Optima

General Insurance Insights 2019

finity

How to get from A to G: Getting AI into GI

In this feature article, Finity's recent acquisition, Deep Logic, look at Artificial Intelligence (AI), its potential use in the General Insurance (GI) industry and how insurers might approach the difficult task of developing viable and practical AI solutions.

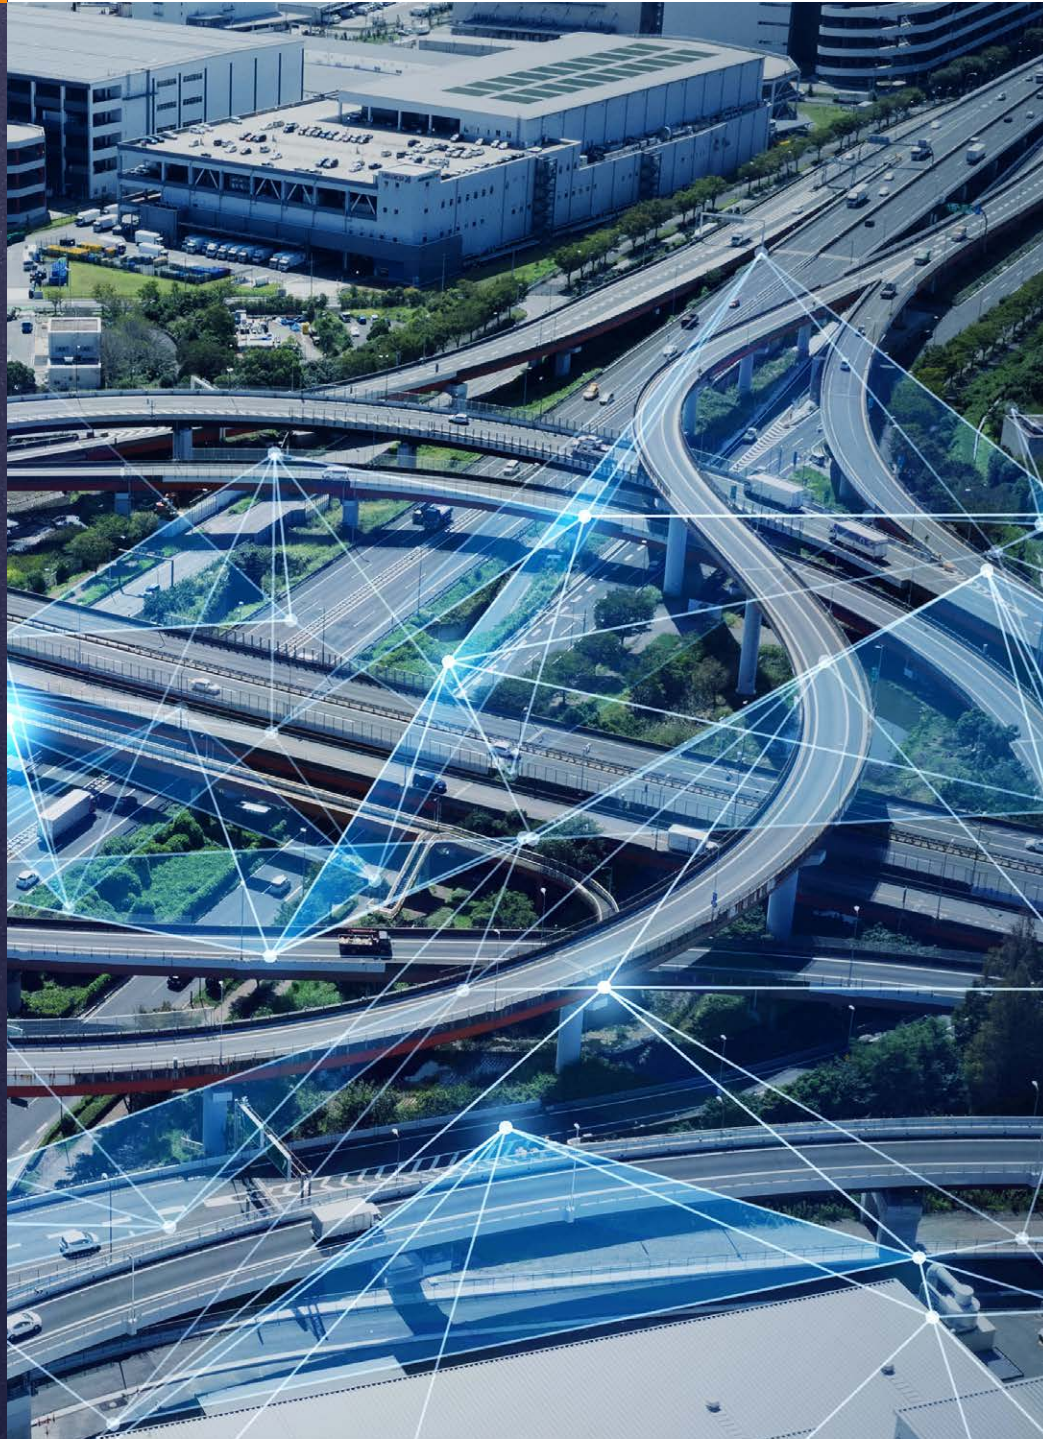

What is AI?

Artificial Intelligence (AI) encompasses a significant set of objectives. The ultimate aim of strong AI is to build a computer that rivals humans in every aspect and which can therefore pass the Turing test.

Weak AI aims to build computer systems that exceed human capabilities, but only within a restricted range of tasks. It is weak AI that has made the most progress, and has resulted in the sort of AI systems that are commonly in use today and within reach of General Insurers.

How AI is being used in the modern world

The very nature of addressing specific tasks means that AI has divided into a number of research areas – from robots to transfer learning.

Autonomous robot research is being conducted in both hostile environments and for cooperative tasks. In the collaborative realm, there is a competition named “the cocktail party”, where a robot attends to a group of people and takes cocktail orders. The robot must remember how each person who ordered a cocktail looked so that it delivers the order appropriately, even after people have moved around.

Computer vision research has been conducted in the military for many years. They have applied this to autonomous trucks and tanks. Commercial applications have now resulted in self-driving cars and facial recognition software.

The imperative to develop reasoning and logic in the AI world has led to the creation of knowledge based or expert systems, via knowledge representation schemes and inferencing capabilities. Developments in the latter include extensions of fuzzy reasoning and the ubiquitous recommender systems.

Today, research has progressed to transfer learning – how what we learn from one area can be used in another. How transferrable the knowledge is depends critically on the difference between the area where it was acquired and the area in which it will be used. While all AI systems suffer from concept drift, transfer learning is particularly susceptible when the inputs to a learning system change over time or its output changes – causing predictions to become less accurate. This is particularly prevalent in social networks, where the topics change quite rapidly and even the shape of the connections alter, as new users become connected and old connections are severed.

