Institute and Faculty of Actuaries

AI and Automation Working Party – Short Term Output

by AI and Automation Working Party

30 January 2020

Abstract

Artificial Intelligence and Automation are buzzwords that are often mentioned in the context of financial services. As a Working Party focussing on these topics, we aim to provide rigorous definitions of these topics to enable a longer term workplan. After defining AI and automation, and considering the potential overlap of these areas, we consider the key application areas of AI and Automation techniques as documented in the literature and according to the practical knowledge of the Working Party. The primary intended audience are actuaries in the insurance sector. Therefore, applications of AI cover mortality forecasting, reserve and capital approximation and investment applications, while the applications of automation focus on regulatory reporting, analysis of surplus, experience investigations and daily solvency monitoring. The paper concludes with a consideration of future topics for exploration.

Contents

Introduction

The Artificial Intelligence ("AI") and Automation Working Party was formed in 2019 to investigate how actuarial practice needs to evolve to make the best use of new AI and automation techniques, particularly in life and health/protection insurance: pricing and product design, reserving, ALM, capital management, and investment. In this report, we document our initial definition of AI and Automation, the aims of the Working Party and provide a review of applications of AI and Automation in Life and Health actuarial practice. We conclude with the future scope that the Working Party will aim to explore.

Over 2019 Artificial Intelligence has continued to build its profile across the mainstream media, with advances gaining widespread coverage, whether that be the GPT2 language model, which generates almost human-like text, released in February (OpenAI, 2019) or the news that DeepMind's StarCraft 2 bot was capable of beating 99.8% of human players it was matched against. On the other hand, recent controversies about potential discrimination caused by "black-box" algorithms have also appeared in the news, for example, the seemingly discriminatory credit limits granted to women by the Apple Card program (Natarajan & Nasiripour, 2019). From a consumer's perspective, it seems that AI is increasingly impactful - whether, driverless cars being available in some cities in the USA (Ohnsman, 2019), computers evaluating x-rays with high levels of accuracy, or machines that can beat humans at poker. But what does artificial intelligence mean for the industries in which actuaries work and the actuarial profession?

Rapid advances in artificial intelligence (AI), automation techniques and data science have already brought changes to industries such as finance, healthcare or retail, and the insurance industry is also experiencing changes. Within insurance, the importance of the new technology is already increasing - according to a relatively recent study from Tata Consultancy Services the insurance sector has spent $124 million in AI systems (Tata Consultancy Services, 2017), compared to an average of $70 million invested by other industries. Given the long history of data driven decision making in the insurance sector, the opportunity to apply AI techniques will likely become a competitive differentiator in the insurance industry. Already, some insurers claim to be using AI within products, for example, Lemonade, a US insurtech company who claim to have processed the world's first insurance claim exclusively using an AI programme.

Focussing on actuarial processes, applications of AI are emerging in experience analysis, non-life pricing, automated underwriting, claims reserving, lapse predictions, mortality modelling and life insurance valuation as well automation of reporting processes, as we review below. All of these may lead to reduced costs, lower risk and competitive advantage. Looking ahead, AI and automation have the potential to revolutionise actuarial work, perhaps leading to new roles for actuaries within insurance companies and/or work in new sectors. On the other hand, traditional actuarial skills may become less valuable over time, and thus, actuaries currently face a challenge to redefine their skillset.

Aims of the Working party

With these challenges in mind, the Working Party aims to highlight the use of artificial intelligence and automation across the full range of work carried out by life and health insurance actuaries. This will be done using the product lifecycle as a lens through which to investigate both emerging best practice and the longer-term changes that will be driven by these new technologies.

Our high-level aims are to:

1. Raise awareness of the current and future uses of AI and automation in life and health insurance

2. Provide an educational resource for actuaries covering the types of techniques used, how they might be applied, and how they can build the skills required to use these techniques
3. Consider the implications of the use of the techniques for all stakeholders

For our first aim above we will use review of publicly available information and augment this with a survey of practitioners in the area.

For the second aim, we will investigate how the use of these techniques might change over the coming years, trying to highlight the areas which will be most affected by these changes and the skills actuaries will need to develop over next decade. Throughout our work we plan to include worked examples illustrating how key techniques are used and helping actuaries to understand these techniques and use them to solve real business problems. This will consist of a GitHub repository containing open source examples of code that actuaries can test as well as implementations of simple automation using open source Robotic Process Automation tools.

The third aim of the Working Party is to consider the implications of these changes both for actuaries and those who rely on our work. Will automation reduce the demand for actuaries? Will we all be replaced by computers? How can a board and customers trust the output from an AI process? What new opportunities are there for actuaries and the Profession? The Working Party will try to answer these questions and look at more of the wider implications of these changes.

Bearing these aims in mind, in the next two sections we have attempted to provide our initial views of definitions of AI and automation, and a high-level review of how these tools are being applied to actuarial work within Life and Health insurance.

Definitions

Introduction

Our view is that it is particularly important to provide rigorous definitions of the terms “artificial intelligence” and “automation”, to ensure that the investigations of the working party are clearly defined. In defining the term “artificial intelligence”, we rely heavily on the textbook on deep learning from Goodfellow et al. (2016). In this section, we first discuss the definition of AI, which we take to mean deep learning techniques (Lecun et al., 2015) and then move on to discuss automation.

AI

The term ‘Artificial intelligence’ was first coined by mathematician John McCarthy in the Dartmouth Summer Research Project held in the summer of 1956 (Wikipedia, 2019). The proposal for the conference said:

“The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.”

The term is said to have been used to differentiate it from related research at that time including automata theory, cybernetics and complex information processing. Even though there is no standardised definition of Artificial Intelligence, it is generally understood as a sub-field of computer science which mimics human intelligence and performs tasks that the human brain is capable of such as visual perception, speech recognition, decision-making, and translation between languages.

