

Straight-Through Processing and Exception Triage Intelligence: AI-Driven Frameworks for Insurance Claims Processing Efficiency Enhancement

Dr. Derek Lichti, Professor of Geomatics Engineering, University of Calgary, Canada

1. Introduction

Claims processing is a fundamental task for the insurance industry and is pivotal in maintaining healthy payer operations. Health insurance companies alone continue spending millions of dollars primarily to manage the high volume of claims received and adjudicate claims from various channels like customer enrollment, doctor billings, laboratories, hospitals, and pharmacies. Over the years, the insurance industry itself has been facing challenges primarily from changes in the proportions of new customers, complexities in products and services, and challenges converting existing customers to new products and services. Auto-adjudication and straight-through processing of claims is where claims are automatically converted to the payer's information system, and a payment or denial is made without manual intervention. An automated system that ensures accurate claims processing can significantly improve operational efficiency, increase customer retention, and enhance customer satisfaction.

Scientific developments in the field of artificial intelligence, machine learning, and data science are universally acknowledged to address these prevailing and emerging industry challenges. This study aims to investigate the important aspects of artificial intelligence, machine learning, and data science, and artificial neural network applications to build predictive models for claims processing. Artificial intelligence and machine learning are seen as the primary technologies that can significantly alter the traditional way of conducting business in just about every industry. The larger aim of the study is to tap the latent potential of this revolutionary technology for improving efficiency in claims processing within the insurance sector. The study is expected to add value by exposing the insurance companies and health insurance regulators to already operating algorithms in the field, evaluate the outcomes of the new algorithms, and

promote learning from existing claims data for supplemented underwriting, fraud detection, and settlement of claims.

1.1. Background and Significance of Claims Processing in Insurance

Claims processing is one of the major operations in every insurer's value chain, as it directly impacts both insurers and policyholders. Manual claims processes are labor-intensive and time-consuming, leading to inefficiencies such as high processing costs and low levels of customer satisfaction. The efficiency drivers of these discrete claims processing steps cannot often be traced or improved using traditional business intelligence tools due to the complex configurations of legacy claims processing software. The growing complexity of insurance products has led to a high volume of claims, making the process more challenging. In order to stay ahead in the market, undergo digital transformation, and simplify complex and intricate claims processing steps, insurance companies have begun to look at leading-edge technologies such as business intelligence, big data technologies, data science, machine learning, and artificial intelligence.

The biggest cost driver for insurance companies is high human workloads in claims processing and adjustment processes. Most of the systems used by insurance companies have not been automated. Significant system automation optimizes the claims process, reduces operational costs, enhances resource allocation, and minimizes the processing times of claims. The delayed processing of a claim can often cause destruction and sometimes offense. Often, if an insurance company delays processing claims or does not pay a claim amount, it can lose significant emotional and legal capital. Customer retention is critical due to today's increasing competitiveness in the insurance industry, and companies must ensure that their brands remain a priority among varied clients' choices. A negative claims experience can seem more normal than typical in cases where an insurance company takes a long time to agree on a resolution claim. As a result, poor quality of claims service would lead to the loss of a consumer in addition to the primary condition that triggered the demand for a credible resolution claim. Regulatory organizations are trying to raise business claims to a higher level of expertise and effectiveness. Regulatory bodies are involved when a number of reports are recorded containing exceptions and compliance with research and compliance demands.

1.2. Role of AI and Machine Learning in Claims Processing

AI has transformative potential in various areas and can serve as a "force multiplier" in improving operational efficiency. AI relies on four fundamental principles: massive calculation power, big data, adaptive algorithms, and domain expertise, which can be applied to enhance various operational functions. AI technologies can assist in analyzing data from a myriad of internal and external sources to identify patterns, generate insights, and keep up with market dynamics. Predictive models, built on historical data, can forecast future states far quicker than a manual or semi-automated process. Automating operational decisions guided by a set of rules can free up staff to concentrate on more complex problems and exceptions. Some products are particularly good at mimicking the human brain in continually learning and addressing evolving behaviors. When faced with a previously unseen pattern, these systems analyze details and decide what action to take with the aid of a human operator. Today, an insurer needs to include AI in a combined offering to remain competitive.

