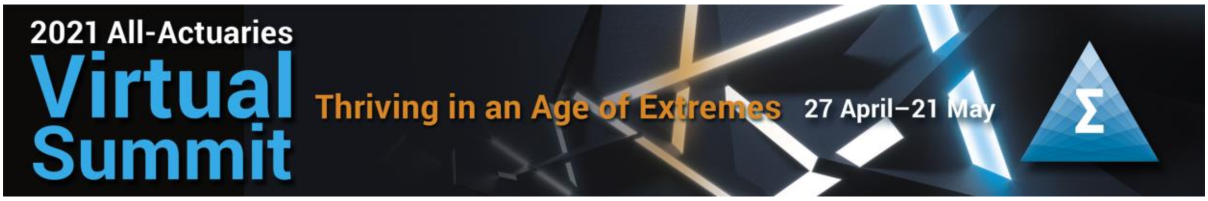

Predictive Total Loss: Using AI and Automation to Improve Motor Claims Outcomes in a Responsible Way

Prepared by Chris Dolman, Hannah Sakai and William You

Presented to the Actuaries Institute
2021 All-Actuaries *Virtual* Summit
27 April – 21 May 2021

Institute of Actuaries of Australia
ABN 69 000 423 656
Level 2, 50 Carrington Street, Sydney NSW Australia 2000
t +61 (0) 2 9239 6100 **f** +61 (0) 2 9239 6170
e actuaries@actuaries.asn.au **w** www.actuaries.asn.au

Abstract

In this paper, we outline a recent project conducted by IAG to use AI and automation techniques to improve the customer experience for motor total loss claims. We outline three aspects which should be of interest to actuaries undertaking similar AI and/or automation projects:

- Technical aspects of the data and modelling process
- Holistic aspects of the project, including effective interdepartmental collaboration and customer delivery
- Ethical aspects of the project, including customer testing and monitoring

We hope that by outlining this project in some detail, we can assist others involved in the delivery of AI solutions in their businesses.

Keywords: automation, AI, model, ethics, customer, claims, governance

Introduction and Background

An insurance claim is often referred to as the "moment of truth" by those in the industry. It is the moment when those many years of paying premiums is tested – and if the experience is not to a customer's satisfaction, policy churn is then extremely likely. It is common, therefore, for insurers to look to make this experience as positive as possible, notwithstanding that a claim circumstance is inherently a negative one to be in.

At IAG, we have recently conducted a series of experiments with the intent of improving the customer experience during a claim. This paper outlines one such successful experiment which uses a collection of technology to dramatically improve outcomes for customers. We hope this will inspire others to explore similar projects in their own businesses. We also hope to demonstrate to actuaries that the technology knowledge required to undertake such projects is within reach for most members, requiring very little upskilling from core actuarial skills.

For a complex, multi-line, multi-brand insurer it may be hard to know where to focus time and attention. We first conducted a high-level analysis to consider this question, with the aim of focussing on the areas with the biggest potential opportunity to lift the customer experience.

Motor total losses was one of the standout areas of opportunity from this analysis. This likely comes as no surprise to those familiar with motor insurance. Total losses by their nature tend to be more severe, which means customers are more likely to be experiencing injury or shock from the incident itself. Any difficulties in the claims process will compound these other challenges being experienced. Additionally, the total loss claims process in many insurers is a little more complex than a typical claim.

For a customer, this means the traditional claims process produces additional uncertainty during a time of stress. The decision to declare a total loss may take some time, and during this time there is an understandable level of uncertainty from the customer. Will I need to go out and buy a new car or not? Following a total loss, common deductions like excesses – even if communicated adequately at sales and renewal - may also be unanticipated by customers and may cause further stress or strain.