Since the focus here is on specific tasks, we note that this definition is relatively narrower than the production of a more general artificial intelligence which could perform a wide range of tasks. Indeed, the recent advances in AI refer to more narrow applications that mimic human intelligence behaviour. These AI techniques are extremely good at handling a narrow range of specific tasks and, therefore, they are referred to as narrow AI.

In contrast, general AI refers to a type of conceptual artificial intelligence that aims to perform the broad range of tasks that the human brain is capable of. It should essentially be able to simulate the breadth of human cognitive systems and therefore should be able to reason, plan, communicate, apply common sense and judgement and integrate all these skills.

We refer the reader to Lamberton et al. (2017) for a high-level summary of what this might entail, and some of the theoretical arguments about general AI. Thus, we don’t know for certain yet whether general AI is possible in practice and or its potential expected timelines. Although all AI techniques that exist currently are examples of narrow AI, they can perform extremely sophisticated tasks, some of which include image classification, automatic text generation and translation.

Various approaches to artificial intelligence have been proposed since the inception of the field, and we refer to Chapter 1 of Goodfellow et al. (2016) for an overview. Perhaps the most successful paradigm is machine learning, which is the approach whereby machines learn by experience and acquire skills without being explicitly programmed by humans.

While explaining the objective of machine learning algorithms, Bengio & LeCun (2007, pg. 1), write that “A long-term goal of machine learning research is to produce methods that will enable artificially intelligent agents capable of learning complex behaviours with minimal human intervention and prior knowledge. Examples of such complex behaviours are found in visual perception, auditory perception, and natural language processing”.

A **machine learning** algorithm learns from a sample data which is known as the training dataset. The choice of the algorithm depends on the nature of the available dataset and the task that we intend to perform. When the algorithm receives new data it automatically redevelops itself. Therefore, as its dataset grows, the algorithm has the ability to perform its assigned task more optimally.
The main types of machine learning are:

- Supervised Machine learning - The algorithm is constructed based on a set of data inputs and outcomes. Historical information is thus used to make a prediction on future outcomes. Some examples of supervised machine learning are:
    - Linear Regression
    - Logistic Regression
    - Support Vector Machines
- Unsupervised machine learning - The algorithm is constructed only based on a set of inputs. Information on outcomes is not provided to it. The algorithm finds patterns within the data and groups it accordingly. Some examples of supervised machine learning are:
    - Clustering
    - Anomaly detection
    - Neural Networks
- Reinforcement learning - This is a type of Machine learning method where the machine constantly learns from past experience by a trial and error method. It starts at an initial state and takes actions, which moves it to a different state and this process occurs recursively.

With each action, there is a reward associated. The algorithm chooses the path of maximum cumulative reward over a time horizon[1].

Machine learning algorithms can continuously refit themselves based on new data without human interference. The human only has control on the variables used by the algorithm but not the relationship between them, which is defined by an underlying model produced by the machine learning algorithm that is not necessarily observable or transparent.

On the other hand, a different approach to AI is **knowledge-based systems**, which have logic (rules or cases) that can be observed and amended by developers. As more information is added to the knowledge base, the rules have to be manually adjusted to allow for it. A knowledge-based system is another approach to AI that relies on the knowledge of human experts to enable its decision making.

A knowledge base and an inference engine are used to achieve this. The following steps are representative of a knowledge-based system's workings:

- The knowledge base, created by human experts, is a storage information that can be accessed by the inference engine.
- A user (not a human expert) interrogates the system from a user interface which provides data to the inference engine.
- The inference engine contains a set of rules or cases that form the basis of the decision.
- The inference engine uses the knowledge base and combines this data with its logic.
- A decision or outcome is provided to the user.

The mains types of knowledge-based systems are:

- Rules Based Systems - The inference engine contains a set of rules and these are used to form the decision-making process. The rules can be in the form of conditional statements that process the information provided in the question asked of the system.
- Case Based Systems - Similar to the above but cases are used to form the decision making process instead of rules.

Both knowledge-based systems and machine learning have calculation engines that can work automatically based on a set of user inputs to calculate an outcome. According to Goodfellow et al. (2016), when dealing with complex types of data, such as images or text, it becomes increasingly difficult to specify manually the rules necessary to build a knowledge-based system, leading modern AI research to focus on machine learning techniques.

The performance of machine learning methods is heavily dependent on the choice of data representation (or features) on which they are applied. For that reason, much of the actual effort in designing and deploying machine learning algorithms goes into the design of pre-processing pipelines and data transformations that result in a representation of the data that can support effective machine learning. This process is called feature engineering, and often relies on human input to create effective features for the task at hand. Some arguments against the approach of feature engineering are the complexity, cost and expert knowledge issues, which are summarized in Richman et al. (2019).

Deep learning methods (a subfield of machine learning) are representation-learning methods which allow a system to automatically discover the representations needed for feature detection or classification from the raw data. This replaces manual feature engineering and allows a machine to both learn the features and use them to perform a specific task. This leads us to a working definition of AI: since most modern applications of (narrow) AI rely on deep learning models, <u>we formulate our working definition of AI as models that rely on deep learning techniques to derive features by</u>

---

[1] . In the €-greedy algorithm of reinforcement learning, this is done by choosing the pth with maximum historical reward with a probability of (1-€) and trying a new unexplored random action with a probability of €, where 0<€<1.

performing representation learning, which are then used within neural networks or other machine learning techniques.

