insurance companies", https://www.toezicht.dnb.nl/en/2/51-237170.jsp

Key consideration 7 - Outsourcing of AI applications

It is essential that insurers monitor their AI applications, regardless of whether they have been developed in-house or outsourced to an external party.

Insurers should compare their outsourcing and partnerships with external parties (e.g. regarding data and models) to DNB's Good practice document for outsourcing by insurance companies (2018). The Guidance on checking Solvency II data quality by insurers (2017) can then be applied to relevant processes.

Insurers must have an outsourcing policy in place, and they need to determine - on the basis of the criteria discussed in DNB's good practice document - which outsourced AI applications are critical. Insurers must also have a process in place for monitoring outsourced processes. The following points should be emphasised:

- Does the insurer possess sufficient expertise to understand how the external application works?
- Have agreements been made with external parties regarding the quality and origin of the data provided, and on how the external models have been trained/calibrated?

If the questions above are not adequately addressed with regard to a specific external party, then the insurer should potentially reconsider the relationship.

Insurers could also take steps individually or, where appropriate, collectively, to prevent outsourcing of AI applications that undermine continuity. Among other aspects, this involves reaching agreements with external parties about the availability of applications and data, as well as preventing excessive dependence on one or more external parties.

Under Solvency II, additional criteria apply for critical outsourcing processes. To determine which outsourced AI activities are critical, the insurer should ask itself the questions set out in DNB's Good practice document for outsourcing by insurance companies:

- Is the activity inherently critical for insurers, e.g. is the activity critical for meeting obligations to policyholders?
- What are the operational effects (reputation, legal) in the event of interruptions?
- What impact could disruptions in outsourcing have on the insurers' income?
- What impact would an outsourcing-related breach of confidentiality have on insurers and their policyholders?

Under Solvency II, insurers must set outsourcing agreements with external parties[7]. This also applies to the outsourcing of AI applications. In addition to the legal obligations, these agreements must also set out the expectations with regard to the quality of the AI applications to be supplied. When for example data collection is outsourced a data delivery contract must be created, which contains agreements about the expected data quality.[8]

3.2.4 Validation

It is important that validation processes are structured in such a way that it can be determined whether AI applications are fit for purpose, even when the applications are frequently or continuously updated or retrained. It is crucial for insurers to formally validate models in order to determine whether they actually do what they were designed for, i.e. if they are fit for purpose. Insufficient validation procedures may create the risk of a model that is no longer fit for purpose. It could

7 Article 274 of the Solvency II Delegated Regulation
8 See also DNB (2017) "Guidance Solvency II data quality management by insurers" www.toezicht.dnb.nl/binaries/50-236703.pdf

even be that the insurer no longer has suitable criteria for assessing whether the model is fit for purpose.

It should be noted that all models are subject to such a risk. However, in the case of machine learning models (that depend on the input data with which they are trained for their calibration) periodic retraining with newly available data is desirable. Such retraining can contribute to keeping the models fit for purpose, but it also raises questions. Retraining can cause the model to change considerably. Such changes can have such a substantial effect that the model may in fact be regarded as a 'new' model, which should be subjected to formal validation. The risk is that validation will not keep up with developments in the model, causing the insurer to lose control.

Here, however, a trade-off must be made between staying in control of the model on the one hand, and the practical feasibility of revalidations on the other. It is important that insurers are aware of this and define what 'major' and 'minor' changes to the model mean. Insurers should determine the minimum frequency for revalidations.

Several questions arise when self-learning algorithms/models are used. In principle, these models are constantly updated. This means that periodic validations are of little to no use. Validation should instead focus on the process through which a model is continuously adapted. Sanity checks on the outcomes of the model are part of this process, focussing on questions as to whether the outcomes are plausible. An option would be to set 'crash barriers', i.e. outcomes that fall outside a predetermined bandwidth. This would automatically generate further manual checks and possible adjustments to the model.