But machine learning is the current star of AI. It aims to develop algorithms that learn from examples in data and automatically produce a model of it. The breakthroughs that led to the modern explosion of machine learning are the deep learning artificial neural network, and ensemble methods.

Artificial neural networks are where layers of neurons are combined, based on simulating brain cells and their interactions. While decades old, it was the development of mixed learning strategies that allowed these networks to go very deep.

Ensemble methods are where simple machine learning models (such as trees) are combined in order to make a collection that acts together and performs much better than any individual model – hence the ensemble term. The two famous ensemble methods are random forests and gradient boosting.

Current AI implementations in insurance and banking

Probably the most well-known AI based tool implemented in the financial services industry today is the Fair Isaac Falcon system. It was an artificial neural network based fraud detection system for credit card transactions, initially developed many years ago – so long ago that deep learning was not yet developed. It has become a worldwide standard and is currently used by many banks around the world.

Other applications in banking include detecting payment fraud, usually perpetrated via electronic funds transfer (EFT) transactions. This typically involves either rules management systems or machine learning via ensembles. We are aware that one of the big four Australian banks has collaborated with one of the universities to create an automated rule optimisation and management system.

In the international life insurance industry, the Lincoln National life underwriting system was developed to automate the underwriting of term life and has been in place for many years. It embodies the knowledge used by human underwriters along with an inferencing strategy to assess both medical and lifestyle data.Insurance companies worldwide have used the Colossus system to handle personal injury claims since the 1990s, and continue to do so today. Colossus is a knowledge based system that uses rules to embody knowledge and inference techniques over those rules. At one time it was handling one in every two personal injury claims in the United States.

Another personal injury handling tool is Claims Outcome Advisor. It handles both personal injury and Workers' Compensation claims. It was developed in the early 2000s and it has an embedded model of the human body. It is also used by insurance companies worldwide and was adopted as standard by a European country.

