IBM **Power Systems**

CONSIDERING THE IMPACT OF AI IN INSURANCE

Abstract

AI is built with data. How can Insurance Companies put their data to work and innovate to build new services to surprise and delight clients and make better, data driven decisions? IBM® Systems have teamed up with Insurance Industry specialists and looked at the impact that AI is having on Insurance Companies.

Sam Jones, Becky Humphreys and Mark Woolnough

Table of Contents

1. Introduction

Our primary focus in creating this white paper is to review the impact that Artificial Intelligence will have on the Insurance industry. IBM Systems have worked in collaboration with UK Insurance Industry domain experts to examine potential use cases and how the adoption of AI can drive innovative services for key target personas.

We've looked at how AI can drive automation, productivity and better decision making to internal lines of business as well as provide new services for a better customer engagement that could attract new customers and increase customer satisfaction to drive brand loyalty.

We have looked at trends in the wider Open Source based AI and data science community and analysed how they could impact the work of Insurance companies.

As a team, we explored methods to create new services incorporating AI, existing use cases from Insurance clients and use case ideas of our own. We then tied these ideas to technology. The foundation of this technology is IBM Power Systems PowerAI. IBM's adoption of Open Source combined with its enterprise know how is a compelling value proposition. It has resulted in an eco- system that offers the building blocks for adoption of machine learning and deep learning from experiment to extreme scale and covers both hardware and software co-optimised for AI workloads.

As we have shone the spotlight on specific use cases, we've provided illustrative examples using the most relevant PowerAI solution. This could be as basic as a notebook written in python, to some of the enhanced tooling and software offerings from IBM and its partners.

Another key aspect we've focused on is the end to end lifecycle of AI. There are many articles about creating and training machine learning models. We have attempted to augment this story with the management and deployment of models for frictionless integration with internal development teams and operations so that models can be easily consumed by application developers working in an agile method.

We hope you enjoy reading this document and find it useful as a guide to understanding the benefits or AI whether you are directly involved in the Insurance industry, or using this particular industry as an illustrative example you can apply to your own domain.

2. What is AI?

To begin, let's define what AI means in the context of this white paper for Insurance. The Oxford English dictionary defines AI as "The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages."

“The New Physics of Financial Services”, a report published in August 2018 by the World Economic Forum, highlights the general confusion around the term AI and what it could potentially encompass. It goes on to state that people don’t necessarily mean a particular technical approach or well defined school of computer science, but rather a set of capabilities that allow them to run their business in a new way. These capabilities can include pattern detection, foresight (probability of future events), customisation (use of data to generate optimal outcomes), decision making and interaction.

2.1. The Importance of Data

Over the last few years, AI has exploded as a new digital era. The advent of Internet of Things, Big Data and Analytics means there are billions of devices and machines generating structured and unstructured data. In 2018, Gini Rometty spoke at the World Economic Forum in Davos about the “incumbent disruptor”. She went on to further elaborate that 80% of the worlds data is not searchable and is owned by existing firms.[1] This is the kind of data that incumbent companies own. Analytics and AI could in theory put this data to use to help incumbent companies transform and pivot into new areas that are beyond Big Technology and new start-ups who have been associated as “disruptors” to incumbents.

2.2. Techniques to build AI – Machine Learning

One of the most successful methods for creating AI has been to adopt Data Science practices to shape data and create machine learning models based on techniques such as linear regression to predict a constant value. Some machine learning algorithms are trained with labelled data to create machine learning models. This is referred to as “supervised” learning.

An example of a Supervised machine learning algorithm would be a deployed AI model that will predict whether a customer is likely to churn based on observed patterns from historical data sets that have been learned by an algorithm developed using a logistic regression classification technique. The historical data used to build the model would have included a target column (the label) that would denote whether the customer churned or not. The algorithm will be developed by training with the historical subset of data and it will learn how to predict the correct target value, 0 (will not churn) or 1 (probably will churn).

Machine learning models can also be trained with non-labelled data sets using techniques to cluster data and categorisation can be performed using this technique of observed historical behaviour. This is referred to as “unsupervised” learning. An example of this would be analysing the energy efficiency of a building by feeding the model historical data on electricity and gas consumption, building size, number of floors, etc and defining a range of clusters to fit that data into. This may leave a distinctive footprint showing one cluster as efficient, and the remaining clusters as inefficient.

Machine learning has also been used to tackle categorisation for audio data, image data and video data from sources like CCTV. Feature engineering techniques can be applied to assist the developed algorithms, but this is timely and expensive to implement.

[1] URL: https://www.weforum.org/agenda/2018/01/new-era-data-responsibility/

2.3. The Advent of Deep Learning

In 2012, the ImageNet Challenge (a competition for image classification) was won for the first time by a Deep Learning Convolutional Neural network, which demonstrated a significant leap in accuracy at classifying images from previously developed algorithms.[2]

These are artificial intelligence models developed using techniques inspired by how the brain works in nature to solve the same problem using a series of connected neurons to identify common patterns and features in the different image categories, making a prediction and repeating the process hundreds, or not thousands of times in a single training run. In each iteration, the network properties are adjusted so it gets better at making the correct categorisation.

In reality, this form of artificial intelligence requires embarrassingly parallel matrix maths multiplication. Running these workloads on traditional CPU based compute platforms requires huge scale to achieve respectable results and has put these workloads beyond the reach of Enterprise and consumer clients. However, as demonstrated in the 2012 ImageNet challenge, applying GPUs to these workloads has revolutionalised Deep Learning and it is now possible, with the right architecture to use Deep Learning to create models for image classification, object detection and natural language processing with greater accuracy than legacy techniques at a price accessible to most small, medium and large businesses.

Deep Learning is now being used to solve other problems such as anomaly detection involving complex, noisy, descriptive data sets with lots of columns that capture an event or environment characteristics. Learning the encoding of the data to understand the prominent features of that data can help identify what is normal, so when data arrives that is outside a normal pattern, it can be flagged as an anomaly.

2.4. Choosing the best approach

It's important to understand that whilst deep learning can solve problems that machine learning has limited success with, its highly likely that an AI application might consist of a combination of machine learning and deep learning models to perform its task. Organisations that have adopted machine learning methods will be looking to augment their machine learning AI with deep learning to achieve better results. The world of AI is evolving and nothing stands still in technology!

2.5. Data Quality drives better results

Generally speaking, when building an AI application, the better the data quality, the better the results will be (in broad terms). However, this is augmented by understanding data bias and then being able to understand how decisions and results have been derived from machine learning and deep learning. An example of data bias would be to build an AI that may not be a good representation of the environment it is running in. An AI for a self-driving car that is trained with data from daytime will be bias towards working in the

[2] Krizhevsky, Alex & Sutskever, Ilya & E. Hinton, Geoffrey. (2012). ImageNet Classification with Deep Convolutional Neural Networks. Neural Information Processing Systems. 25. 10.1145/3065386.

daytime and not at night time. To remove data bias, you would add data from night time as well.

2.6. The right platform matters!

At IBM Power Systems, we have developed a platform specifically designed for data centric AI workloads.[3] We work with innovative partners to bring the platform to life with offerings designed to make data science, machine learning and deep learning more consumable, understandable, explainable and scalable. We also have the benefit of calling in the wider IBM Cloud and cognitive capabilities to create a compelling value proposition for our clients. The foundation to everything is our collaboration with Open Source – Open Power. According to a KDNuggets Poll in 2018, python topped this list with 65.6% share for machine learning software with many other open source tools very prominent in the survey.[4]

Cognitive services and APIs can also be called upon to complete the wider IBM story. Watson on IBM's Public Cloud or IBM Cloud Private can rapidly accelerate the time it takes to bring AI to production and can augment models trained on premise where the data lives. This is particularly true for the Insurance industry where sensitive information is stored and isolated behind the firewall.