Machine learning can also use data to detect insights but has the capability to adapt to new patterns. This capability is particularly useful in the detection of fraud. Drawing a straight line between innovation and the business results associated with operational excellence in claims processing is an easy thing to do. People get much too excited about product-driven innovation, as important as that is. Faster claims payments—or declining ratios of claims handling expenditure to incurred losses—are where the real action is: a direct line drawn between innovation and real business results. In fact, efficiency gains from higher automatic processing rates can account for as much as 50–70% of the payback of a large post-merger technology project when a new system should primarily focus on growth, markets, and underwriting. With the advanced capabilities available today, insurers are equipped to deliver an efficient, top-notch claims service to customers without unduly exposing themselves to new types of risk.

1.3. Purpose and Scope of the Study

Purpose of the Study While the previous section has presented a broad overview of current and anticipated future capabilities to improve the claims handling process through the use of artificial intelligence and machine learning, the objective of this study is to quantify these improvements. The research will address the following proposition: AI, based on both NLP and OCR, is capable of increasing the efficiency in unstructured

claims handling by 50% in terms of time and/or cost and increasing the accuracy compared to the current intercoder agreement. Scope of the Study This study will focus primarily on three areas. Firstly, it is necessary to understand the current technology available for introducing and applying AI capabilities for claims processing. Secondly, a realistic view on the exchange of information between technology providers and insurance companies is required, allowing the implementation of AI. Finally, the outcomes of this study will be focused on the improvement of claims processing efficiency from the availability and use of AI, particularly around the improvement in the speed and accuracy of the process. Other proportionate or lesser efficiencies need to be considered within a detailed implementation and impact study. Some potential areas that will not be addressed within this current scope are briefly discussed. Standing alone, a purely technological study of AI in the context of insurance claims processing could be the focus of an academic research project. This study, however, prefers to compare the 'as-is' and 'to-be' states, building on the existing knowledge base and ensuring the output and recommendations are relevant and beneficial to industry stakeholders. Expected Outcomes and Research Impact Measurable outcomes will include the expected claim processing times for the current, to-be defined, and potential ways to finalize the claims handling workflow. The standard operating procedures for adjusters and claims handling workers to cover the newly defined workflow will be delivered as handbooks. Finally, some recommendations for future research, industrial adoption, and further investment in a lengthy, collaborative project will be identified.

2. Foundations of Machine Learning

Machine learning is an important part of the discipline of artificial intelligence (AI) and concerns itself with the aim of finding patterns in data that can be used to make inferences in new data. While machine learning has a mathematical underpinning, the relevance of AI technologies to insurance has not been found solely in actuarial and statistical practices; the ability of AI to understand and interact with language, images, or predictions relevant to sensor drives applications in underwriting, claims, distribution, and other business areas. Before we discuss the application of machine learning today, let us go over some basic terminology. Terminology is the language of a technical area, and using it properly allows participants in the discussion to understand both other participants and their own thoughts better. Machine learning can be divided into three main categories regarding the role and type of human-machine interaction:

supervised, unsupervised, and reinforcement learning. Often, it is ambiguity in objectives or the availability of trained data that drives the choice of algorithm, and different applications may necessitate different types of learning. Given their respective strengths: supervised learning methods in claims are useful for identifying those claims which deserve closer investigation, unsupervised learning methods can be used to group together similar data, e.g., clusters of policyholders, and reinforcement learning methods might be better suited to optimizing what to do next. Supervised learning in claims might be used to work out which claims are fraudulent, which warrant closer inspection because they are more likely to be fraudulent, and which can be paid without further activity. In the first case, historical cases where we know the outcome can be used to train the model; in the second case, a new set of claims could be grouped based on characteristics, while in the third case, a reinforcement learning model could, based on the latest data and information at its disposal, recommend the next best action. Machine learning is a rapidly evolving field, and its work is diverse, with research into the methods typically presented at annual conferences. It used to be that the knowledge of expert systems had to be hard-coded, almost as if we already knew the answers. Now there is substantial potential to solve some unanswered questions. Despite advances in science, we have yet to understand all of the practical implications of generative models and deep learning, and both have only received attention in the last two years. Claims processing is ripe for the application of machine learning, and the initial steps are straightforward. While some of the models described here can be complex and computationally demanding to fit, external vendors are also able to support and provide models under agreed service-level agreements. In the future, the first two steps are preventative; accidents occurring in the first instance would eliminate two or three out of the three opportunities for fraud. AI systems could help customers with preventative maintenance or route them towards repair centers where their house or car is likely to be fixed properly.