Overseas insurers use what used to be EagleEye Analytics' Talon machine learning platform for insurance, which is now Guidewire PA. It comprises a number of both statistical and machine learning techniques, especially modified for insurance datasets and incorporated some pioneering work on the science of data.

Closer to home, one Australian insurer contracted a local AI group to build a 'contract works' underwriting and pricing system. The system created a representation of construction projects and then used geolocation to add supplemental information to the project by finding weather related risks for houses, large buildings, roadways, bridges and more. Originally developed in the mid-1990s, this system was used until ten years ago.

Australia has a long history of systems development within the insurance industry. A number of policy and claims administration systems that became globally successful were developed here.

Potential uses of AI implementations in General Insurance

Underwriting

In many classes of business, underwriting means automatic policy acceptance, and this is certainly the case for most personal lines. The majority of the cost has been driven out from personal lines underwriting by eliminating paper proposals, limiting phone access, and encouraging internet based engagement. This means these lines rely on the pricing to be very accurate, since this is virtually 'no touch' acquisition.

Two issues arise from this automation. The first relates to each and every insurance application: Which ones are representative of the "usual distribution" of applications and which ones are not? To answer this question we need to understand how likely or unlikely particular applications are. Then we can identify those that have never been seen before and those that occur very rarely – and separate them from the applications that are more common. We need a process that allows the vast majority to be assessed automatically while accurately identifying likely anomalies for handling by humans.

The other issue is the nature of this likelihood, and its transition over time. Of course, if competitors change prices then this should, and does, impact on the distribution of policy applications. Detecting this change in the distribution between time periods informs the business about what is happening to the flow of new customers and how this is changing. This allows formulation of a targeted response and represents new capabilities for insurers.

In commercial lines there is some scope for human underwriter discretion in vetting applications. As with personal lines, there is a push to remove costs from the underwriting function. For the small to medium enterprises (SME), most commercial policies are no longer seen by an underwriter. For the riskier or more unusual small enterprises, as well as for larger enterprises, underwriter approval is still the norm.

Using knowledge based systems it is possible to automate the underwriting of many of these risks, and to dynamically adapt the application with risk pertinent questions. Automation, using knowledge, can remove the need for underwriter review in many cases. It can also enhance the effectiveness of the underwriter when they are required because pertinent information will be available for each risk individually. This can improve efficiency and focus review on the most pertinent cases.

Machine learning based methods are now a feature of many pricing approaches for personal lines insurers. The ensemble methods allow interactions between predictive variables to occur locally, and only when they are needed. The resulting models perform significantly better than the statistical methods previously employed, and there are opportunities to make appropriate changes to machine learning algorithms to tailor them for commercial lines books also. This would bring improved data driven pricing to portfolios that previously had to rely on more ad hoc methods.

Claims

Long tail claims present unique challenges for insurers. Matching the complexity of the claim to the experience of the claims adjuster is critical to reducing claims leakage and improving outcomes for both the claimant and the insurer – this issue is particularly relevant for Workers' Compensation.

A solution adopted by a US insurer was to augment the structured data that is captured for these claims with text mining of adjuster notes. They wanted to identify potential "jumper claims" – claims whose estimate at 45 days was low but which ultimately cost more than $50,000. The adjuster notes contained the most valuable information for identifying these claims prospectively. The models that resulted were so effective that the US state fund estimated a reduction of 4% of overall claims cost through improved attention to these problem claims. Several other US states have subsequently implemented this same approach. This is new capability aimed at reducing claims cost through reduced leakage.

Improving long tail claims cost outcomes with the Artificial Immune System

The Artificial Immune System (AIS) is an advanced anomaly detection algorithm that is loosely based on the architecture of a biological immune system. Just as a human immune system detects foreign bacteria and viruses, and then remembers these signatures so that they can be easily detected (and acted upon) in the future, the AIS detects anomalies in a data set and then creates rules to easily identify these in the future. For long tail classes of business, the AIS can identify changes in the behavior of claims by comparing the joint distribution of data between two time periods and provide very granular details, such as which claims are:

- Completely new, with novel characteristics;
- No longer occurring;
- More common than before; and
- The sorts of claims which are becoming more prevalent – together with generating a description of them.