Building data centric applications means that data gravity is key – the location of a large data set will influence where the AI will be built and trained. If the large data set lives in the cloud, hosted on a cloud providers infrastructure, then it makes sense to build and train the AI there. If the data resides on premise either in a private cloud or an enterprise data warehouse and data lake, then it would not make sense to shift all that data to the cloud and then train your model.

2.7. The AI lifecycle

Data wrangling, data manipulation building and training models is only half the story. How do you manage the life cycle of the model? How do you continuously improve the model? How do you integrate it into the application stack? How do you manage it as a micro service? How do you ensure that Development teams can access the model with zero friction and build the killer app that puts it into production for consumption? We have shone a light on the Insurance industry to illustrate how these problems are solved in this particular domain.

3. Impact of AI in Insurance

Artificial Intelligence is driving significant change in business, and insurance is no exception.

AI has the potential to transform the business model of an insurer by:

[3] URL: https://www.ibm.com/it-infrastructure/power
[4] URL: https://www.kdnuggets.com/2018/05/poll-tools-analytics-data-science-machine-learning-results.html

- **Improving the speed at which tasks can be carried out**; with Robotic Process Automation (RPA) being used to take away simple, repeatable tasks from Operational teams, and more complex actions now either being informed or carried out by trained AI models
- **Optimising the service, or 'next best action',** insurers can provide to customers, brokers, and other external third parties, based on their relationships, preferences, and past interactions
- **Providing new insights** that can be used to adjust, and eventually optimise, the way insurers price and distribute their products and services, and manage risk
- **Fundamentally changing how they operate**, both day to day and in the long term. Here there will be opportunities to move from the traditional coding of complex processes to an iterative use of trained AI models against large (enterprise) datasets.

This type of transformation has been made viable by the recent explosion of data in the world economy (there are many articles that attest that roughly 90% of today's data has been produced in the past 2 years alone) and notable advancements in deep learning techniques (neural networks) and supporting architectural frameworks.

Resulting products and services from tech leaders have since raised expectations amongst an insurer's customer and adviser base, who now may expect a service or interaction with their insurer to be just as fast, smart, and convenient; however infrequent this may be.

As alluded to above, successfully leveraging this technology requires new code frameworks, change methodologies and, ultimately, a cultural shift for insurers. However, if applied successfully, AI stands to benefit everyone; from the call centre handler, to the underwriter, to the customer. The next section elaborates on how this can be applied in practice, using select use cases across an insurer's value chain.

4. Use Cases

In this section, we have taken some ideas around use cases for AI in Insurance and have offered a further explanation on how they could be implemented. It's interesting that all these use cases augment existing capabilities, rather than replace human flow, even though many are associated with automation. They all deliver benefits like better decisions from data, increased decision speed and a better customer experience.

4.1. Inspection use case

Within the Commercial sector insurance companies carry out inspections to validate their underwriting decisions based on the exposures presented by that risk; pre-cover, post-inception or line-in with the renewal cycle. This enables them to identify any existing or potential risks and support their clients in risk management, thereby reducing their exposure. Due to the complexity and/or size of the job, inspections can often be time consuming. This use case focuses on the reduction of inspection time and increased surveyor productivity with the use of AI.

By adopting a tool like IBM's PowerAI Vision, insurance companies can create deep learning image recognition models and train them using labelled and augmented datasets to classify risks. Training the model on a large and augmented dataset can help improve model accuracy and remove risk from data bias that could be factored with smaller data sets. Once trained the model could be converted to Core ML for use with a mobile app written with Swift using frameworks like Vision to interpret Core ML models with object detection capabilities.   This would provide Surveyor with the ability to process the recorded drone footage and evaluate on the spot. An alternative would be to collect the recordings and upload and process them using an application on a cloud platform as a batch job in the office. A combination of these methods might be preferable.

Introducing the use of drones, either automated or surveyor controlled, and/or body

cameras to inspections could help reduce inspection time. Reports could be created using the parsed results of the object detection models to categorise by severity and assigned a timeframe for resolution.

Combining AI with inspections yields richer data for the insurance company. This not only benefits the surveyor whilst collating the inspection report but could assist the Underwriter in making informed underwriting decisions and ultimately help clients with their risk management.

4.2. IOT use case

One of the biggest opportunities of AI within the insurance industry is with the Internet of Things (IoT). Insurance companies are increasingly partnering with InsurTechs in order to improve their customer propositions by taking advantage of innovative solutions.

Another advantage offered to insurance companies by partnering with InsurTechs is the opportunity to cross sell to existing customers. Insurance companies could offer discounted insurance to the existing customer base of the InsurTech's or its manufactures.

There are a number of IoT smart home devices that have been developed that can detect and alert customers when there are issues within their home or commercial property, for example, leak/moisture sensors. By combining IoT and AI Insurance companies could offer a superior service to their customers.

If insurance companies could leverage the data being collected by the smart home devices, alongside additional location and weather datasets they could use predictive analytics to assist their customers, not only in detection but preventative action. For example, if a customer had installed leak detection sensors in their property, predictive analytics models could be built using the aforementioned datasets to predict which customers might be vulnerable to a leak. Insurance companies could then proactively send out repairers to replace faulty pipes before they burst leading to claims.

With skilled python developers these models could be built manually using a Jupyter notebook and machine learning libraries such as scikit-learn. Alternatively, if there wasn't the required skillset available in-house, this could be created and published using a data science workbench like IBM Watson Studio.

4.3. End to end Automation Use Case

Alongside the use cases highlighted in this section, it's also important to consider the use of AI in helping insurers automate complex processes, end to end. This typically involves the use of RPA, in tackling the simpler and repeatable tasks, alongside trained AI models for (near) real time inference.

Consider, for example, the claims assessment process. Whilst RPA, or other supported application frameworks, will be able to automate the process by which a claims assessor receives evidence, or shares the resulting outcomes, when it comes to the assessment itself, something more specialist is needed. This is where more advanced AI based techniques can be applied, for example:

- Natural Language Processing (NLP) could be used to automatically and intelligently extract insights from large written documents or audio clips, by searching for and interpreting key statements in the correct context;
- AI Vision (Image Classification and Object Detection) could be used to instantly provide a measure of risk or damage from a photograph or video clip – for example involving motor collisions;
- Select outputs could then be passed to an AI model that's been trained, using internal guidelines, policy terms and deep learning techniques, to instantly provide a coverage or settlement figure back to the claims assessor;
- A chatbot (virtual agent) could be used to guide the claims assessor or customer through the assessment process; providing a simple, straightforward overview of what's happening, and the opportunity to easily collect any supporting information about the claim.

The value of this approach has been long acknowledged, with many leading suppliers of RPA software, like Blue Prism, now supporting common integrations with popular platforms (like IBM Watson and Salesforce) that, amongst other things, can be used to train and host these types of models, for inference purposes in a more complex process.

RPA can therefore not only be seen as one, more simple and straightforward, AI based technique, but something that can be used well with other advanced AI based techniques to help insurers move forward on their automation journey.