Key consideration 8 - Validation

It is crucial to establish a validation procedure for AI algorithms/models. Answering the questions below may contribute towards developing such a procedure:

- How important/critical is the model in terms of impact on the customer and the stability of the insurer, and does the validation procedure focus on the significance of the model?
- How does the validation procedure differentiate between various types of machine learning technologies, between different training methods for models (supervised, unsupervised learning) and between self-learning and non-self-learning algorithms?
- How are 'major' and 'minor' changes defined in the model? How big do the changes need to be before formal revalidation must take place?
- What criteria and situations are used to determine whether the applied model, data or assumptions are no longer considered appropriate?
- What is the role of scenario analysis – where the performance of the model is tested under extreme scenarios (extreme input data) that are not incorporated in the training data – in the validation procedure?
- How is the quality (accuracy, completeness, suitability) of the input data taken into account in the validation?
- For a non-self-learning algorithm/model: what is the minimal frequency of the validations?
- For a self-learning algorithm/model: how can the training process as such be validated? What role do continuous sanity checks and output restrictions play?

3.3 AI and the consumer

What are the key considerations regarding the duty of care when applying AI?
AI applications offer insurers many options for a more detailed analysis of consumer characteristics. The importance of being technically able to trace and explain the results of AI applications has been discussed earlier. Just as important, however, is that AI applications are used in a way that is socially acceptable and explainable.

3.3.1 Online decision environment

It is important that AI applications that are used in decision environments encourage consumers (either consciously or unconsciously) to make decisions that benefit their financial well-being. A decision environment refers to the context in which consumers make purchasing decisions depending on the price offered (premium). The insurer can influence decision environments in various ways. Examples are expanding or limiting the product range or by nudging tactics, in which consumers are indirectly guided towards certain choices.

Behavioural pricing – also known as dynamic pricing or pricing optimisation – is another instrument that is commonly used in the decision environment. This involves pricing an insurance product based on consumer behaviour. This behaviour is largely independent of the consumer's risk profile, and relates more to the likelihood that a consumer/customer will take out or cancel an insurance product (elasticity), or an individual's expected value for the insurer. The increasing availability of (personalised) data and analysis methods make AI applications more readily usable for behavioural pricing. This poses the risk that unreasonable variables will be used in the decision environment, or that the environment will direct consumers towards choices that are not necessarily in their financial interest.

Key consideration 9 - Designing a decision environment

- It is important that AI applications being used in an insurer's decision environment, encourage consumers either consciously or unconsciously to make decisions that benefit their financial well-being.

The following three considerations may be taken into account when developing decision environments:

- What may be expected from an 'average', rational consumer, e.g. in terms of effort to compare providers or gathering information about the product and pricing?
- What will be the impact on consumers if they make bad choices?
- How much effort does it take for an insurer to protect consumers from making choices that are not beneficial to their financial well-being?

These questions also show the importance of dynamic decision environments that focus on the characteristics of different products and different groups of consumers. What may be expected from one group of consumers – e.g. young people or university-educated people – is not necessarily reasonable for other groups, e.g. the elderly. In this regard, differentiation between products is also a consideration. For instance, the potential negative impact of making a wrong choice for vehicle insurance will probably be less significant than for disability insurance or life insurance.

This by no means implies that AI applications should not be used in a decision environment. AI applications may be deployed in the product development process, which is governed by product development standards, to better connect the product and the target group, supported by a suitable decision environment. AI applications will in some cases make it so much easier to protect customers from making wrong choices that an insurer who does not use such applications may run the risk of failing to fulfil its duty of care.

3.3.2 Social explainability

It is important that AI applications and their outcomes are both technically and socially explainable and socially acceptable.
Social explainability goes beyond being able to explain the technology or the outcomes of AI models. It touches on the question of whether the outcomes of insurance models can be considered socially and ethically acceptable and fair.

In this context, the type of input parameters should be considered. Insurers generally use input parameters that are statistically highly significantly correlated with the risk of a particular type of claim. There is also often an intuitive cause-and-effect relationship, so that it is clear to consumers why an insurer uses that particular parameter as a proxy for its claim likelihood. Such an intuitive relationship enhances the social acceptability of that parameter. When using AI applications, it is also important to look at the intuitive relationships between parameters and the risk to be determined.