2.1. Basic Concepts and Terminology

Due to the complex and, to the uninitiated, cryptic language inherent to the field of machine learning, it is perhaps best to preface this with a clear delineation of terminological conventions. In the world of computing and data science, models are representations of functions that take inputs, called features, and produce output values. The process of "training" a model involves presenting the model with various inputs,

"features," and revealing to the model the outcome, called the label. The collection of known feature-value examples and their matching known labels can be referred to as "training data"; it is by seeing many of these example pairs that machine learning systems learn. As in any practical application of statistics, models are only as reliable as the data used to train these models. Although the technical requirements of data quality can become quite complex, in short, machine learning models only learn to produce good output patterns when good patterned inputs are provided.

Once a model has been carefully trained to achieve its highest potential, it will need to be tested, run with a new set of input, known as "testing data," to see if the patterns it has learned are being accurately applied. One of the main distinctions between training data and testing data is that while models have seen the training data during training, they have never been shown the testing data, and indeed the testing data never acts as a part of a model's learning process. Thus, testing data globally represents an entirely external dataset, whose main purpose is to simulate how the model would perform given entirely new and unpredictable real-world examples. Data can disrupt the world of machine learning not only through the misprovision of low-quality examples but equally through an overabundance of low-quality data or by providing data that presents the wrong examples. Neural networks, for example, can demonstrate model behaviors known as overfitting and underfitting, where the models may lose the ability to generalize from training to testing data by embedding too much training data noise or being overly conservative with what patterns they learn, respectively.

2.2. Types of Machine Learning Algorithms

Based on training data, machine learning algorithms can be classified into three types: supervised, unsupervised, and reinforcement learning. In supervised learning, models are trained using labeled data to learn the mapping function between input features and output labels, such that the trained model can replicate that function to make predictions on new unseen data based on input features. The practical cases of regression and classification are popular types for prediction purposes using supervised learning methods. Supervised learning algorithms can be used in numerous applications where predictions are needed for the outcome of unseen examples based on the trained model, such as in marketing, credit card approval, fraud detection, medical diagnosis, and risk evaluation. In unsupervised learning, the goal is to discover hidden structures or

patterns in the input data, which means the data is not labeled. Clustering of data is one of the practical applications of unsupervised learning, which groups the dataset records into clusters based on some similarity or dissimilarity measures defined between record data points. Another form of unsupervised learning is known as association, which finds interesting relationships between the dataset attributes. Based on different patterns or trends, unsupervised learning techniques can be applied to reveal useful information about the input data characteristics. Reinforcement learning is very different from supervised learning, where an agent learns by trial and error to make sequences of decisions based on responses by interacting with its environment. In reinforcement learning, valuable feedback about the quality of the agent's action is not always provided. For instance, if an animal is learning to adapt to its surroundings, the reward of actions it has learned is likely to be delayed. Starting from an initial observation, reinforcement learning requires an agent that notices its environment, selects an action, and receives feedback at discrete time steps. This approach makes the agent part of the decision-making processes required to achieve a problem's or task's primary goal. Healthcare, finance, and robotics, to mention a few, make use of reinforcement learning techniques for decision making. For periodic payments or claims, a suitable learning model can be selected that improves automated learning time and builds models based on the prediction task in performance evaluation.

3. AI Applications in Claims Processing

Artificial intelligence is poised to transform many areas of business, but in few sectors does it stand to have as great an impact as in insurance. This is particularly true for claims processing, the area of the insurance value chain where customer reputation is "won or lost." By automating routine tasks in the processing of claims, AI has the power to change insurance companies from a lean perspective. That's because companies can effectively reduce losses or the cost associated with avoidable, unreasonable claims processing expenses. For example, AI can "read" and extract the necessary information from accident reports, claim forms, estimates, bills, etc. Furthermore, AI systems can assign an exam lift classification at claim opening or review, identifying needed documentation by vehicle damage.

AI systems can detect fraud by assigning a fraud probability index. Applied across the entire portfolio, this kind of process can reduce a company's loss ratio substantially.

Automating the detection of fraudulent claims provides strategic benefits to insurers. The first advantage is reduced costs, as fraud link rate detection performance is enhanced and fraudulent claims go down to the industry average. The second major advantage is that applying this degree of automation in fraud detection can minimize the loss ratio radically. The investigation of fraudulent claims costs significantly, and so reducing the overall share of fraudulent claims is a way of preventing indirect and potential costs to companies. The savings mainly result from preventing any "unreasonable" claims processing, which otherwise occur. Additionally, automating time-critical business decisions, for example, at the claim adjudication phase, expedites payouts, which enhances the company's public image, as well as the policyholder's perceived service value.