4.4. Claims acceleration use case

AI could be applied to automate and accelerate claims. The way a client interacts with the Insurance company can be accelerated with new AI capabilities like chatbots with natural language processing capabilities such as Watson Assistant. When this is combined with visual based AI image categorisation and machine learning to estimate the cost and impact of the claim, a decision can be accelerated with full AI based automation leading to claims being paid in hours or days rather than weeks.

It is likely that this kind of automation for claims acceleration will only work in low impact claims, but even the effect of using AI to handle these claims types could lead to more manual and process driven complex claims being augmented with some of the AI techniques used in the automated process, or even just the impact of automating the claims process for these low impact incidents could free up agents to prioritise work and spend more time working on the complex claims.

It's possible that some claims like medical insurance claims could be accelerated using AI techniques to summarise forms either hand written to text (see Planet AI capabilities later on in this document) or using natural language processing to categorise and summarise a body of text and assess the impact on the claim so that evidence can be parsed faster to make informed and accurate decisions.

4.5. Pricing Sophistication use case

Insurance companies use techniques like GLMs (Generalised Linear Models) for price optimisation for sectors like car and life assurance. Pricing optimisation techniques allow insurance companies to gain a better understanding of their customers and allow them to balance capacity with demand and drive better conversion rates.

Applying machine learning and non-linear models to review pricing optimisation could represent an evolution that could challenge traditional linear models. However, there are challenges around adoption of these techniques and trusting the AI which can be viewed as a black box that cannot be fully trusted. Data journalism and model interpretability may go some way to solving these issues (see section on H2O Driverless AI below).

Adding non-traditional data sources such as weather data or travel data and unstructured data like hand written medical reports can also augment price optimisation and help insurers make better decisions.

4.6. Risk Identification and Classification use case

Machine Learning and Deep Learning techniques can be applied to historical data sets to help identify clients who might be at risk of generating "large loss" scenarios for car insurance or home insurance. Given that an insurance company is likely to have a considerable body of data associated with large loss incidents, putting this data to work and feeding it to a machine learning algorithm that will learn to detect a client who poses a risk as a large loss candidate is an intriguing use case for adopting AI, augmenting existing methods of identifying risk. Machine learning techniques might be probabilistic or decision tree based and arrive at a level of accuracy worthy of including as part of this analysis.

There could be a considerable number of parameters in a data set that might indicate patterns of large loss scenarios that just aren't detectable or intuitive when using traditional data analysis techniques. Applying deep learning to learn the "encoding" of data is another potential technique that can be applied to detect and categorise large loss scenarios accurately.

It's possible that deep learning techniques such as using auto encoders could deliver better accuracy than machine learning techniques with larger, more complex data sets that require more compute power.

As with many of the use cases discussed here, incorporating non-traditional data sources such as weather or travel as extra dimensions to traditional data sources and models can help assess and classify risk.

5. Developing an AI Strategy

Here are some key topics, challenges and suggestions that will contribute to an organisation's AI strategy:

5.1. Data

AI is built with data and is data centric. There are two aspects to AI solutions when you consider the data aspect. The first is building your AI. The second is the deployment of the AI model into production and building the AI pipeline.

When it comes to building AI, you need to ensure you gather and collect your data in a passive way. Typically, you will build AI with data at rest. It may reside in a data warehouse, or a big data Hadoop platform, or a combination of both. It is essential that interrogation of data happens in a non-intrusive way. A key consideration might be "do I want to run a large ad-hoc query on my production enterprise data warehouse". The likely answer to this is "no" and the solution might be to create a separate mirror system that is isolated from production.

When it comes to deploying AI, the next consideration is how you integrate it into the existing data flow. Sometimes, AI will run ad-hoc or batch jobs based on demand. However, some AI will involve near real time analysis and be engineered into the existing data pipeline.

IBM have teams that are dedicated to help with these scenarios and a methodology that can tease out these potential problems and provide solutions. An example of this is the "Data First Method" offered by the IBM Analytics team.[5]

5.2. Skills

AI can be built by a cast of many. Building and deploying an AI application could include:

- A Data Engineer: Somebody skilled in manipulating data using tools like python, or enterprise ETL tools like Datastage.
- A Data Domain Expert: Somebody who understands the data features. This person would be closely related to the person consuming the AI or even the same person.
- A Data Scientist: Somebody who understands statistics and algorithms and data modelling techniques.
- A Developer/DevOps: Somebody who integrates the AI into an application.
- A Business Analyst: Somebody who can measure the value and impact on the business of the developed AI application.
- An IT Architect: Somebody who will collaborate with the above team and define the requirements for the best platform and software available to the team that will host the development and deployment of the solution.

When an organisation embarks on an AI journey, they might consider moving people with an Analytics background to cover some aspects of the data engineer and data science role. This can be augmented with software accelerators that are designed to make data

[5] URL: https://www.ibm.com/analytics/datafirst

science more consumable and delivering results based on embedded best practice that masks the complexity of building from scratch. Examples of this are IBM PowerAI Vision, H2O Driverless AI and Model Maker in IBM Watson Studio.

Another approach is to recruit and deploy specialists in data preparation and data science techniques. This is potentially expensive, but will be factored into the return on investment AI offers and can start small and grow and scale depending on demand.

There is a vibrant and varied community of businesses who specialise in Data Science and AI. They can help with all stages from workshopping ideas (the art of the possible), building minimum viable products to assisting clients to build their data science team.

A holistic approach could deliver the best balance to solve skills issues: identify software accelerators that can get you started, identify enthusiastic and skilled resources internally who can transition to a new AI centric role, whilst use consultancy to develop your practice and understand the value you will deliver and be measured by internally in your organisation with a good understanding of timescales to accomplish all this.

5.3. Culture and Adoption

As evidenced in the KDNuggets survey mentioned earlier, many data scientists and data engineers have adopted Open Source. Similarly, many enterprise organisations have transitioned to Open Source based Linux solutions, utilising distributions such as Red Hat Enterprise Linux and SUSE Linux (SLES).

For AI, the benefits of Open Source projects are that they are community maintained with rapid innovation triggered by the needs in the user community. Many IBM solutions incorporate elements of Open Source and this is particularly true when it comes to offerings like PowerAI and Watson Studio Local on Power. IBM offer enterprise support for the PowerAI Open Source package to further assist enterprises adoption strategies for Open Source based AI.[6]

Within the enterprise, the adoption of Open Source tooling can seed the concept of Shadow IT, with small practices experimenting with AI and machine learning in isolation, getting lost in large and complex organisations. Given the breadth and depth of potential Open Source AI tools and frameworks and the almost tribal nature these projects foster, this can lead to a confusing and incoherent strategy when choosing to adopt AI. This could result in different lines of business doing their own thing, a lack of visibility and real results, materials scattered on laptops and GitHub repositories.

[6] URL: https://developer.ibm.com/linuxonpower/deep-learning-powerai/faq/

IBM have created an enterprise offering for a data science workbench that solves many of these issues. Watson Studio Local incorporates advanced Analytics workflows, machine learning models and deep learning models. It incorporates the key concepts of data science as a team sport, community, collaboration and IBM enterprise know how, provided as a private cloud platform using the same technology that delivers the IBM Cloud and Cloud Private experience. [7] It can also help solve problems such as how you publish and deploy machine learning models and integrate with development and manage the life cycle of your published models.