Key consideration 10 - Testing applications for social explainability

When using AI applications, it is important to consider whether the outcomes of the application are justifiable in social terms:

- To what extent are the patterns and proxies found and used by the AI application fair and explainable from a social point of view? And how has this been tested?
- How strong is the correlation between the patterns and proxies and the insured risk?
- To what extent is there an intuitive link between the patterns and proxies found and used on the one hand, and the claim likelihood and risk for the insurer on the other?
- To what extent and how is it possible for individuals to identify, demonstrate and draw attention to any deviations from the peer group in which they are placed?
- To what extent is the choice for more far-reaching or less far-reaching micro-segmentation socially explainable?

It is also essential to look at how individuals are subdivided into risk groups. Creating such categories is one of the core competences of insurers. If an individual's risk profile differs in certain aspects from that of the risk group where the insurer has categorised the individual, the outcome may be found to be unfair or socially unacceptable, even if it is both technically explainable and based on intuitive causal relationships. This particularly applies to cases where individuals are excluded from coverage based on a single characteristic, for example their occupation.

AI offers insurers the opportunity to reveal more patterns and thus to categorise people into risk groups based on a larger number of dimensions. Insurers can make sure that the outcomes of AI applications are socially explainable by comparing the outcomes of models against social desirability (e.g. solidarity), by helping consumers to reduce their risks, and by helping consumers understand why their risk profile might differ from that of their peers.

As such, AI can also help to enhance social explainability: it gives insurers the opportunity to subdivide 'macro' risk groups into smaller groups, almost down to the individual level. Micro-segmentation like this may mean that customer acceptance and pricing decisions are better aligned to customers' personal risk profiles. It may also enable individuals to demonstrate that their risk profiles differ from that of their peer group. A recent example concerns the premium increase or narrowing the policy conditions for taxi insurance. More detailed risk profiles for taxi drivers based on the use of AI (e.g. analysis of IoT data) may result in a better distinction between low-risk and high-risk drivers. This may not only lead to better policy terms and/or lower premiums for drivers with a low risk profile, but will also give drivers the opportunity to influence the outcome by taking measures to lower their risk profile.

However, more sophisticated risk segmentation may also have a negative effect on the solidarity principle. This is discussed in more detail in Section 4.

4. Effect of AI on solidarity

Depending on the application, AI may have both a positive and a negative impact on insurability and solidarity. The anticipated growth in the deployment of AI underlines the need for a more comprehensive debate on solidarity in the insurance sector.

The increasing use of AI in the insurance sector harbours the risk of pressure on solidarity between different risk groups. This dilemma is not new and not exclusively the result of AI applications: the desired degree of solidarity is the subject of a broader and long-standing debate.

Extensive use of AI may have an impact on this dilemma: if insurers are able to create risk assessments at a more personal level by using more data and more powerful models, insurance premiums may be affected. The differences in premiums may increase if individual risk profiles can be determined at a more granular level. However, a development like this does not necessarily have to result in decreased solidarity. First of all, AI applications can reduce information asymmetries, which in future could improve insurability for groups that previously had difficulty finding an insurer that would cover them.

Nor is it necessarily unfair when people who run higher risks pay more. This applies in particular if an individual has a demonstrable influence on his or her risk profile and can thus also influence the premium. Here, deploying AI may even strengthen solidarity: it may increase confidence that others are doing what they can to reduce overall risk in the risk pool.

Nor does AI for pricing purposes necessarily have to result in a substantial decrease in solidarity for insurance products where individual risk cannot be easily reduced. In the case of occupational disability insurance, for example, it is a well-known fact that some professions are riskier than others, and this is taken into account in terms of pricing. AI has the potential to enable insurers to look more closely at all aspects of a person (not just at the 'occupation' parameter), which means that for some customers the premium could be relatively lower than under current pricing methods.

When it comes to products such as occupational disability insurance or life insurance, however, there is a risk that more personalised risk assessments – at least for some consumers – will lead to higher premiums or even exclusion. This may undermine solidarity in the sector, especially when individuals have little to no influence on their risk profile.