Deep learning algorithms, are usually implemented as **artificial neural networks (ANN)**, which are algorithms that were originally inspired by the structure and function of the human brain. We note that this analogy no longer carries much relevance in modern deep learning research, and thus, the best way of conceptualizing these algorithms might be as a collection of highly flexible statistical techniques. These algorithms contain a collection of connected nodes called neurons representing the inputs, output and hidden or abstract features. Each neuron transmits signals to other neurons similar to the working of neurons in a human brain and the output of each neuron is computed by some non-linear function of the sum of its inputs. Each input has a weight assigned that adjusts as learning proceeds. These non-linear functions transform the features at one level (starting with the raw input) into features at a higher, slightly more abstract level. Typically, there are multiple layers. Different layers may perform different transformations on their inputs. Signals travel from the first layer (the input layer), to the last layer (the output layer) and they can also possibly go through the layers multiple times.

A diagram of an ANN is shown in Figure 1, generated using the NN-SVG software (LeNail, 2019).

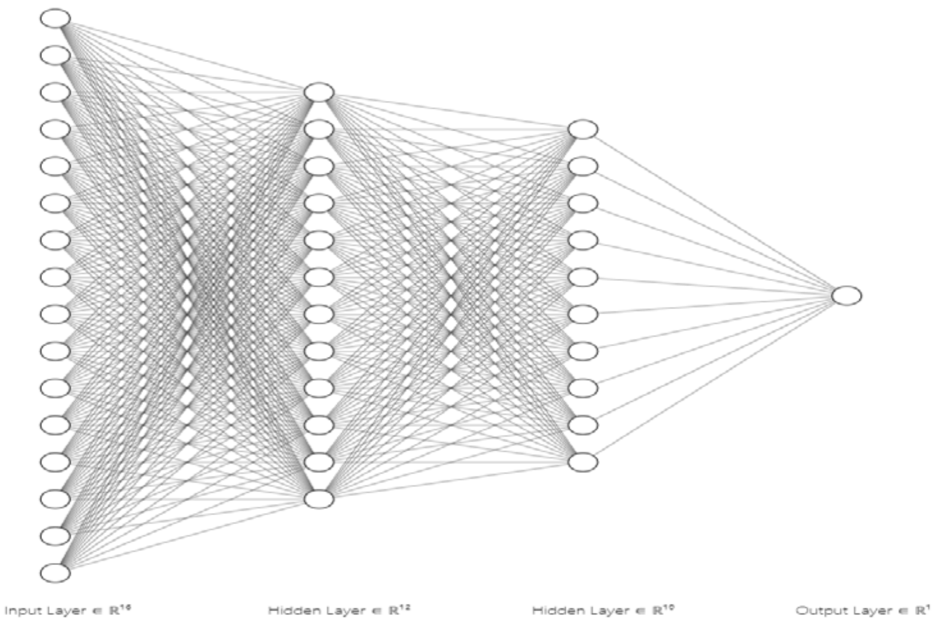

Figure 1: Diagram of a feed-forward artificial neural network.

In addition to their ability to perform automatic feature extraction from raw data, another often cited benefit of deep learning models is their scalability. With the increase in data availability, there is a need to build models that scale. As we construct larger neural networks and train them with more and more data, their performance continues to increase. This is generally different to other machine learning techniques that reach a plateau in performance. Results for deep learning get better with more data, bigger models and more computation. Deep learning excels on problem domains where the inputs (and even output) are analogue or unstructured. Meaning, these inputs are not a few quantities in a tabular format but instead are images of pixel data, documents of text data or files of audio data.

Given these algorithms help to create autonomous learning from data and information, this kind of modelling technique requires a lot more data to train on and more computational resources in comparison to other machine learning techniques. Some generic use cases of ANN are parsing texts in electronic medical records, speech sentiment analysis, recommender systems for ads on social media. These use-cases are also transforming the insurance industry which is generally data-rich. Insurance companies have data living in various silos, including text, image, and voice, but by extracting it and integrating it into a data lake, firms can use deep learning to automate processes that have traditionally relied on human involvement.

This is affecting the many areas within the insurance value chain, such as:

1. Pricing - assess customer risk, as well as optimize price based on customer segments. Analyse policyholder data from Fitbit/apple watch to offer discounts on pricing, use image processing deep learning techniques to assess biological age from photos uploaded by customers
2. Claims Management - deploying chatbots to interact with customers, and using deep learning to predict the total loss and process the claim
3. Fraud Detection - distinguish between normal and fraudulent behaviour and adapt over time based on variations of fraud patterns in the data
4. Customer Analytics - improve customer retention, identify customer purchase behaviour and develop more sophisticated marketing programs with targeted offers and upgrades for high-value customers

In closing this section on AI, we note that the models produced using deep learning are often considered to be like “black boxes”, meaning to say that the workings of these algorithms are not easily understood or inspected. Recently, so-called “Explainable AI” algorithms are becoming prominent, since these algorithms often achieve similar accuracy to deep learning models while still being transparent. For some examples of these algorithms see Vaughan et al. (2018) and Caruana et al. (2015). The different approached to explainable AI seeks to address different elements of the explainability, such as, clarity on the algorithm used, how different input variables affect the output, how influential these inputs are in predicting the output and how the input data needs to change to achieve a different output. For example, Layerwise relevance propagation (LRP) is used for determining which features in a particular input vector contribute most strongly to a neural network's output. In the context of machine learning, techniques such as partial dependence plots can help to explain the relationship between the different levels of the predictor and outcome variable in a somewhat similar way to GLM relativities.

Automation

Automation is not a new concept; insurers were early to adopt the use of computers in order to perform actuarial liability calculations that had been previously performed by scores of actuaries by hand. The goal at that time was to reduce the time and costs associated with actuarial work, a theme that resounds still today. In fact, simple forms of automation are so embedded in the day-to-day work that actuaries do (e.g. scripting in VBA or using actuarial modelling software) that all actuaries will have had some exposure to it, even if they are not aware.