Many AI stories finish when an algorithm is built that can cleverly predict an outcome. It

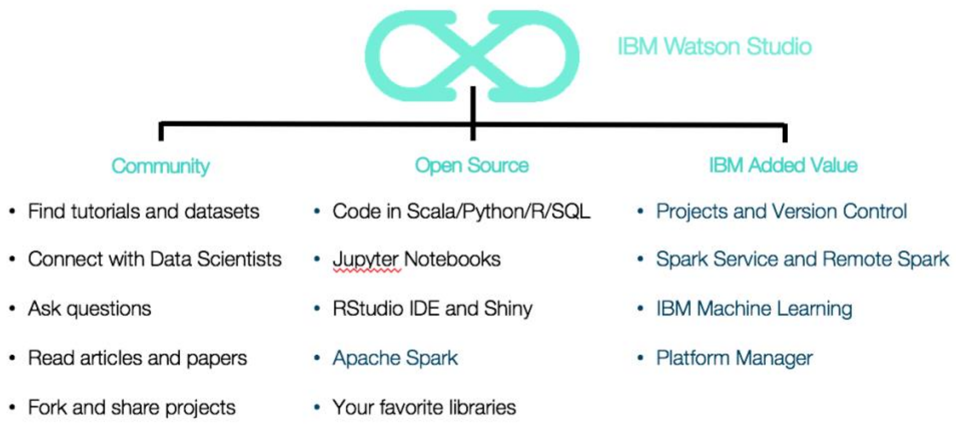

might be too complicated to incorporate this into an application though. This would limit the impact and take up of the algorithm, making it niche and less transformational. The algorithm may be obsolete in 2 months when the patterns of data have changed significantly enough to render the algorithm obsolete. It might be too difficult to alter, build and deploy the algorithm again. IBM Watson Studio Local includes a Watson Machine Learning service that could help solve both these problems with a monitoring facility and cut and paste code snippets that can be easily placed into applications.

[7] URL: https://www.ibm.com/uk-en/marketplace/watson-studio

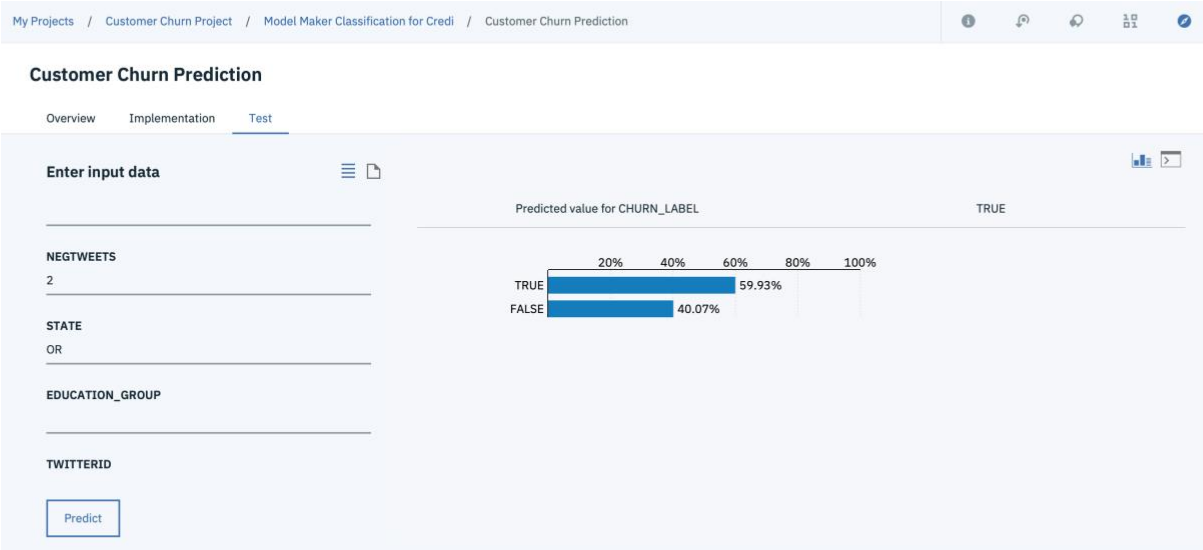

Another key challenge for many Insurers is trusting AI. Having the skills to interpret a machine learning or deep learning model is a real challenge, so some AI solutions are viewed as a "black box" and there can be a fundamental lack of trust in the developed AI if you cannot understand how a decision was made or how it arrived at a particular conclusion or predicted value. "Data Journalism" could be the next frontier for interpretability, and this is something that H2O are taking very seriously and incorporating into their offerings for AI (see section 5.5 for more details).

5.4. Infrastructure and Technology Innovation

As previously mentioned, AI is a data centric workload. Developing AI includes the manipulation of data, building a machine learning model and training it with data. This can be a time-consuming business. It's important to have a combination of the right skilled resources working with the right software on the right platform. Data gravity is a key consideration for where the workload lives. If you are training a machine learning model using a large dataset that is hosted on premise, it would make sense to host the training in the same location. The same is true if the data is located in the cloud – it makes sense to train your AI in the Cloud.

The IBM Cloud has a host of solutions for Data Science and AI including Watson services and Watson Studio.[8]

Cloud needn't be viewed as a specific location or destination for these services anymore. Cloud increasingly means a capability. Services can be hosted on a private cloud using infrastructure in a client's data centre. They can also be hosted on a public cloud. It is prudent to consider a multi cloud model, with different cloud vendors and different on premise and hosted solutions based on requirements like data security, access and control. It seems that most Insurance companies have strategies that span both Private and Public Cloud. Therefore, it might make sense to build AI applications using data hosted on premise, whilst publishing and deploying the AI on the cloud for consumption.

[8] URL: https://www.ibm.com/cloud/data

Deep Learning is a recent evolution of AI and its popularity is growing thanks to Open Source frameworks like TensorFlow. As mentioned in the opening sections of this document, if we are moving towards larger data sets to build better AI for certain scenarios leveraging GPU accelerated workflows, then it would be worth considering the impact this will have on a typical data science workbench and the kind of personas who would consume this kind of workload. IBM Watson Machine Learning Accelerator (formerly IBM Systems PowerAI Enterprise) is an offering designed to help teams of Data Scientists build, manage and scale deep learning training workloads. It includes Spectrum Conductor with Spark from IBM's Spectrum Compute family of offerings. This can help data scientists and administrators by distributing and prioritising jobs, whilst sharing a common platform and includes tooling designed to analyse and optimise deep learning jobs. [9]

PowerAI and Watson Machine Learning Accelerator is designed to be underpinned by the