3.1. Document Classification and Extraction

Volume of Documentation Insurers tend to a vast array of policies each and every day; each activated policy generates several documents such as policy data, claims forms, repair invoices, photographs, police reports, loss of earnings assessments, legal invoices, general correspondence, etc. If these claims are contested, or if claimants do not recover within 2 years of the accident date, it is likely that additional paperwork will be generated before resolution, in the form of medical records and hire agreements, engineering reports. Once a claim is ready for the insurer to make a decision, many more documents can typically be generated, especially from law firms. Average claims take about 3.5 years to resolve in interests court, and the life of a case will typically be 7-9 years and sometimes longer. Consequently, excellent document management systems are needed. Document classification is a relatively simple way to automate the splitting of documents into categories, and most small operational improvements are likely to arise in the areas of data recognition and document relevance classification, i.e., triaging claims emails automatically. Data entry or re-entry; otherwise, the entire process has to, in the main, take place manually, including the scanning, the reading the text, and entering data manually. The various elements are expensive, particularly the human labor element. Data Recognition Technology In an era of developments in computer technology, an alternate solution is data-recognition technology, such as Optical Character Recognition and Natural Language Processing, whose purpose is to work from the electronic image itself, rather than the individual elements of data information with manual handling. Keyed data entry is typically 95% accurate because human

operators have limited attention spans and may key in wrong information, but also because electronic images generate noise. Incorporated into a document management system, therefore, and using rules of group acceptance, or where the meaning of data tells the system that it is wrong, the requirement for manual handling can be reduced. Claims information must, however, still be verified, particularly when there would be a consequent policy liability. A brief comparison of the currently available AI-driven tools can be found in the text below. A Case Study Another successful program is at a major insurance company, which handles over 2 million documents each month. The company boasts that they have reduced the time it takes to handle urgent accident and repair claims from 9 days to 24 hours with their document management system. Also, documents usually take 6 days to get around all the desks when it went by one of ten possible individuals in the system. Using the system, priority documents can then be sent immediately to the next desk, usually with a turnaround time of 1 day. Linking the system to call center staff has also meant that staff have not had to increase the number of staff in tandem with the increasing number of telephone claims. Call center staff can now answer some of the questions themselves by looking at the same electronically held documents and images as the claims handlers. Operational Imperatives One of the key operational imperatives for the insurance business is to remove unnecessary costs. The lack of consistency in the interpretation and classification of personal injuries at the first notify stage, and then during the life of a claim, is also a key irritant for the public. One of the key potential improvements comprises fully one-quarter of the entire case file. If these can be properly and fully managed and instructions properly fed through to all the interested parties on a case, then this will have a major impact on operational effectiveness and on customer satisfaction.

3.2. Fraud Detection

Being part of the insurance claims process, fraud detection significantly benefits from the abilities of AI algorithms to process big data efficiently. Fraud is known to cause approximately 10% loss of all property and casualty insurance, and it is estimated that in the US, non-health insurance fraud costs companies over \$40 billion annually. Given this financial problem, industry practitioners and academic researchers have started investigating new machine learning techniques to help identify unusual claim behaviors. The key focus of research has been on the application of machine learning and data

mining – the two paramount techniques that enable organizations to explore large volumes of datasets easily and spot problems not identified in the past.

One machine learning technique predominantly used in the claims process is expert systems/decision trees, which enable insurers to detect unusual behavioral patterns, e.g., assessing the relationship between the first notice of loss and the date of claim payment, which insurance agents cannot do easily. In such situations, AI algorithms can crawl through a large quantity of datasets and establish interconnections that have meaning. One of the major players in AI is the use of neural networks to detect fraudulent claims and credit card transactions. Neural networks are replacing previous statistical modeling techniques to verify such transactions, basing decisions on finding hidden correlations in vast datasets. Some insurers have explored the use of text mining to assess the written statements made by clients, then measuring their connotations against clients' loss history. This allows AI algorithms to quickly reveal potential fraudulent claims, reducing the load placed on claims analysts. The use of AI algorithms in the claims process has proven to increase the number of frauds detected while at the same time reducing the organization's reliance on excessive human resources. Finally, urgent action must be taken to conclude the claim if there is any indication of fraud. By acting early, the claim payout decreases significantly. A study found over a five-year period that the average saving was \$73,000 per case of fraud identification prior to decision conclusion.