However, increasingly insurers have been looking to explore and adopt the use of robotic process automation (“RPA”) software, which is the use of automated rules-based software that executes pre-programmed tasks across platforms, in order to “industrialise” the processes and procedures of the insurer in general, and in some instances, also the actuarial function. Perhaps the distinguishing characteristic of RPA compared to traditional automation tools familiar to actuaries, is that RPA software can automate processes that are driven in multiple software platforms, mimicking the actions of a human user (whereas tools such as VBA are constrained to operate only in the Microsoft Office environment). For a detailed review of automation within insurance, we refer to Lamberton et al. (2017).

Not all of the tasks performed by actuaries requires the professional and academic knowledge that actuaries have and so the goal of some insurers use of RPA has been to automate actuarial workflows and streamline any operational activities that are either regularly done, that are repetitive, or that need to be done in tight timeframes (or a combination of these). In doing so it frees up expensive actuarial resources to perform more value-add activities. RPA also brings other benefits such as reduced error rates by removing manual steps in a process, and the ability to manage system load e.g. by scheduling RPA processes to run during the night or at weekends when system demand is lower. Logs and audit trails created by RPA helps to keep track of users, runtime and any changes made. This, in turn, helps to implement proper checks and controls to minimise compliance and operational risks.

AI and Automation

To this point, we have discussed the topics of AI and Automation as independent concepts. We now consider to what extent there may be overlap between AI and Automation.

First, RPA software is increasingly utilizing AI algorithms to automate processes that previously had to be performed by a person. This software may be referred to as Intelligent or Cognitive RPA. Lamberton et al. (2017) categorize this type of software into two main streams: firstly, “use-case” specific solutions incorporate AI algorithms into a particular part of the automated process, for example, incorporating an optical character or speech recognition algorithm and secondly, combining analytics platforms with RPA.

A second, and different link between the concepts may be specific to actuarial work. A significant component of actuarial work relates to the building of predictive models of mortality, morbidity and lapses, among others. These models can be built using deep learning techniques, potentially requiring less time than the traditional actuarial approach. Thus, AI techniques may lead to the automation of tasks previously performed exclusively by actuaries.

Finally, we consider the recent emergence of so-called AutoML techniques, which are software approaches that automate the application of machine learning techniques to data, with minimal human intervention. Whereas the traditional application of machine learning requires substantial human expertise to prepare the data, including splitting the data into training and testing sets, choose a suitable machine learning algorithm and then select hyperparameters, AutoML presents an intelligent system that performs some, or all of these tasks.

Review of Applications of AI and Automation

In this section, we present a view of recent applications of AI and Automation that we are aware of. Our view of applications of AI is shaped by recent publications showing how these techniques might be applied to actuarial tasks, whereas for automation, we focus more on the practical experience assisting insurers in implementing automation techniques.

AI

The application of AI within Life and Health practice might be viewed as less mature than its application with GI firms. The recent Bank of England survey of Machine Learning (ML) within UK financial services firms (Bank of England, 2019) shows that Life Insurers have generally less mature machine learning deployments, compared to GI firms, and that fewer life insurance firms have active machine learning activity across the lifecycle of machine learning applications. Furthermore, the insurance pricing example provided in the survey relates to GI pricing, and not to Life/Health pricing. Although the survey focusses on machine learning, one could infer that applications of AI are likely distributed in the same manner.

Why might there be fewer applications of AI within the Life/Health sector compared to the GI sector? To provide rigorous answers, a survey focussing on actuaries working within Life/Health insurance might uncover greater insights and could be considered as a workstream for part of the working party.

Intuitively, though, it makes sense that that the type of models widely used by actuaries within Life/Health practice are different from those used within GI practice, and thus applying AI to these models might be less easy. In particular, much of the actuarial modelling activity within life insurance involves cashflow projection models (valuation and pricing models) whereas many of the models used within GI are predictive models (for example, IBNR, loss cost and price sensitivity models, examples of which can be found in Kuo (2018), Schelldorfer & Wüthrich (2019) and Dutang & Petrini (2018)).

However, in the phase of assumption setting, one might expect to see more applications of AI to life insurance. A consideration of the academic publications focussing on the application of AI to actuarial work shows that, indeed, this is where most applications have been proposed. On the other hand, once results have been derived from cashflow models, these can be approximated using regression techniques, including deep learning, and several applications of this technique have appeared. Finally, on the asset side of the balance sheet, several applications of deep learning have recently been proposed.

Compared to long term life and health products, short term health products such as private medical insurance (PMI) are more similar to GI products. This is because PMI presents an annual contract between the insurer and the customers and therefore the pricing and reserving is done on an annual basis, like GI products, i.e. no cashflow projection models are needed. This gives an opportunity to use AI/machine learning techniques more immediately for these annually renewable products, for example, to derive the expected cost of future claims, either traditional generalised linear models or machine learning regression models or a combination of these could be used.

A different consideration relates to the increasing access of insurers to fitness and health data for example through wearable devices, which can be incorporated to augment the traditional pricing and reserving process. While so-called real-time pricing factors are more commonly associated with GI products, such as Motor Insurance, increasingly Life and Health Insurers are incorporating these data.

Seemingly, the use of AI models to optimally utilize these news types of data will be important as they become more widespread (see Dong et al. (2016, 2017) for some examples of these models applied to telematics analysis).

In the rest of this section we focus on describing these applications at a high level.

Mortality Modelling

A significant amount of research into applying deep learning models to derive mortality assumptions has taken place. This research can be classified into two main categories - first, experience analysis, which is the derivation of mortality rates from a set of claims and exposure data, and secondly, mortality forecasting, which is the projection of mortality rates into the future.