In the debate about AI and its effect on solidarity, it is primarily the sector that must take the lead: insurers can and must consider how the use of AI applications affects solidarity (also see key consideration 10). Initiatives taken by the sector – such as the Solidarity Monitor, which is published annually by the Dutch Association of Insurers and which aims to monitor developments in the area of solidarity – are very welcome. The AFM and DNB are open to join the discussion on this topic and are prepared to continue the dialogue.

It is essential, however, to have a broader social debate on solidarity and not only focus on AI applications: the topic to be discussed is the desired degree of solidarity in the insurance system as a whole. This concerns choices that affect all of society and consequently need to be made by society as a whole.

Glossary

Artificial intelligence: Applications that are based on analysing varied and large amounts of data using technologies such as machine learning.

The technologies used in AI applications (such as machine learning) are 'intelligent' in the sense that they are able to optimise rationally: within a given task they are able to choose the best action to achieve a certain goal, taking into account set criteria.

In addition to the use of intelligent techniques, AI applications also rely on data. The development of AI is being fuelled by the ever-larger volumes and diversity of available data. In this exploratory study, the term 'big data' is used to refer to these massive volumes of usable data.

Machine learning: A broad field of research that is part of artificial intelligence, focussing on improving the performance of a system by training that system based on optimisation algorithms and input data.

Unsupervised learning: A form of machine learning where the machine learns from input data that is not classified, labelled or categorised. Instead of responding to feedback, the machine identifies similarities in the data and its response is based on the presence or absence of such similarities in each new piece of data.

Supervised learning: Another approach to machine learning involves training the machine with readily available input and output. This data serves as a learning basis for the machine to perform a future task with similar input data.

Reinforcement learning: A form of machine learning inspired by behavioural psychology, where a machine learns by being rewarded for the correct performance of tasks and is punished for incorrect output. Without human intervention, the machine learns to maximise reward and minimise punishment.

Deep learning: A subset of machine learning, the difference being that deep learning solves problems in a non-linear way as opposed to regular machine learning, where a linear process is followed. The term 'deep learning' refers to the fact that the process goes through several layers before an outcome is generated.

Deep neural networks: Deep neural networks (DNN) are networks in which the statistical output of one layer is converted into input for the next layer, hence the term 'deep'. Whether such input is used in the next layer also depends on the model's threshold values.

Clustering: Clustering techniques are examples of unsupervised learning. These techniques are used to find similarities between data points such as correlations and then group these data points together. An example would be clustering emails based on patterns or specific words used in the text.

Anomaly detection: An automated process for identifying data that does not belong in a certain set or pattern. Insurance companies mainly use this technology to detect fraud in customer claims. Claims with a deviating pattern are filtered out and forwarded to a staff member who assesses the claims.

Gradient boosting: This technique is used to discover patterns ('bias') in the 'error term' of other models or predictors, thereby reducing the error term and thus improving predictions.

Random forests: As the name suggests, this type of model combines a large number of decision trees to ultimately arrive at the best possible outcome. Such models are used relatively frequently, partly because they can be used with relatively limited amounts of data. To a certain extent each of the parameters can also be 'weighted'.

Natural Language Processing: Refers to an application, based on artificial intelligence, dealing with the interactions between computers and humans. The focus is on how to program computers for processing and analysing large amounts of natural language data.

False positives: The outcome does not correspond with the actual facts. An example of a false positive in claims management would be a rejected claim that should actually have been approved based on the claim data used.

False negatives: The outcome does not correspond with the actual facts. An example of a false negative in claims management would be an approved claim that should have been rejected based on the claim data used.

Underfitting: The algorithm (model) is unable to describe trends in the data. It thus involves a substandard 'fit' for the data set used.

Overfitting: The model has been 'overtrained' (trained too specifically) on a certain dataset. The model may have a good fit for one particular dataset, but not for other data points.

Spurious correlations: A situation where two data points may correlate with each other while depending on another data point. This produces a false connection between two points, because there is no clear cause-and-effect (causal) relationship.