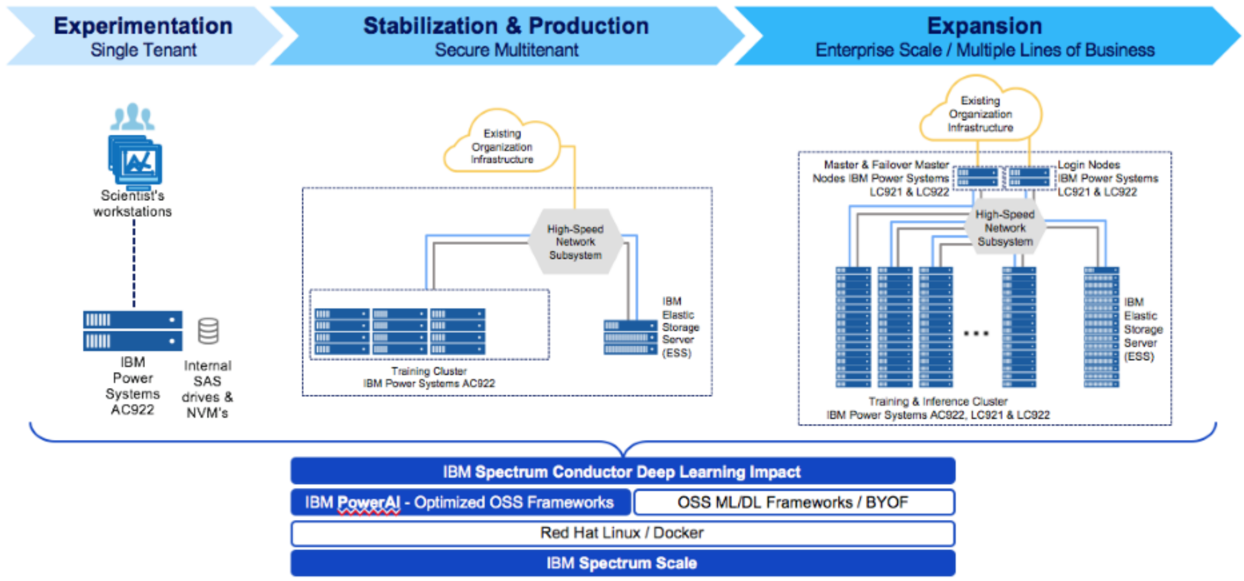

One software stack from experimentation to expansion

POWER9 accelerated compute servers (AC922) built with Open Power Foundation partners like Nvidia and Mellanox.[10] The combination of embedded Nvidia V100 GPUs with Mellanox Infiniband can scale to Super Compute workloads. In November 2018, Top 500 published its Top 10 ranking Super Computers with Coral Summit topping the list. You can read how this IBM built Super Computer performed by searching the Top 500 website which includes a high level summary of the architecture and performance.

The chart above shows the journey from the start and the capability to grow an AI Grid dedicated to multitenancy Machine and Deep Learning illustrating that the journey from single node to scale out cluster has much in common with the building blocks for Coral Summit.

5.5. Software Acceleration

IBM Systems have tooling to accelerate Deep Learning and Machine Learning. PowerAI

[9] URL: https://www.ibm.com/uk-en/marketplace/deep-learning-platform
[10] URL: https://www.ibm.com/uk-en/marketplace/power-systems-ac922/details and Power Systems AC922 Overview - IBM

Vision is a point and click solution for developing Deep Learning image classification and object detection models and has been referenced in some of the use cases in this document. Its capabilities include automatic image pre-processing, data augmentation for better accuracy, automatic labelling for rapidly building data sets, and using transfer learning to accelerate training times with high model accuracy and includes tooling to deploy models on a device or as a local API.[11] It can be further integrated to video management systems when partnered with IBM's Intelligent Video Analytics software and used to build custom models for new analytic object definition.[12]

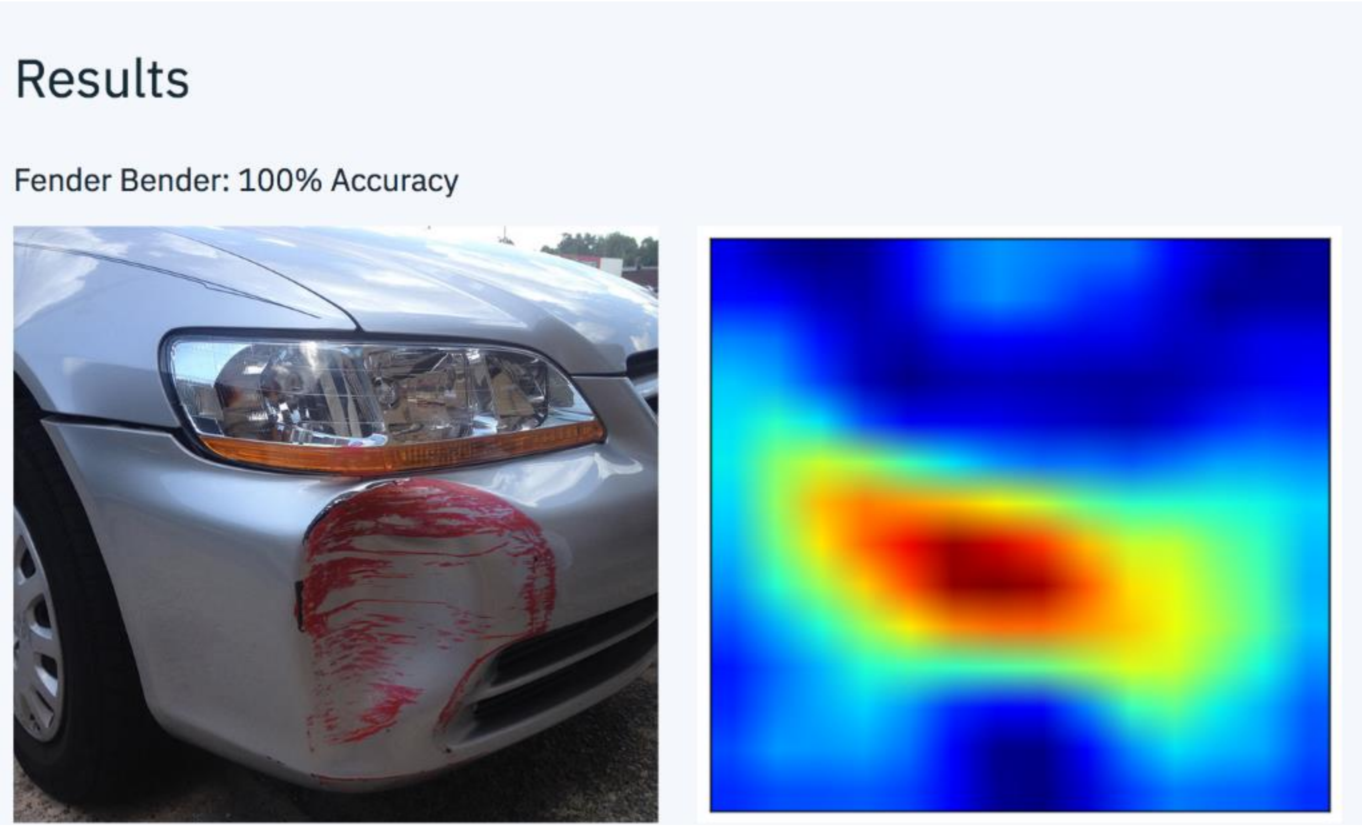

Another exciting software accelerator is H2O's Driverless AI. When you supply a labelled time dimensioned data set, you can configure H2O Driverless AI to employ several different techniques in a single "experiment" to find the most accurate model, and deploy it as a package at the click of a button. It also expands on the concept of "data journalism" so it will provide a report of the prominent features of the data set and techniques and feature engineering used to build the model.[13] It could be described as AI for developing AI, or an expert data scientist in a box.