IAG's project team determined that a major underlying cause of the challenges above was the uncertainty over the total loss declaration, and the time taken to make a determination. Notably, on average a total loss took around 15 days from lodgement to finalisation, and a sizeable portion of this was time taken to assess the claim as a total loss. A project was initiated to explore ways to create greater certainty earlier and integrate this knowledge into the claims process. We also explored process re-engineering to cut down on the overall claim cycle time. It was believed that this would improve customer satisfaction and drive a reduced need for customers to contact us, saving money for IAG – a win-win.

This paper first briefly outlines the problem in more depth, then considers the "data science" aspects of the project: how we constructed a model to give greater certainty, earlier, around motor total losses. We then outline the practicalities of how such a model was used to intervene in the claims process, and the impact this had on outcomes. Finally, we consider ethical aspects of the project, noting that IAG has used this project as a pilot case study for the Australian AI Ethics Principles published by the Department of Industry, Science, Energy and Resources[1], complementing IAG's standard AI ethics process.

Creating Greater Certainty, Earlier

The Problem

In a traditional insurance claims process, it can take a considerable amount of time for a total loss to be declared, after a claim is lodged. Essentially, this arises due to the human expertise required, which is at the end of a chain of events, each of which can take some time:

1. A claim is lodged
2. The car is taken to a suitable repairer, either by the customer or by towing, if it is undrivable
3. The car waits in turn for assessment by a suitably skilled person
4. After assessment (which may involve considerable work such as removal of panels, etc), a recommendation on repairability is made
5. The insurer may accept or challenge that recommendation, as appropriate. The latter may lead to a repeat of steps 3 and 4
6. An initial determination of repairability is made. Total losses are taken for salvage at this stage; repairs are scheduled if vehicles are deemed repairable
7. During an attempted repair, further damage may still be found, which may occasionally result in a change of decision and ultimately a total loss outcome

This chain of events can take weeks, in the worst cases. Clearly this is not a good outcome for customers, who can be unaware of the potential complexity of their claim or the chances of a total loss occurring.

IAG conducted qualitative and quantitative customer research around this process to explore potential improvements. This research concluded that the core issues were:

- Lack of information about the potential for a total loss, early in the process
- Lack of information about progress during the claims cycle
- Surprises for things like excess and registration deductions

In order to remedy some of these issues, we needed information about the potential for a total loss earlier in the claims cycle. Waiting for an assessment was simply taking too long. Hence our analytics team explored whether a model could be built to predict a total loss outcome early in the claims cycle. If this were to be possible, various

[1] industry.gov.au/ai-ethics-framework

interventions could be designed using this model which could improve the customer experience.

We believe it is important to embed ethical considerations into the entire AI project lifecycle, rather than treating it as a standalone exercise. We discuss ethics in more detail below, but we note that it is at this very early point in a project: problem scoping, definition and/or ideation, that certain crucial ethical questions can often best be asked. Most important here is consideration of the intent – whether that intent genuinely benefits customers, IAG, or both, and whether that intent would be in line with reasonable customer expectations. Generally, customers would expect that we would use information provided directly by them to handle claims as efficiently and effectively as possible – depending on the final intervention, this predictive model could well be an example of that. Our general intent with this project was to improve the customer experience – whilst there might be gains for IAG as well, our primary intent was not to drive insurer value. Overall, this intent seemed an appropriate one, and we proceeded with the project in more depth.

Modelling

IAG is fortunate to have access to a rich dataset of historic claims and policy information, including all information provided by customers at claim lodgement, and the final claims outcomes.

Much of the data was structured data collected at policy sale or at claim lodgement, including vehicle age, sum insured, reported vehicle damage areas and various other similar items of data one might expect to be collected in structured form.

One potentially important item of data was a freeform claims description. This may be entered by the customer online, or, more commonly, entered by a call centre consultant on the customer's behalf. This data field has also been historically underutilised within analysis due to its unstructured nature.

It was anticipated that this freeform text data could contain important information to help predict a total loss. We explored various techniques to extract useful data from these freeform text descriptions. The words from the sentences needed to be standardised (encoded) into a structured format in order to be used in a model along with the other pieces of claims and policy data.