Bias: A distortion of outcomes due to systematic or incidental errors. People may also exhibit biases, which could result in an irrational, erroneous line of thought. Such biases may occur both unconsciously and consciously.

Explainer model: This model makes clear which choices a model has made to arrive at the outcome it produces.

Partial dependence: This visualises how each variable or predictor influences the predictions of the model.

Black box: A complex system or algorithm, the internal operation of which is hidden or difficult to understand.

Proxy: A variable that may be a derivative or statement of another, non-included, variable.

References

Actuaries Institute (2016), "The impact of Big Data on the Future of Insurance", November 2016

Balasubramanian, R., Libarikian, A., McElhaney, D. (2018), "Insurance 2030 — The Impact of AI on the future of insurance", McKinsey

Bank of America & Merrill Lynch (2018) Robotics and artificial intelligence (AI) worldwide market size estimates, based on 2018 to 2030 forecasts, by segment (in billion U.S. dollars)

Cai, E. (2014), "Machine Learning Lesson of the Day - Overfitting and Underfitting", StatBlogs

Netherlands Institute for Human Rights (2018) "What is discrimination", available in Dutch at https://www.mensenrechten.nl/nl/discriminatie-uitgelegd

The Geneva Association (2018) "Big Data and Insurance: Implications for Innovation, Competition and Privacy", March 2018

Gan, G. (2013), "Application of data clustering and machine learning in variable annuity valuation", Insurance: Mathematics and Economics, Vol. 53 (2013), pp. 795-801

Grand View Research (2017), "Artificial Intelligence Market Analysis By Solution (Hardware, Software, Services), By Technology (Deep Learning, Machine Learning, Natural Language Processing, Machine Vision), By End-use, By Region, and Segment Forecasts 2018 - 2025".

Helveston (2016) "Consumer Protection in the Age of Big Data", Washington University Law Review

Joint Committee of the European Supervisory Authorities (2018), "Joint Committee Final Report on Big Data", JC/2018/04, 15 March 2018

Kearns, M. (2017), "Fair Algorithms for Machine Learning", 18th ACM Conference on Economics and Computation, 26-30 June 2017

Mullainathan, S., Spiess, J. (2017), "Machine learning: an applied econometric approach", Journal of Economic Perspectives Vol. 31-2, Spring 2017, pp. 87-106

Nian et al. (2016), "Auto insurance fraud detection using unsupervised spectral ranking for anomaly", The Journal of Finance and Data Science

O'Neil, Helfand, Kochenburger (2017) "Big Data: Looking Under the Hood and Into the World", Pullman & Comey Attorneys LLC, 18 September 2017

Orbis Research (2017), "Global Artificial Intelligence Market Analysis and Forecast 2022 by Size, Share and Growth Rate".

Roy, R. and Thomas George, K. (2017), "Detecting insurance claims fraud using machine learning techniques", presented at 2017 International Conference on Circuit, Power and Computing Technologies.

Seely, S. (2018), "Eight use cases of for machine learning in insurance", Microsoft Azure

Tata Consulting Services (2017), "Hoe Kunstmatige Intelligentie de financiële dienstverlening verandert"

TechEmergence (2017) Share of companies investing in artificial intelligence (AI) worldwide, by industry, as of 2016, and Tata Consultancy Services

Tractica (2017), "Artificial Intelligence Software Market to Reach $89.8 Billion in Annual Worldwide Revenue by 2025".

Dutch Association of Insurers (2018), "Voertuigdata: de klant aan het stuur"

Dutch Association of Insurers (2016) "Grip op Data"

Vetzo, Gerards, Nehmelman (2018) "Algoritmes en grondrechten", Boom Uitgevers, The Hague

Wang, Y, Xu, W. (2019), "Leveraging deep learning with LDA-based text analytics to detect automobile insurance fraud" Decision Support Systems, Vol. 105 (Jan 2018), pp. 87-95

Zhang et al (2017), "Understanding deep learning techniques rethinking generalization", Cornell University Paper