[11] URL: https://www.ibm.com/uk-en/marketplace/ibm-powerai-vision
[12] URL: https://www.ibm.com/downloads/cas/GJBJQM4Y
[13] URL: https://www.h2o.ai/products/h2o-driverless-ai/

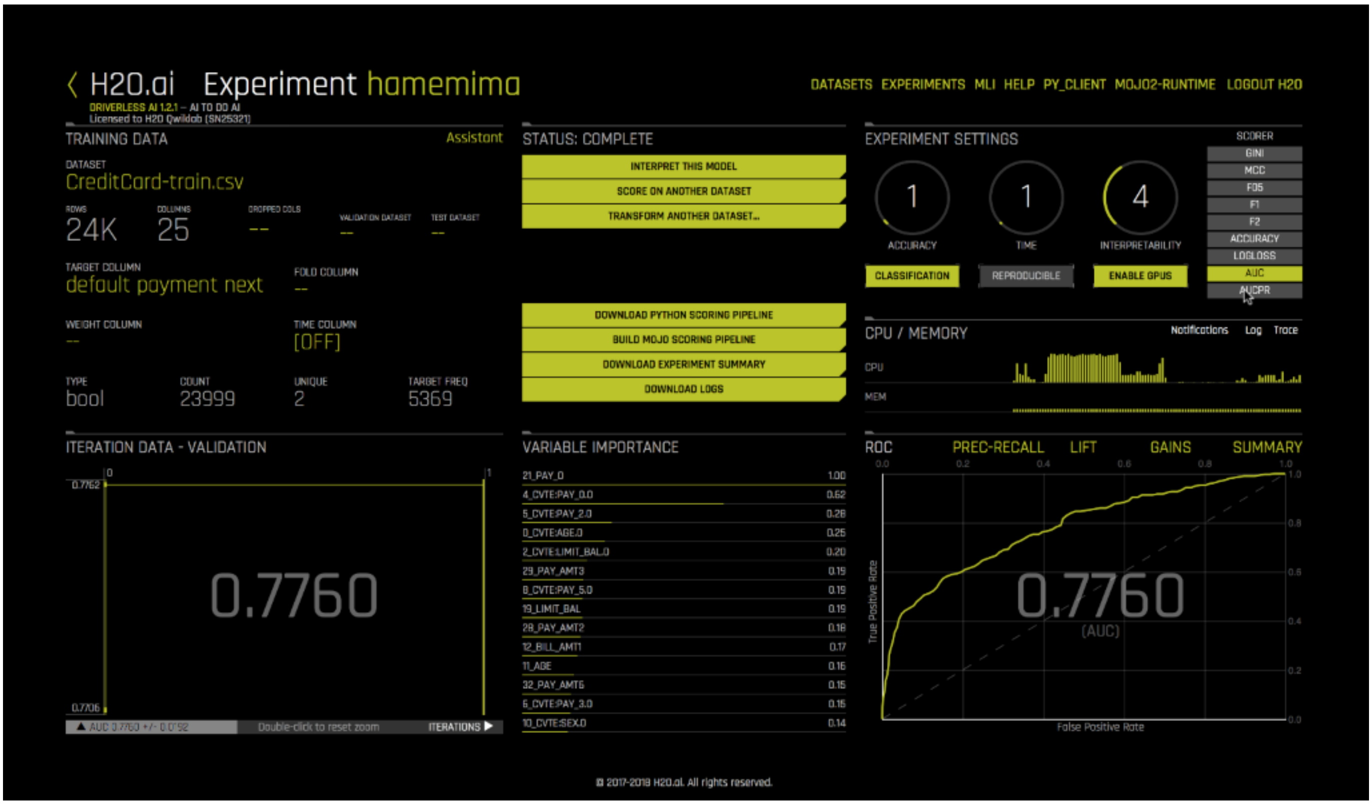

There is a vibrant Software Vendor ecosystem out there developing AI solutions. As part of this engagement, we looked at an independent software vendor called Planet AI. Using AI techniques, they have developed an automated text recognition solution. Planet AI's Intelligent Document Analysis (IDA) solution is able to convert heterogenous input data (e.g. images, PDF, handwritten, machine printed) into an electronic standard format (e.g. PDF, JSON) in a fully automated process. [14]The extensive use of the latest state of the art and multiple award winning Artificial Intelligence for text recognition and document understanding in combination with GPU based compute enables an unrivalled accuracy and speed.[15]

Within the insurance industry, this could be useful for ingesting machine printed hand written documents like medical reports and searching for key phrases and words. Access to all types of information within the documents is made possible utilising Planet's IDA framework, access to information held in the documents is made easy. Powered by Planet's PerceptionMatrix a new and unique technology that can be used for printed and cursive text recognition, all types of documents can be transformed into searchable and electronic formats.[16]

IDA is able to process more than 50,000 pages per hour on a single 4xGPU server such as the IBM Power AC922 and can leverage platform scale out capabilities for even larger solutions.[17]

---

[14] URL: https://www.planet-ai.de/ida/
[15] URL: http://www.planet.de/files/planet/dokumente/Intelligent_Document_Analysis_Framework.pdf
[16] URL: https://www.planet-ai.de/wp-content/uploads/2018/10/IDA_EN.pdf
[17]URL: http://www.planet.de/files/planet/dokumente/Intelligent_Document_Analysis_Framework.pdf

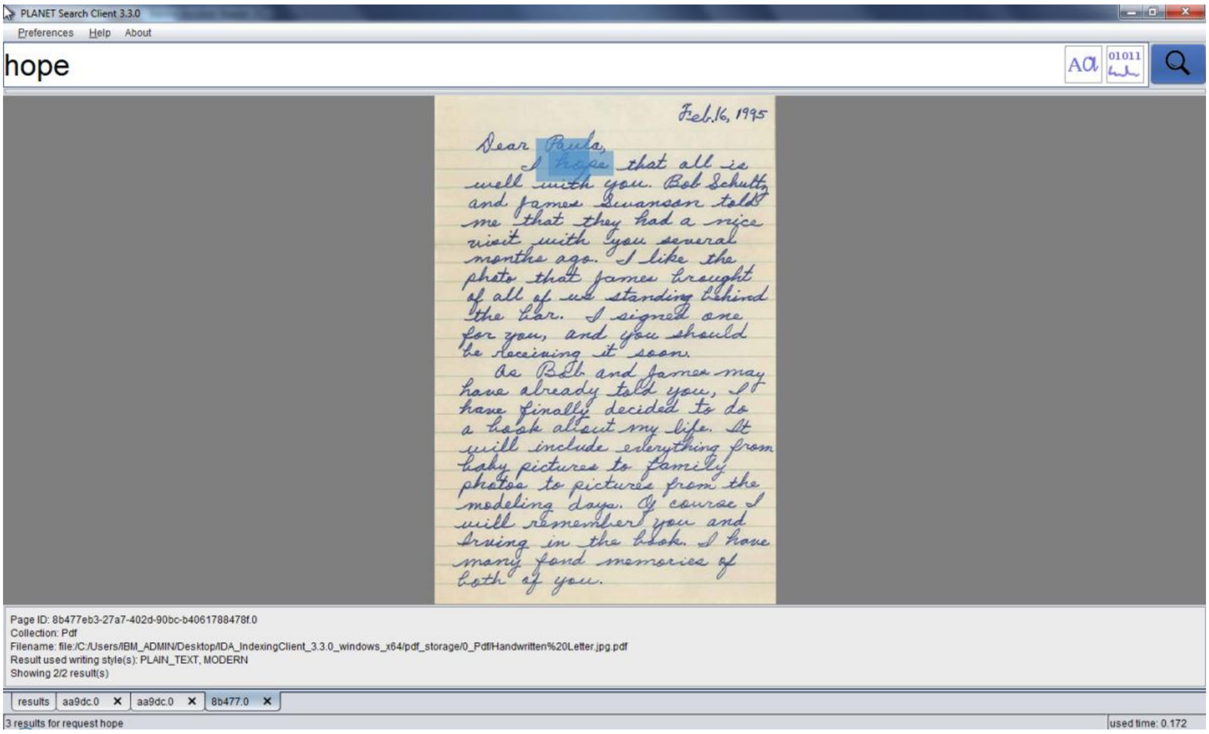

5.6. Use Cases

Another key element of understanding the art of the possible is industry use cases and developing new use cases. There are generic high level use cases that many Insurance organisations can review and adopt.