The text descriptions needed to be cleaned up and prepared before encoding, as illustrated in Figure 1 below. We first attempted to correct common spelling mistakes and consolidated some notable synonyms. We then applied a process known as 'lemmatisation'. Lemmatisation groups various forms of an underlying word so they can be analysed as a single item. For example, 'crash', 'crashing', 'crashes', 'crashed' are all forms of the same underlying concept and, for the purposes of analysis, can be grouped. This is particularly useful when the number of training samples is relatively small compared to the overall number of unique words across the training set.

Figure 1: Preparation of text data

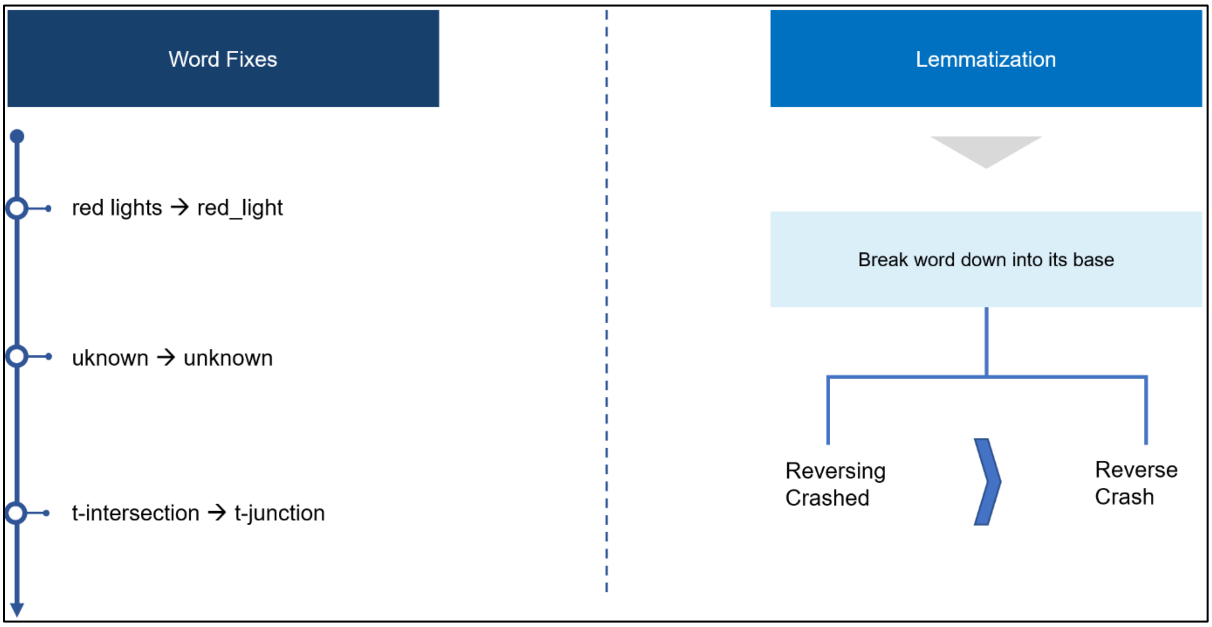

We explored several encoding techniques and ultimately went with the TF-IDF encoder. The TF-IDF encoder transforms the sentences into fixed length vectors of numbers by counting the number of words appearing in each observation, weighted by the frequency of observations in which the word appeared in the overall dataset. This gives a higher weight to words used an unusual number of times in a claims description, and a lower weight to words that are common to all descriptions, such-as common connecting words like 'a', 'and' or 'the'; or words like 'car' which would appear regularly in motor claims.

One of the more interesting techniques which we also tested was using neural nets to encode the descriptions. Embedding methods such as the Word2Vec model and the BERT model were each used to encode the claims descriptions. The result was that the encoded data was not as effective for prediction as the TF-IDF encoded data. The key reason why these state-of-the-art neural nets performed worse than the simpler encoders was the lack of data to properly tune the neural nets. We expect that with a larger sample of data, these techniques may have been valuable. Likewise, the corpus used to pre-train the BERT model is different to the somewhat unique language used in claims descriptions, requiring a much larger sample size to adequately fine tune the model for this context.