A recent example of the experience analysis using deep learning techniques is in Rossouw & Richman (2019), who apply several machine learning models, including deep learning, to a large life insurance dataset, received from a reinsurer, containing claims and exposures from four ceding companies and covering both life and critical illness insurance. Since reinsurers may experience substantial delays before claims are notified by the ceding companies, the models are designed to both estimate Incurred but not Reported claims (IBNR), as well as predict mortality rates based on the usual covariates, such as age and gender, as well as more unusual ones, such as whether or not a policy is joint-life. Notably, the machine and deep learning methods outperform other models, such as generalized linear models. Of all the models considered, the deep learning model produces the best predictions on the unseen data.

Several recent examples of projecting mortality rates using deep learning have been published (Hainaut, 2018; Nigri et al., 2019; Ronald Richman & Wüthrich, 2019a, 2019b). Instead of using simple linear methods to model the lifetable, as is the case in the Lee-Carter mortality projection method (Lee & Carter, 1992), Hainaut (2018) uses neural networks to produce a summary of mortality rates, and then projects these summaries using time series methods, finding that this method produces more accurate forecasts than the Lee-Carter model. On the other hand, Nigri et al. (2019) model the lifetable in the same way as in the Lee-Carter model, but then project mortality rates using recurrent neural networks, finding that the predicted mortality is again more accurate than if simple time series models were to be used. Richman & Wüthrich (2019b) apply recurrent neural networks to project the population mortality rates of Switzerland, and find that the predictive accuracy of the forecast rates is better than corresponding predictions using the Lee-Carter model. This was validated by measuring the predictions made with both models on unseen data and observing that the neural network model made smaller errors. Finally, Richman & Wüthrich (2019a) show how a deep neural network can be built to forecast the mortality rates of multiple populations and genders simultaneously, achieving greater forecast accuracy than several other multi-population mortality forecasting models.

This research into mortality modelling using deep neural networks shows that for the purpose of assumption setting, deep neural networks might provide greater accuracy to actuaries than possible using traditional techniques. Of course, in practice, the outputs of the mortality models just described would probably not be used directly within pricing or reserving models, but some modifications to the rates, such as smoothing the raw outputs of the models, or applying expert judgement to adjust the forecasted rates, would occur before being used.

While these applications relate to mortality modelling, similar models could also be built to set other assumptions, such as lapse, withdrawal and morbidity rates.

Reserve Approximation

The valuation of most life and health insurance products can usually be performed relatively easily using discounted cash flow models, built in industry standard software packages. In some specific instances, though, it becomes difficult to run these models, for example, variable annuity products which are hedged daily require valuation and sensitivity results on a daily basis, and the calculation of capital requirements under Solvency II for with profits products with guarantees require specialized techniques such as nested Monte Carlo.

Hejazi & Jackson (2016, 2017) deal with the problem of the valuation of a portfolio of variable annuity contracts on a daily basis, to derive the option sensitivities (the “Greeks”) necessary to dynamically hedge the portfolio. Rather than trying to predict these sensitivities directly, they use a neural network to predict how similar each contract in the portfolio is to a representative group of contracts for which the sensitivities are known. The values for the individual contract are then derived using a weighted average of the sensitivities in the representative group, where the weights are determined by the neural network.

A different way of approximating reserves is in Castellani et al. (2019), who focus on the valuation of with-profits products under Solvency II. An ideal manner of performing this valuation would rely on nested Monte Carlo simulations. The first set of simulations, or the outer simulations, would derive the (real-world) distribution of the market risk factors at the end of the valuation year, and the second set of simulations, or the inner simulations which are run for each simulation position, would then derive the risk-neutral value of the portfolio. Performing this in practice becomes problematic, since full nested Monte Carlo takes too long to run in practice.

One accepted method of reducing the time taken to is using Least Squares Monte Carlo (LSMC) simulation, which approximates the risk-neutral value of the portfolio. In this technique, only a few inner simulations are run (reducing the time requirements), and then a regression model is fit to the inner simulations conditional on the outer simulations. Castellani et al. (2019) investigate whether a deep neural network might outperform the traditional regression approaches usually employed for this modelling, and find that the SCR is modelled with greater accuracy using deep learning.

Another related approach is in Kopczyk (2019).

Investment Applications

Although several applications of deep learning to investing have appeared, these have mainly been outside the actuarial literature; for an overview of industry applications see Cao (2019). Some of these approaches involve applying reinforcement learning, which is type of machine learning that trains agents to perform optimal strategies in an uncertain environment (Sutton & Barto, 2018), to the problem of portfolio composition (Park et al., 2019). Other approaches seek to forecast financial statement variables using deep neural networks (Alberg & Lipton, 2017). An interesting application of deep neural networks and reinforcement learning for Asset Liability Management (ALM) is in Fontoura et al. (2019), who trains an agent to determine an optimal investment strategy for a pension fund over multiple periods.

Of the three broad areas of applications that have been considered, investment applications appear to have received the least attention from actuaries.

Automation in the UK life/health space?

UK insurers use automation solutions

To date, we are aware of automation solutions being either trialled, or fully embedded in the processes of UK life and health insurers, with the automation solution designed either:

- In-house by centralised business information / information technology teams comprising of data scientists, developers and solutions architects, with input from the actuarial function. Using a third-party RPA software platform offered by a number of actuarial consultancies and software companies that are either designed specifically for actuarial processes, or more can be used more generally.

Although solutions vary, broadly speaking automation solutions involve the automated collection, transformation, storage and transfer of data, with interaction with a range of software at various points along the automated workflow.