This is a new type of information about evolving trends in long tail claims books and it is information which can be used to improve claims cost outcomes.

Improving fraud detection with the Artificial Immune System

Fraud detection in claims is a top of mind issue for insurers, and gaining increased attention. Known fraud patterns, once the frauds are found, are usually embodied in some form of supervised model. These models can be logistic regression based, rules based or ensemble based, with rules being the prevalent approach. All of these existing methods are based on supervised learning; the models are only possible once the data is labelled.

Most insurers would confirm that they believe there is significantly more claims fraud than they are able to detect with their existing rules. Furthermore, fraud is a dynamic problem – once rules are developed that detect a certain kind of fraud then the perpetrators change their approach. The AIS is designed to look for these new forms of anomalous (that is, fraudulent) claims. The output of the AIS can be used to triage claims into two classes. One class is for normal claims, the ones that occur often. The other class is for claims that seem strange, and which should be sent to more experienced claims handlers for special attention. If any claims are confirmed as fraudulent then this information can be fed back into the AIS in a virtuous cycle that further improves the detection rate of its algorithms.

Roadmap for AI implementation

There isn't one obvious roadmap for implementing AI that applies to all businesses. Every company has a unique history, which colours the data that they have access to. Similarly, each company writes a subtly different subset of the population. Knowing the business imperatives is key to determining whether any AI approaches will achieve them and which to deploy. For example, a large telco wanted to deliver their best advice to their business clients, but they were unsure how to achieve this efficiently and on a large scale. This advice was the best practice expertise of their most experienced staff in shaping a telco solution that best met the clients' needs. The AI solution was straightforward – a recommender system. AI experts built a system and trained the telco's staff in the technology, achieving strong skills transfer and the deployment of a successful recommender system.

What if you don't have a clearly defined imperative, but you want to explore what could be done to transform your business? In this case, collaboration between your business expertise and AI experts can work. Find AI people who have a track record of successful business implementations and collaborations. Be prepared to share knowledge – the business needs to acquire some knowledge of what AI can do, and the AI experts need to acquire knowledge about how the business functions. The idea is to match the technical potential with the business opportunity; that's why cross pollination is critical. After a proposal is agreed, a conservative approach is to conduct a proof of concept project. After that, it can be turned into a project that combines AI, project management and implementation – and eventually business transformation.

Potential traps

AI is growing in popularity and availability, making it more accessible and affordable. There is a lot of AI code out there, with much of it comes in free packages such as R and Python. This narrows down the skillset required, providing businesses with access to a sufficiently sizeable pool of potential expertise. However, it is important for coders to truly understand how the algorithms work. This is essential in order to deliver solutions that can be tailored and adjusted throughout a project and which have a high likelihood of ultimately meeting the business' needs. Businesses therefore need to tread carefully in selecting AI coding/developer resources.

Then there are the many companies that offer specialist AI functionality, using more in-depth AI techniques. These are not free (like R and Python) and it is the algorithms themselves that are the valuable IP that is being licensed (rather than sold). This approach offers the prospect of a shorter development time and potentially more sophisticated functionality, but do not expect the offerings to be fully transparent and the workings will generally not be available for inspection. A proof of concept approach might be the best pathway here.

No matter which option you choose to progress your AI strategy, it is important that you understand what your business imperatives are. Start small, check often on progress and ensure that the problem being solved is the one you actually want to be solved!

There have been a lot of data science initiatives that have never been implemented. One noteworthy example in recent times relates to a very large global insurance company which hired 100 data scientists. The aim was to transform the company from one based on an actuarial heritage to a more modern data driven enterprise. After 18 months not one thing from any of the data scientists had been implemented. Then they were all fired, including those who originally thought up this initiative.

Clearly no one gets any value if it's not implemented. It doesn't matter how smart the people are, if it isn't implementable then there is no value (unless the goal was purely about insight). Going from proof of concept to implementation requires augmenting the AI skillset with deep domain knowledge, operational understanding and advanced technology skills. It also means evaluating the proof of concept to see what needs to change in order to make the technology implementable so as to extract maximum business value.

There have been previous waves of enthusiasm for AI; one could refer to them as hype. In the 1980s it was expert systems. In the 1990s it was artificial neural networks. More recently it has been deep learning. Each has been over-sold, and it will happen again. Take care to understand what the business value is for you, how to approach a proof of concept, and finally how to implement the result – that is where the true value lies.