Another key aspect of the data which we had to consider when choosing features was label leakage[2]. As we wanted to utilise the model early in the claim lodgement process this meant that any useful features which were added post claims lodgement were not suitable. Similarly, any potential for items of data to change post lodgement

---

[2] Kaufman, S., Rosset, S., Perlich, C., & Stitelman, O. (2012). Leakage in data mining: Formulation, detection, and avoidance. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 6(4), 1-21.

needed to be considered. To illustrate: data is recorded and regularly updated for "damage areas" (essentially, which general part of the car is damaged) – for this model we needed to only use the vehicle damage areas reported by the customer rather than the (more accurate) vehicle damage areas subsequently reported by claims assessors during vehicle inspections.

Once the encoded claims description text field was combined with policy details, the resulting structured dataset can be used by traditional machine learning regression techniques to predict if a claim was a total loss or not. We experimented with various model forms, including GLM, GLMnet and GBM. We performed hyperparameter tuning with cross validation to select the best performing model and set of parameters. Ultimately, we found that a GBM model gave satisfactory performance with an AUC of 0.937.

We noted that a model of this form was able to be readily deployed into production once suitable interventions had been designed around it. We outline this production system architecture at a high level below.

Generally, we consider that this modelling approach is something that is open to all actuaries. The techniques used are relatively simple for technical professionals like actuaries to grasp. Tools are free and open source, and individuals can freely train themselves in their use. Any actuary should be able to construct a model of this nature, provided suitable data is available.

Intervention Design and Deployment Process

A model which drives no action is interesting at best, wasteful at worst. Our challenge now was to design an intervention which used this model in a responsible way – helping customers but not creating undesirable side effects.

It is notable that whilst the model is of reasonably high accuracy, it is not close to being perfect. This residual uncertainty limits the design space considerably. It would not be appropriate, for example, to make a final decision about a claim using this model alone.

Hence the initial intervention was designed to present useful information to customers, rather than make decisions. When a claim looked reasonably likely to result in a total loss (the threshold for which we discuss below), an SMS message was sent to customers, outlining this possibility and referring them to a website which contained information on what would occur if a total loss was the outcome.

This intervention was tested with real customers by IAG's design team, and the design and form of the message adapted iteratively to ensure it was received and understood correctly. This step is critical for success of a project of this form, and too often overlooked by data science teams.

We then designed production software to take this action, as illustrated in Figure 2 below:

Figure 2: High level Conceptual Architecture

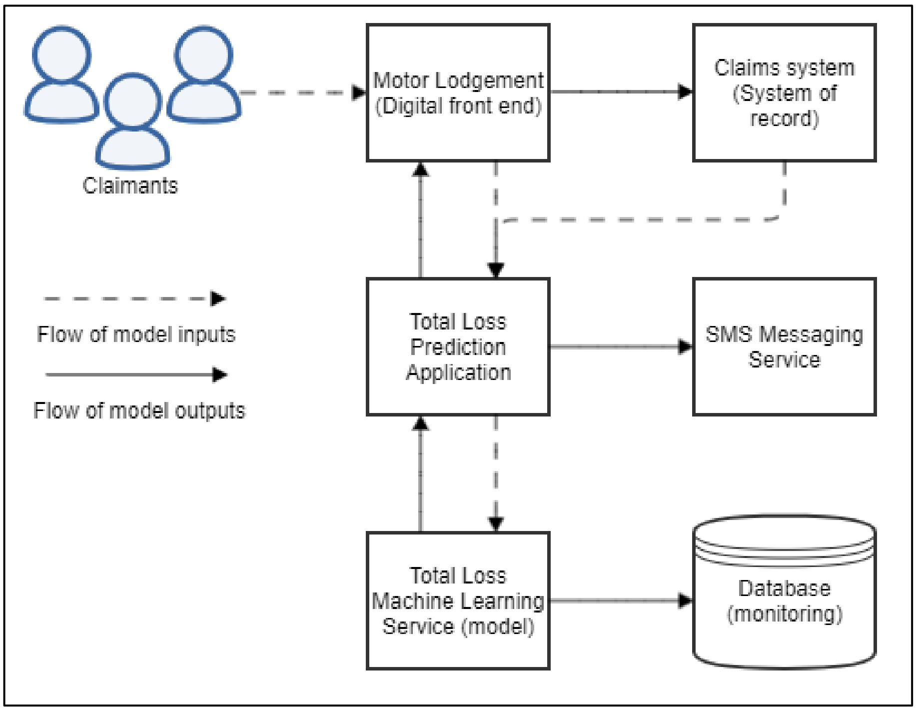

As stylised above:

- Claimants first lodge claims via the web portal (digital front end). For claims lodged over the phone, an analogous process follows via the call centre agent.
- The web portal calls the Total Loss Prediction Application layer (API) which retrieves the lodgement data as well as any necessary customer data from core systems.
- The Total Loss Prediction Application layer then forwards the input data to the machine learning service which scores the model and returns the output.
- The claims which return positive model responses are passed to an SMS Messaging service which queues the claims for SMSs to be sent to customers.
- The predicted values are also sent to the core systems in order to ensure appropriate records are kept of the decision to send a message and the reasons for it.
- The machine learning service also sends the inputs and outputs from the model to the enterprise data warehouse for monitoring purposes.

Several teams within IAG had to be engaged before this form of change could be taken to market. Whilst this may vary in other organisations, we make this observation as we anticipate similar forms of delivery complexity would exist elsewhere, and consideration of this is important for project success. Teams involved included:

- Data Governance: engaged to formally review that the use of the customer's data did not breach any data privacy laws and to formally review the ethics of the overall data-driven system.
- Legal: reviewed the contents of the SMS together with the language used in the linked website to ensure our communications were appropriate and compliant with our legal obligations.
- Digital: this team is responsible for the digital front-end systems, hence had to integrate the new system to ensure the new capabilities were triggered correctly when the customer lodged a claim.
- Group Technology: this team manages the core back-end systems, hence had to help test the new capability to ensure it did not negatively impact those core systems.
- Direct Claims: this team were the key front-line stakeholders for this project, and critical to its success. The reduction in touchpoints for the customer would reduce the time spent by staff in answering questions and doing other administrative work. It could also give rise to new questions the customer may ask the service consultants and so their team had to prepare for the process change. Change management for frontline teams is an important consideration for any project of this nature. This would be more complex and involved than some analytics professionals would anticipate!

Following this design and build, a small-scale A/B test was conducted. Whilst we were confident that this intervention would be well received, particularly as it had been tested by the design team, an initial small-scale deployment of this nature is something we consider to be responsible for projects like this, as:

- An A/B test allows us to robustly measure the (assumed) gains of a project against a randomised control group, and
- A limited deployment means that any unintended downsides have limited effect, and can be resolved prior to a full-scale deployment

The A/B test also provided an opportunity for us to test the balance between model prediction accuracy against the volume of claims impacted by the intervention. By adjusting the intervention threshold closer to a certain "total loss" model prediction (i.e. setting a threshold close to a value of 1), model precision would increase, however less messages would be sent: the volume of false negatives would also increase.

In this case, our evaluation of potential harms (outlined in more depth below) suggested that false positives were a greater concern than false negatives. Hence, we set an initial decision threshold at a relatively high level: a threshold of 0.85 giving a precision of 0.9. This still resulted in a considerable volume of interventions, but de-risked the project from a large number of false positives. We intended to review this threshold over time, particularly as the impact of false positives became better understood.