The adoption of automation technology by UK life and health insurers have been largely driven by:

- Changes in insurance regulation (e.g. Solvency II, IFRS 17) that requires more frequent regulatory reporting, with increasingly burdensome time pressures, and greater volumes of information to be produced by actuarial teams;
- Internal pressures for actuarial teams to manage their actuarial systems, processes, and resources in a more cost-effective manner;
- The availability of cheaper storage and computing power and the ability to scale up and down computing requirements using cloud computing;
- A drive to reduce or remove model risk by using end-to-end governance and controls made available through automation technology such as scheduled workflows, user access control and timestamping; and
- Changes in data regulation (e.g. Data Protection Act, GDPR) that necessitate more robust controls on where data is stored, access restrictions, and the type of data stored.

How are UK life and health insurers using automation?

The degree to which firms operating in the UK life and health industry have adopted automation technology varies greatly, from those who are in the process of considering potential use cases, those who have implemented RPA for select processes (most commonly data processes), to those who have a single streamlined technology platform for its actuarial operations that incorporates data workflows from a number of different sources from business units across the organisation (e.g. policy admin systems, financial ledgers, claims, manual input files provided by actuaries), third-party data sources or systems (e.g. market data), and software (e.g. actuarial modelling systems, economic scenario generator ("**ESG**"), capital modelling tools).

Commonly the final step in automated workflows are to interact with data visualisation tools, such as Power BI or Tableau, in order to communicate results or key findings in dashboards and reports that stakeholders can view or interact with (e.g. drill down).
Since RPA is the use of automated rules-based software that executes pre-programmed tasks, it lends itself well to certain actuarial processes that are reasonably standardised such as valuations, experience investigations, analysis of change of surplus, and compliance.

Actuarial processes

Regulatory reporting

Changes in financial regulation over the last 10 years, in particular Solvency II, drove some UK insurers to make drastic changes in order to modernise their actuarial operating model. Solvency II requirements mean that insurers must annually complete an ever-increasing number of financial projections, creating and managing greater volumes of data in a more robust way than ever before, with increasingly tight deadlines (i.e. from the start of 2019 solo undertakings in Europe will have to submit their Solvency II results within 5 weeks of the valuation date for quarterly submissions).

Consequently, automation has been used by UK life insurers to ensure they can meet these reporting demands, with a number of insurers having in place fully automated valuation workflows, that can be scheduled to run at certain times or continuously, that can include:

- Data (policyholder and asset) collection from various systems;
- Data checking, cleansing and model point creation for actuarial models;
- Production of economic scenarios using ESGs;
- Collection of demographic and economic assumption inputs;
- Asset and liability cash flow modelling (under various stochastic or 'what-if' scenarios, and various levels of granularity (e.g. product, fund, entity level) that uses cloud computing to allow faster runs;
- Economic and regulatory capital modelling;
- Reporting of reserving and capital results.

These workflows then populate Solvency II quantitative reporting templates ("**QRTs**"), national specific templates ("**NSTs**") and financial stability reports that can be viewed in data visualisation software. Stakeholders can then view these reports that include a comparison of the figures since the last valuation date, and to see results both at an aggregate level, but also with the ability to "drill-down" into the results to see them at a fund or product level. Workflows can also automate the population of other disclosures such as the ORSA or SFCR.

More recently, the demand to adopt RPA is developing as insurers have been require to review their systems and processes for IFRS 17. Due to the nature of this accounting regulation, it requires both the actuarial and accounting systems of insurers to speak to one another in a way that was not really required before, and the storage of a vast amount of historical data. As a result, insurers that already have an automation platform in place are utilising the platform to extend the scope to include IFRS 17, with others considering the use for the first time.

Analysis of surplus

A natural progression of the automation of actuarial balance sheets is to then put in place workflows that take into account the analysis of change in that balance sheet from one valuation period to the next, in particular automating the actuarial runs required in the analysis of change in Solvency II Own Funds, and the presentation of the results of the analysis of change. This allows actuaries more time to focus on storytelling and explaining the key changes to stakeholders.

Experience investigations

Experience investigations of demographic assumptions, such as lapses and mortality, occur regularly (i.e. typically quarterly or annually) by insurers and so are a good candidate for automation by UK life insurers.
This automation leverages data from databases on claims and policyholder admin systems, then performs any required calculations for the experience analysis (e.g. exposure calculations, expected claims or lapses) before displaying the results in tables, charts and dashboards created using data visualisation software.

Daily Solvency Monitoring/ Proxy modelling

Many firms have recognised that the environment they operate in moves so quickly that a monthly or quarterly solvency snapshot (as required by regulatory reporting) is not sufficient for their risk management, as it does not allow them sufficient time to respond to the changing environment.
In order to remain competitive, certain UK life insurers are utilising automation to allow them to monitor solvency on a daily basis, enabling management to see and respond to solvency-related risks in hours instead of in weeks or months.

Other actuarial processes

Other areas where actuarial functions have utilised automation is in asset and liability analysis (informing investment and hedging strategies), economic scenario generation (often as part of a valuation workflow) and in supporting its risk management.

Automation is also increasingly found in pricing teams e.g. to produce daily MI dashboards, for example in pricing annuities where daily yield movements can affect key profitability metrics.