AI Insurance Segments Being Transformed By Contemporary AI

Business Segment	Applications of AI	Current Techniques
New Business/Underwriting	Scoring, Chatbots	Case based reasoning, Rules, ML, NLP
Claims	Validation, NLP, Scoring, Scene Understanding	Case based reasoning, Rules, ML, NLP, SVM
Product Development	Learning from diverse modalities of data	Mostly Text and NLP
Policy Servicing	Customer service, next best action	Descriptive analytics
Customer Experience	Personal Lines for cross-sell, upsell	Various ML, NLP
Actuarial	Add predictive power to computational models	Currently being investigated

However, the true value of AI to an Insurance company could be beyond standard use cases and be viewed as a way to augment each individual company's data assets. These will have a relatively unique footprint, so teasing out use cases based on these assets requires a different approach.

As part of the research for writing this paper, we went to the IBM Austin Briefing Centre and participated in a 2 hour Design Thinking taster. Design Thinking is a methodology that can be used to develop product and service design.[18] It is persona driven and there are several workshop types that can be used and adopted to tease out use cases and measure their impact and feasibility.

[18]URL: https://www.ibm.com/design/thinking/

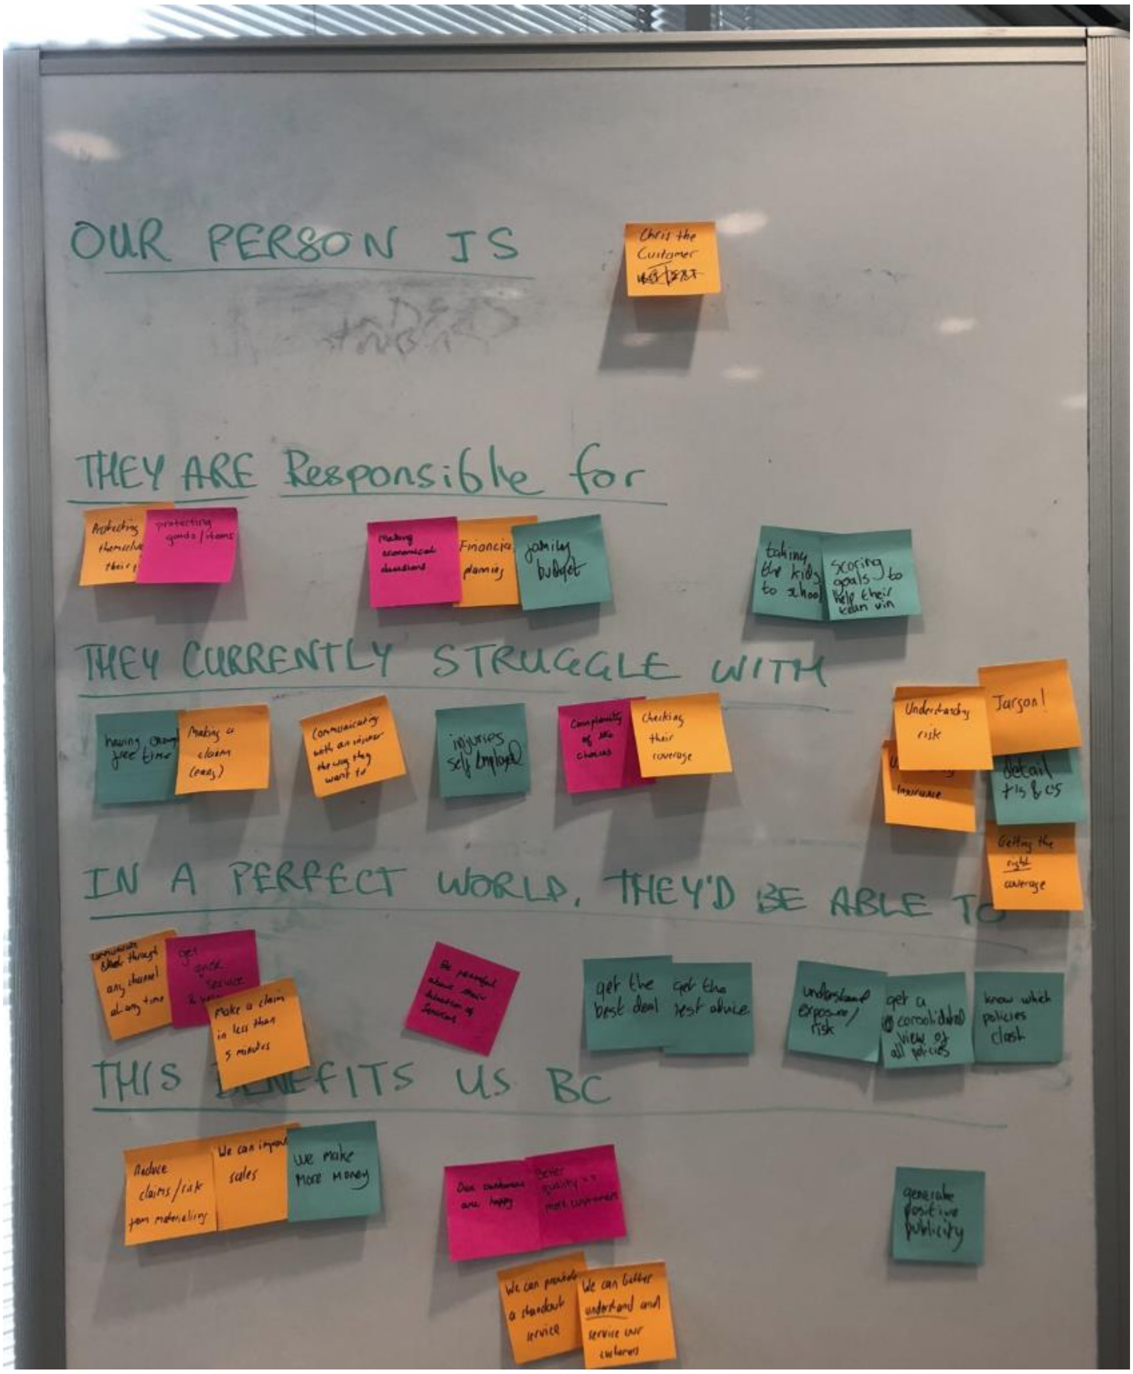

A typical engagement would involve all key stakeholders from a business unit or customer, developers, architects, engineers, data scientists and executives. It's a fun engagement and encourages creative thinking no matter how outlandish the ideas might be! In our session, we came up with an Insurance Claims time machine (low in the feasibility scale!) and a Terms and Conditions filter bot (towards high impact and high feasibility).