Results, Monitoring and Further Implementation

The initial test achieved outcomes which exceeded our expectations. Average gains of around 10 points in NPS were measured across a sample size of 200. No material issues were found. This gave us confidence to proceed to a full deployment. At the time of writing, around 5000 people have now received text messages warning them of a potential total loss, allowing them to better prepare for this likely outcome, early in the claims process. At the time of writing, there have been no significant complaints made in relation to false positives, which was our main area of concern prior to launch.

A small amount of model degradation was discovered shortly after launch. This was remedied by a small adjustment to the decision threshold, to maintain the desired precision.

Following this initial intervention, further interventions in operational processes making use of this model have been explored and prototypes built. However, these are still in the early stages of execution. Again, this demonstrates that there is far more to analytics project delivery than the "analytics" step – where customer facing process are being adapted, IT and change management considerations are critical. Building models is often the simplest part of a project!

Alongside model driven interventions, digitisation and automation of various aspects of the total loss process was also introduced. Whilst these aspects are not the focus of this paper, we note that these interventions dramatically improved the settlement process for total losses, reducing the number of customer touch points and improving cycle times. Under the simplified process as soon as an assessor determines the claim is a total loss the customer is notified immediately and asked to provide required information online. Once the customer uploads this information, they automatically receive a formal settlement offer and payment is then processed automatically.

Ethics

As noted above, IAG used this project as a pilot case study for the recently published Australian AI Ethics Principles. Below, we consider each of these published principles and (at a high level) our project's response to it, some themes of which we have already touched on above. We note that ethics considerations were an important and continuous part of the design process, rather than a standalone "thing to do". We consider it important that ethics considerations are embedded within a project, at all stages from design to delivery. This helps foster collegiate outcomes and regular discussions, rather than encouraging a combative "ethics signoff" process immediately prior to project deployment.

Principle 1: Human, social and environmental wellbeing

"Throughout their lifecycle, AI systems should benefit individuals, society and the environment."

The project's overall intent was to benefit IAG's customers, with no other conflicting objectives. However, the team considered it might also improve IAG's operational expenses in handling claims, and might also improve metrics like churn rates, as a by-product of improved satisfaction. They clearly documented this intent at the project outset.

Principle 2: Human-centred values

"Throughout their lifecycle, AI systems should respect human rights, diversity, and the autonomy of individuals."

When developing new solutions for customer, IAG adopts a IAG's human-centred-design process. This ensures human values are actively considered in new initiatives. An important step in early development is to test potential designs at a small scale, as outlined above. This means the impact of the system including any unintended harms can be understood and addressed before a solution scales up. This gives us confidence in the system's positive impact.

A critical design decision was to give the system limited power. No final decisions are made about claims by this system, it merely offers targeted information to assist customers. So, it doesn't fundamentally challenge human autonomy or other core values. If the model outlined was used to make claims determinations or other material decisions, errors could have been far more problematic.

Principle 3: Fairness

"Throughout their lifecycle, AI systems should be inclusive and accessible, and should not involve or result in unfair discrimination against individuals, communities or groups."

For new AI systems, at IAG we consider fairness by a detailed analysis of potential harms and benefits. Notably, we identify the circumstances that could give rise to harms, the potential degree of harm caused, and how these harms are distributed across the population. Where possible, this is measured using objective data. For this project, fairness was judged by a qualitative process of identifying the potential harms, followed by the collection of quantitative data after launch. In this case, a message being sent in error was judged to have the most potential for harm. As a result, and as discussed above, the initial decision threshold was set relatively high. This meant fewer notifications would be sent, but at a smaller risk of a message being sent in error. This design decision will be regularly reviewed, particularly in light of system performance data. We also reviewed the expected distribution of potential harms and benefits to for various vulnerable or protected groups but did not find any material areas of concern.

Principle 4: Privacy protection and security

"Throughout their lifecycle, AI systems should respect and uphold privacy rights and data protection, and ensure the security of data."

The prediction uses data the customer provided directly at the point of claim and at policy inception and did not use any additional data sources. This meant privacy and security risk was minimal, as this data is already subject to IAG's privacy and security controls. We concluded that using existing customer-provided data in this manner would align with customer's reasonable expectations that IAG would assess claims to the best of their ability using data provided. If the model used additional data, there would have been a greater need to consider appropriate privacy and security protocols and whether these additional data sources would be reasonably expected.