Governance and compliance

Automation has been used by actuaries to ease the burden of governance and compliance. For example, the management of with-profits business. In the UK, insurers who write with-profits business are required to have a With-Profits Actuary, someone who is independent of the Board of the firm, who publishes an annual report on compliance with the Principles and Practices of Financial Management ("**PPFM**") for a with-profits fund, a document that outlines how the fund is to be managed. The PPFM will have some information on how the company determines what benefits will be payable to policyholders on surrender, death or maturity, and may specify this as a percentage of asset share (e.g. 100% of asset share, or the aim is to ensure that 90% of all claims pay out between 80% and 120% of asset share). Automation has been used to extract data from claims systems (i.e. pay-outs), actuarial systems (e.g. historical asset share prices) and manual inputs (e.g. roll-forward of asset share since last valuation date), to perform calculations of claim / asset share on a policy, product, fund level, and to present this in Power BI dashboards that can allow users to easily see all key information, but also to drill in where there are outliers.

Data processes

There are cases of automation being used in the data processes of UK life and health insurers in the entire product supply chain, with actuarial data processes in scope of this. Insurers use automation to collect data from external sources such as Bloomberg, SNL, or from internal sources such as databases (Oracle, Microsoft SQL), automating the flow of that data by running extract-transform-load ("**ETL**") steps on the data, with improved levels of governance such as a full audit trail of where the data was sourced, timestamps on when the data passed certain points in the data workflow and when scripts were automatically run.

Insurers are also automating the review of actuarial policy data by automating tests on the data for completeness, exceptions to pre-defined rules (e.g. blank entries, non-standard entries, values that are out of expected bounds), and comparison to historic data sets for unusual changes to a policy's details (e.g. gender, sum assured, premium) that may be highlighting an error in the admin system. These checks can be made bespoke, taking into consideration the unique product features set out in the terms and conditions and other policyholder literature. The results of these checks can be presented in custom dashboards or reports that can be downloaded or automatically deposited in folders on the network for actuaries to then review.

Automating data processes means that actuaries can spend less time putting these checks together, and more time reviewing the quality of the data, investigating any issues (e.g. errors or omissions), and rectifying them before the data is used in actuarial work.

Operational processes

The work of actuaries often contributes to other key operational processes of insures and automation is also seen as a key opportunity to support this work. For example, UK life insurers who write with-profits business have looked to automate processes that had previously been done by actuaries in Excel, including designing an automated process for:

- Determining policy-level scaling factors to ensure surrender values are commensurate to policyholder's asset shares.
- Deriving final bonus rates for the following calendar year based on targeting smoothed asset share pay-outs.
- Deriving regular bonus rates to support annual policyholder statements that provide illustrative benefit pay-out projections for policyholder documentation.

Where to from here?

In this report, we have explained the longer-term goals of the AI/Automation Working party, defined our understanding of the key concepts of AI and Automation, and provided a preliminary view of the applications of AI and Automation within the Life and Healthcare practice areas. To build and develop the working party's long-term output, the following further work items will be considered:

- To raise awareness, a survey of current actuarial practice will be performed to understand the impact of AI and Automation techniques and recent patents will be searched to understand emerging commercial applications
- To provide an educational resource, a software workstream will provide open source code for several key applications of AI and Automation
- Finally, considering the implications of these techniques for all stakeholders, a survey and review of Explainable AI (considered in the context of the GDPR) and Fairness will be provided.

References

1. Alberg, J., & Lipton, Z. C. (2017). Improving Factor-Based Quantitative Investing by Forecasting Company Fundamentals. Retrieved from http://arxiv.org/abs/1711.04837
2. Bank of England. (2019). Machine learning in UK financial services. Retrieved from https://www.bankofengland.co.uk/-/media/boe/files/report/2019/machine-learning-in-uk-financial-services.pdf?la=en&hash=F8CA6EE7A5A9E0CB182F5D568E033F0EB2D21246
3. Bengio, Y., & LeCun, Y. (2007). Scaling Learning Algorithms toward AI. In L. Bottou, O. Chapelle, D. DeCoste, & J. Weston (Eds.), Large-Scale Kernel Machines. https://doi.org/10.7551/mitpress/7496.003.0016
4. Cao, L. (2019). AI Pioneers in Investment Management - An examination of the trends and use cases of AI and big data technologies in investments (R. Preece, Ed.). Retrieved from www.cfainstitute.org
5. Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., & Elhadad, N. (2015). Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2015-Augus, 1721–1730. https://doi.org/10.1145/2783258.2788613
6. Castellani, G., Fiore, U., Marino, Z., Passalacqua, L., Perla, F., Scognamiglio, S., & Zanetti, P. (2019). An Investigation of Machine Learning Approaches in the Solvency II Valuation Framework. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3303296
7. Dong, W., Li, J., Yao, R., Li, C., Yuan, T., & Wang, L. (2016). Characterizing Driving Styles with Deep Learning. ArXiv, arXiv:1607. Retrieved from http://arxiv.org/abs/1607.03611
8. Dong, W., Yuan, T., Yang, K., Li, C., & Zhang, S. (2017). Autoencoder regularized network for driving style representation learning. IJCAI International Joint Conference on Artificial Intelligence, arXiv:1701, 1603–1609. https://doi.org/10.24963/ijcai.2017/22
9. Dutang, C., & Petrini, L. (2018). Machine Learning Methods to Perform Pricing Optimization. A Comparison with Standard GLMs. In Variance (Vol. 12). Retrieved from http://www.variancejournal.org/issues/12-01/69.pdf
10. Fontoura, A., Haddad, D., Bezerra, E., & Haddad, D. B. (2019). A Deep Reinforcement Learning Approach to Asset-Liability Management. Retrieved from https://www.researchgate.net/publication/334761255
11. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.
12. Hainaut, D. (2018). A neural-network analyzer for mortality forecast. ASTIN Bulletin, 48(2), 481 508. https://doi.org/10.1017/asb.2017.45
13. Hejazi, S. A., & Jackson, K. R. (2016). A neural network approach to efficient valuation of large portfolios of variable annuities. Insurance: Mathematics and Economics, 70, 169 181. https://doi.org/10.1016/j.insmatheco.2016.06.013
14. Hejazi, S. A., & Jackson, K. R. (2017). Efficient valuation of SCR via a neural network approach. Journal of Computational and Applied Mathematics, 313, 427 439. https://doi.org/10.1016/j.cam.2016.10.005
15. Kopczyk, D. (2019). Proxy Modeling in Life Insurance Companies With the Use of Machine Learning Algorithms. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3396481
16. Kuo, K. (2019). DeepTriangle: A Deep Learning Approach to Loss Reserving. Risks, 7(3), 97. https://doi.org/10.3390/risks7030097
17. Lamberton, C., Brigo, D., & Hoy, D. (2017). Impact of Robotics, RPA and AI on the insurance industry: challenges and opportunities. The Journal of Financial Perspectives: Insurance, 4(1), 13.
18. Lecun, Y., Bengio, Y., & Hinton, G. (2015, May 27). Deep learning. Nature, Vol. 521, pp. 436–444. https://doi.org/10.1038/nature14539
19. Lee, R. D., & Carter, L. R. (1992). Modeling and forecasting U.S. mortality. Journal of the American Statistical Association, 87(419), 659 671. https://doi.org/10.1080/01621459.1992.10475265