3.3. Automated Decision Making

Automated decision making or automation in decisions is not something new. It has been a part of many organizations and is even more relevant in dynamic environments like the insurance industry. The insurance industry has seen a transformation in the way claims are processed. Robotic process automation is used to speed up the creation, underwriting, and even the claims collection process. Many companies do not require any human intervention for claims, as up to a certain point they are processed automatically. In motor insurance, for example, with the help of AI-based systems that analyze camera images, the damage to a vehicle is estimated and an automated decision is sent back to the insurance company about whether the claim can be accepted or not.

Indeed, the immediate deep analysis of a huge amount of related data may generate decisions and create a file without any human intervention. While this could save time

and money, the speed of AI system use may not allow enough customer-specific information to be filtered into the decision-making. Although AI speeds up the decision-making process, it could be argued that manual decisions allow some discretion, flexibility, and empathy, and allow the handling person to appropriately balance policy rules with fairness and even justice. For now, AI systems can only make decisions if they have been trained on data that has been labeled by humans. Removing humans from the claims process is also not easy—and not legal. In any case, a totally automated claims process would pose challenges in terms of preserving accountability. Automating decision making or parts of it is taken forward by automating the analysis of a product, behavior, or event. Dynamics like the effects of the ruling and the increase of the premiums could lead to an adverse impact on customer retention that could longer term compensate for the cost efficiency of rapid settlement handling.

Nonetheless, the actual moment of an automated decision can help in giving an additional competitive edge, especially in customer delight and quicker treatment. Companies are becoming very effective at quickly analyzing and adjudicating digital claims by using AI and machine learning technology to automate decisions. For example, using real-time data analysis, customer claims made through an app can be quickly processed—improving the pet owner's experience and reducing average claims handling time from several days to just a few hours. Need a replacement home right away? An app linked with another service uses a chatbot to quickly have claims for lost, stolen, or damaged personal possessions automatically approved within seconds. Decisions for a larger, unique loss will move to a human handler. Policyholders in certain regions are now able to receive payouts within five days with an artificial intelligence system that covers everything from underwriting to claims.

4. Case Studies and Success Stories

Case Study: Farmers Insurance

Due to a high volume of auto claims, Farmers turned to machine learning to reduce the resources associated with hiring adjusters and rekeying data. They also wanted to improve the speed and accuracy in making a decision on a claim. Farmers wanted to predict when auto claims could move directly to a total loss before they were inspected. Using ML, Farmers has decreased the percentage of calls that result in a new assignment by 24 percent and increased the efficiency of their adjusters, with the closure rate of new

assignments increasing by 23 percent. Farmers now auto-detect 125 percent more total loss decisions without an inspection than before, while also automatically settling these claims 1.5 days faster.

Case Study: Liberty Mutual

Liberty Mutual wanted quicker processing of property estimates and digital photo apps, accuracy in estimating cycle time, and touchless auto claims. Liberty Mutual implemented a solution and has seen a decreased cycle time for the completion of estimates. A limited release was done, and estimates started to come back as fast as 20 minutes, with an average time of one hour compared to the 40-hour cycle time of traditional methods. Moreover, the customer satisfaction score of the platform currently stands at 4.6 out of 5. Liberty Mutual also implemented a digital auto photo app and has seen an immediate drop in cycle time, comparing close to 80,000 digital app claims to a similar number of traditional photo app claims. "Technology is helping Liberty Mutual evaluate the potential claims damage faster and more accurately than ever before, and our customer satisfaction shows it. The solution integrates seamlessly into our estimation and claims platforms and aligns closely with our ever-onward goals."

4.1. Implementation of ML in a Leading Insurance Company

4.1. Implementation of ML in a Leading Insurance Company. This subsection explains, in the context of a case study, the successful implementation of machine learning in claims processing within a leading insurance company in the United States. The insurance company is one of the three largest in its segment with over USD 25 billion in annual revenues. Claims processing was a major operating issue in the company before the implementation of machine learning. It was slow, taking up to 30 days for a full cycle, and frequently customers were dissatisfied because they believed the long process was inefficient. The workers spent upwards of 700 hours a day doing the process, and it cost USD 4,782,000 annually for 4,000,000 claims processed. Due to inefficiencies in the process, the workers frequently made mistakes; a 13% error rate on claims syntax cost the company USD 249,700 annually for resubmissions and errors. The system was improved by implementing a machine learning solution.