Principle 5: Reliability and safety

"Throughout their lifecycle, AI systems should reliably operate in accordance with their intended purpose."

As outlined above, IAG tested the system at different stages of its development and deployment. This included testing with customers, in both a qualitative and quantitative manner. This process gave us confidence that the system operated in line with its intended objectives, and that any problems or unintended harms had a good opportunity to be found and rectified. When we had to make trade-offs, for example around decision thresholds, we thoroughly documented and presented these considerations to system owners for a clear decision. Per this process, the person responsible for the system signs off on any system updates.

As observed above, we noticed a degradation in system performance shortly after launch, from our monitoring process. This allowed us to quickly adjust the decision threshold in order to preserve the desired accuracy. Without automatic monitoring, this degradation may not have been noticed for some time. This demonstrates the importance of system monitoring to help ensure ongoing reliability and safety.

Principle 6: Transparency and Explainability

"There should be transparency and responsible disclosure to ensure people know when they are being significantly impacted by an AI system, and can find out when an AI system is engaging with them."

This system doesn't significantly impact customers; it provides information only and doesn't make any material decision regarding the claim itself. Given the limited negative risks to customers, we didn't consider it necessary to make the model transparent or explainable. Instead we prioritised accuracy.

In a modern world of information overload, we don't consider it appropriate to provide advance notifications to all customers about benign systems like this. If individual customers wish to query the notification or complain about it, standard mechanisms are still available to do so, including formal complaint channels for any serious issues which may occur.

Principle 7: Contestability

"When an AI system significantly impacts a person, community, group or environment, there should be a timely process to allow people to challenge the use or output of the AI system."

As noted above, the system did not significantly impact customers or assist in material decisions. The system provided information only. This meant contestability didn't need rigorous consideration. We note in any case that contestability is a core part of existing insurance regulation - customers can use internal and external dispute resolution schemes for any complaint they might have – including about claims processing.

Principle 8: Accountability

"Those responsible for the different phases of the AI system lifecycle should be identifiable and accountable for the outcomes of the AI systems, and human oversight of AI systems should be enabled."

The development and deployment of the model had a clear business owner who took responsibility for the overall objectives, design and any problems arising from the end to end process. This aligned with the principle of accountability.

The team who developed the system will conduct regular monitoring of the system outcomes. Any material updates to the system must include re-evaluation under IAG's existing AI ethics framework and signoff from the responsible business owner. Monitoring of this system has already resulted in a minor adjustment to decisioning thresholds in response to a small degradation in model accuracy, as noted above. This demonstrates the value of monitoring and accountability.

We encouraged customers to contact us if they had any questions about the notification, and we would action any feedback appropriately. To date there have been no material complaints or major items of feedback which would warrant major change, though we remain vigilant to this possibility

Conclusions

The greater availability of data and low-cost software and technology tools allows many opportunities for insurers to improve traditional processes. Our aim in outlining this work is to show actuaries what can be accomplished, with relative ease, in a large, traditional environment. This is a dramatic shift from only a decade ago, when such interventions were generally far more costly and complex to implement, though not impossible.

We also want to highlight the integration of ethics considerations into a project of this nature, in a deliberate, targeted manner. We hope that this gives actuaries practical tips on how this might be done in their own organisations.

Finally, we hope that this showcases the great outcomes which can be achieved by an AI and automation project without resorting to "rocket science". Nothing presented above is beyond the ability of a suitably determined member of our profession. There is great excitement around AI, with press coverage every week outlining new and exciting developments which have the potential to dramatically change how we live. Whilst exciting, this risks blinding us to the potential for relatively simple projects such as that outlined. To gain the benefits of AI and modern technology, we must focus on achieving real world outcomes of genuine benefit, not be led by excitement of the latest technology.