20. LeNail, A. (2019). NN-SVG: Publication-Ready Neural Network Architecture Schematics. Journal of Open Source Software, 4(33), 747. https://doi.org/10.21105/joss.00747
21. Natarajan, S., & Nasiripour, S. (2019, November). Viral Tweet About Apple Card Leads to Goldman Sachs Probe. Bloomberg.Com. Retrieved from https://www.bloomberg.com/news/articles/2019-11-09/viral-tweet-about-apple-card-leads-to-probe-into-goldman-sachs
22. Nigri, A., Levantesi, S., Marino, M., Scognamiglio, S., & Perla, F. (2019). A deep learning integrated Lee Carter model. Risks, 7(1), 33. https://doi.org/10.3390/risks7010033
23. Ohnsman, A. (2019). Waymo Says More Of Its Self-Driving Cars Operating 'Rider Only' With No One At Wheel. Retrieved November 10, 2019, from Forbes website: https://www.forbes.com/sites/alanohnsman/2019/10/28/waymos-autonomous-car-definition-if-you-need-a-drivers-license-its-not-self-driving/
24. OpenAI. (2019, February). Better Language Models and Their Implications. Retrieved November 10, 2019, from https://openai.com/blog/better-language-models/
25. Park, H., Sim, M. K., & Choi, D. G. (2019). An intelligent financial portfolio trading strategy using deep Q-learning. Retrieved from http://arxiv.org/abs/1907.03665
26. Richman, R, von Rummell, N., & Wüthrich, M. V. (2019). Believing the Bot - Model Risk in the Era of Deep Learning. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3444833
27. Richman, Ronald, & Wüthrich, M. V. (2019a). A neural network extension of the Lee-Carter model to multiple populations. Annals of Actuarial Science. https://doi.org/10.1017/S1748499519000071
28. Richman, Ronald, & Wüthrich, M. V. (2019b). Lee and Carter go Machine Learning: Recurrent Neural Networks. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3441030
29. Rossouw, L., & Richman, R. (2019). Using Machine Learning to Model Claims Experience and Reporting Delays for Pricing and Reserving. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3465424
30. Schelldorfer, J., & Wüthrich, M. V. (2019). Nesting Classical Actuarial Models into Neural Networks. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3320525
31. Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction, Second Edition (Vol. 1). MIT Press.
32. Tata Consultancy Services. (2017). Getting Smarter by the Day: How AI is Elevating the Performance of Global Companies. Retrieved from https://sites.tcs.com/artificial-intelligence/#
33. Vaughan, J., Sudjianto, A., Brahimi, E., Chen, J., & Nair, V. N. (2018). Explainable Neural Networks based on Additive Index Models. Retrieved from http://arxiv.org/abs/1806.01933
34. Wikipedia. (2019). Dartmouth workshop. In Wikipedia. Retrieved from https://en.wikipedia.org/w/index.php?title=Dartmouth_workshop&oldid=916260002
35. Wikipedia. (2019). Dartmouth workshop. In Wikipedia. Retrieved from https://en.wikipedia.org/w/index.php?title=Dartmouth_workshop&oldid=916260002
36. https://royalsociety.org/-/media/policy/projects/explainable-ai/AI-and-interpretability-policy-briefing.pdf
37. https://en.wikipedia.org/wiki/Explainable_artificial_intelligence

Institute and Faculty of Actuaries

London

7th Floor · Holborn Gate · 326-330 High Holborn · London · WC1V 7PP
Tel: +44 (0) 20 7632 2100 · Fax: +44 (0) 20 7632 2111

Edinburgh

Level 2 ·Exchange Crescent · 7 Conference Square · Edinburgh ·EH3 8RA
Tel: +44 (0) 131 240 1300 · Fax +44 (0) 131 240 1311

Oxford

1st Floor · Park Central · 40/41 Park End Street · Oxford · OX1 1JD
Tel: +44 (0) 1865 268 200 · Fax: +44 (0) 1865 268 211

Beijing

6/F · Tower 2 · Prosper Centre ·5 Guanghua Road · Chaoyang District · Beijing · China 1000020
Tel: +86 (10) 8573 1000

Hong Kong

2202 Tower Two · Lippo Centre · 89 Queensway · Hong Kong
Tel: +11 (0) 852 2147 9418 · Fax: +11 (0) 852 2147 2497

Singapore

163 Tras Street · #07-05 Lian Huat Building · Singapore · 079024
Tel: +65 6717 2955

www.actuaries.org.uk