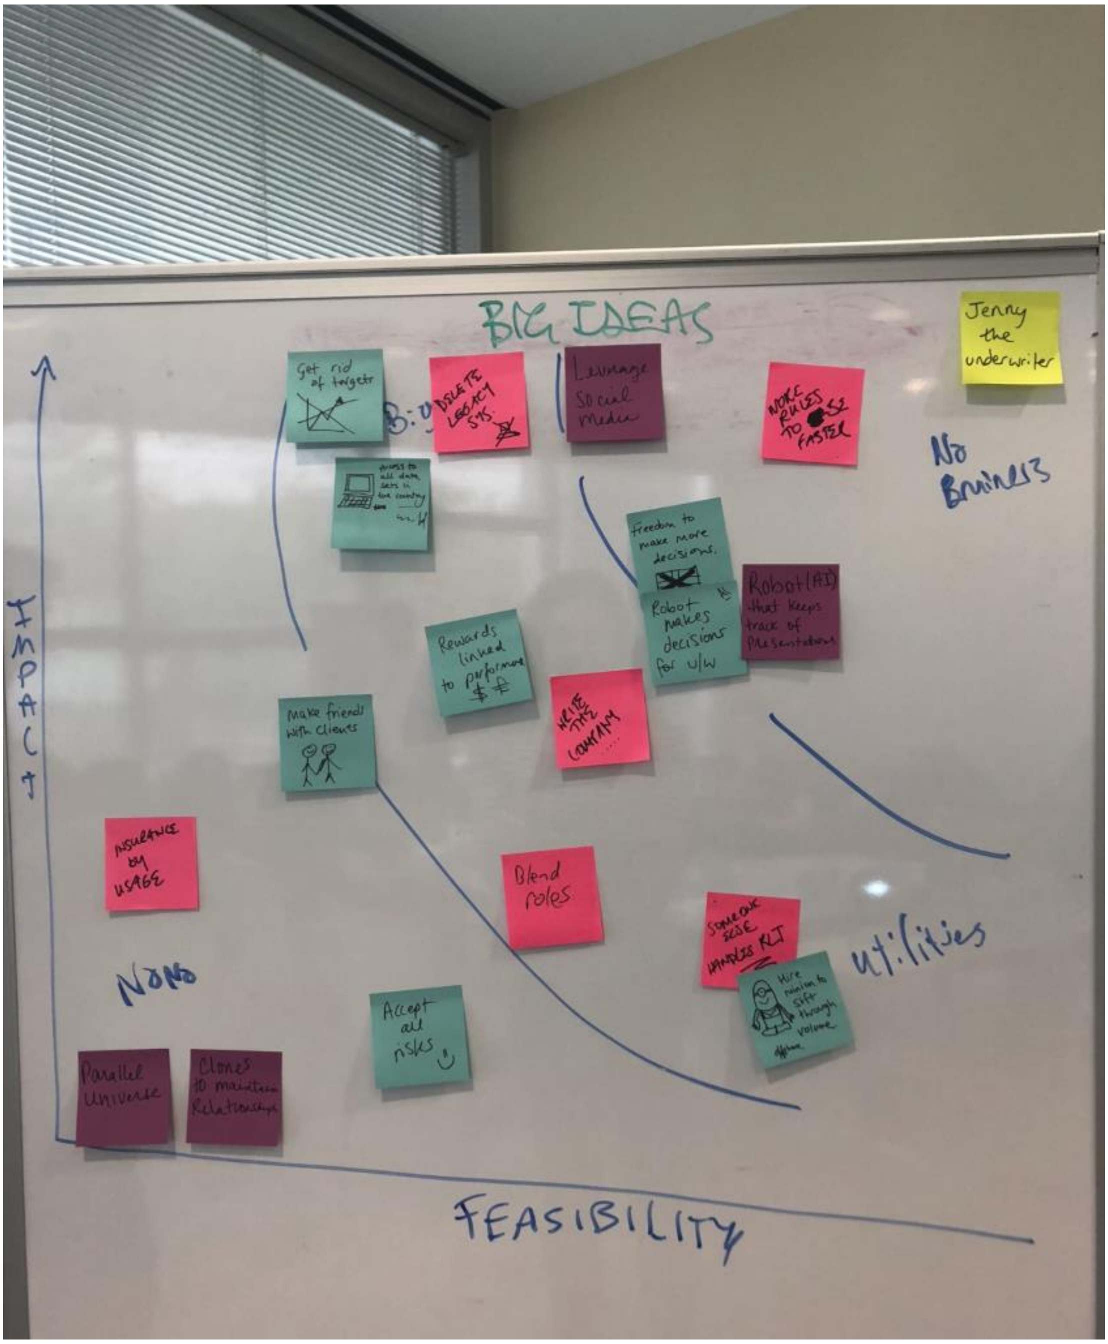

Use cases are likely to evolve over time, so the race is on to innovate. As Machine Learning and Deep Learning AI is built with data, we hope that this document has been helpful in suggesting potential use cases as an appetizer, but there must be vast potential for Insurance Companies to put their data to use with AI.

6. Acknowledgements

The following people and organisations were instrumental in helping us put this document together. We have listed their name, contact details and a description of the services they provided:

Mike Blake, mike_blake@uk.ibm.com: Mike hosted and facilitated a fascinating session on the Internet of Things and its practical application at the Hursley Client Briefing Centre in the UK. He also gave us a tour of the IBM Museum. Highly recommended!

Nelson Fernandes, nelson.fernandes@planet.de: Nelson kindly flew over from Germany and met us at the Hursley Briefing Centre. Nelson provided us with a comprehensive overview of Argus Search and Planet's adoption of IBM Power and their plans to adopt natural language processing as they build their capabilities and expand their solution.

Kevin Lancashire, klancashire@uk.ibm.com: Kevin is a Technical Specialist for IBM Analytics and provided us with a brilliant overview of Watson Studio Local and Data Science. Kevin has extensive SPSS knowledge and really helped us understand the art of the possible with a data exploration related to an Insurance data set.

Cathy Cocco, ccocco@us.ibm.com, Desiree Strom, dstrom@us.ibm.com, Cassie Yellin, cyellin@us.ibm.com: Cathy, Desiree and Cassie are the team at the Austin Briefing Centre and facilitated all activities for us in Austin. Cathy put together a comprehensive and interesting agenda for us and ensured everything ran smoothly. This part of the internship was absolutely essential to much of the content in this document. We highly recommend a trip to the briefing centre!

Tony Sawan, tsawan@us.ibm.com: Tony provided us with an interesting and fun tour of the Power Development Lab in Austin.

Rajaram Krishnamurthy, rbkrishn@us.ibm.com: Raj provided us with a really interesting presentation on enterprise AI architecture and we have used some of his slides in this document. Thanks Raj! Brilliant session.

Indrajit Poddar, ipoddar@us.ibm.com: IP gave us a fantastic presentation about Watson Studio Local on IBM Power, which is IBM's enterprise Data Science workbench and gave us an overview about its capabilities and architecture.

José Paez, Jose.Paez@ibm.com, Gabi Campagna, Gabriella.Campagna@ibm.com: José and Gabriella provided us with a fantastic workshop themed around ideation for AI using IBM's Design Thinking practice. This was a fun and illuminating engagement and we all found the session inspiring.

Mike Hollinger, mchollin@us.ibm.com,
Carl Bender, cabender@us.ibm.com,
Ana Santos De Leon, aesantos@us.ibm.com: Mike, Carl and Ana gave us a fantastic overview and deeper dive of IBM PowerAI Vision. They were helpful and extremely

knowledgeable and provided us with a comprehensive overview of current and future plans for IBM PowerAI Vision, the IBM Power Systems point and click solution for building image object and categorisation Deep Learning models.

Ploy Buraparate, ploybura@us.ibm.com: Ploy gave us a brilliant tour of the Austin Design Studio!

7. Document Authors:

Becky Humphreys, rebecca.humphreys@allianz.co.uk: Becky is an IT Business Systems Analyst working for Allianz based in the UK.

Sam Jones, sam.jones@uk.zurich.com: Sam is an IT Architect working for Zurich Insurance in the UK.

Mark Woolnough, mark.woolnough@uk.ibm.com: Mark is a Power and AI Architect working for IBM Systems in the UK and Ireland.