The insurance company took several steps to involve their business unit in integrating these tools with the existing claims process. First, investment in training and

development was made by the nomination of a “Claims Processing Transformation” leader with the support of high-level executive vice presidents. Through guidance from a strategic change steering committee formed by all the company’s vice presidents, the workers were taught about the new process and received the proper motivation to implement the technology. Management hosted a town hall meeting to endorse the leadership’s commitment to the learning process before the workers began their multi-day training in person. Face-to-face interactions with management positioned this significant change as a worthwhile new opportunity for employees instead of just a new policy. The classroom sessions guided workers through an overview of the project, outlining the physical changes, the support systems, and the expected shift in operations. The company provided emotional assistance to those who had difficulty adapting. In addition to ongoing training, workers received an investment in technology through a new barcoded case identification that makes every piece of mail look unique and new. This improves the recognition process in part and package claims and results in a 99.99995% read accuracy of scan data, further reducing errors in the process. These investments amount to a process that enhanced outcomes. While the processing cycle stayed the same at 30 days per cycle, the processing output of the four million claims doubled between the old, inefficient system and the projected level of production with the new machine learning solution in place, processing claims every 21.4 days. Staff who were reassigned to other tasks performed the extra work, creating an even greater efficiency. The company assumed that if workers were reassigned to perform the claim rejections, resubmissions, and errors that should have occurred at the 13% rate, they could reduce the claims that were incorrectly filed from a 43% validation rate down to a 30% rate of claims filed. The current response time, however, to inform insureds of the resulting recovery rate might eventually decrease from 54% to 45% over the course of ninety days from when the process resumed. While workers reallocated to this “wrong jobs” proposed incorrectly calculated data for the insureds on the rejections, workers who received the proper claim data continued to process the 30th day of the 85.2% filed claims at the 100% validation rate.

4.2. Impact on Efficiency and Accuracy

As detailed in prior subsections, expert respondents generally measure efficiency gains using various metrics prior to implementing machine learning. The overall objective of most of these measurements is to reduce the time taken to complete the claims process,

to reduce the error rates associated with decision-making, and to improve the overall satisfaction of their customers. Achieving these will result in fewer inquiries, fewer appeals, and higher scores on satisfaction surveys. Notable results to be contrasted against this target performance measurement data include a large commercial insurer that reported initial claims processing times of 90 days and error rates of 30%.

They reported the ability to reduce staffing from one skilled claims adjuster for every 4,000 claims to six employees supporting the work of two claims adjusters, or 1,625 claims per adjuster using technology, up to 2,500 claims per adjuster with technology and process improvements only. A second, smaller survey respondent applied advanced analytics on their property team's work, reviewing video and building records of every catastrophe claim, and found that 15% of the catastrophe claims were fraudulent and 30% of the claims filed were such minor, repetitive claims that they were replaced by an algorithm within the first month. There was some level of file examiner review for all catastrophe claims throughout the history of the department before implementing machine learning. Almost overnight, the 15% of the most universally fraudulent claims were flagged by the technology's black box for an adjuster's review rather than all claims being reviewed by an adjuster. In one of the departments, the company used the model outputs to determine which claims had to be processed by humans based on the exceptions, essentially automating its fraud detection process. The results of this research indicate that the application of machine learning and data analytics led to a better customer experience with an improved process and decision-making accuracy.

5. Challenges and Future Directions

Clearly, however, several challenges and potential risks associated with the use of AI technologies remain, which may inhibit the uptake of AI for insurance. First, there is the fundamental issue of transparency, where some 'black-box' AI systems may be too complex to explain to customers or regulators. This lack of explainability may also be one of the factors leading to a lack of trust and possible consumer resistance to relying exclusively on AI. Second, the ever-present possibility of accident or error by the system may technically suggest systematic risk to the insurance company as the use of AI becomes more widespread. Third, there is the ongoing challenge of designing systems that are fair and non-discriminatory in the sense of being able to treat all policyholders equally and design pricing that reflects an individual's risk level.

It is also necessary to address concerns about the privacy of data and the potential use of AI for surveillance that are central to a monitoring and feedback approach, where the individual is offered, has access to, and uses driving advice for a cheaper motor insurance premium. Finally, insurers have the pressing economic challenges of competing and keeping up with the market as more and more AI systems are increasingly deployed for customer service operations and other financial products. Removal efficiency, then, becomes an important consideration in the quest to adopt AI systems. One challenge faced by insurers in developing AI-powered claims systems is that many insurance companies involve a human in the claims process to accurately judge whether a claim is, in fact, valid. This, of course, is particularly true where fraud needs to be carefully verified and agency theory.

While increased efficiency in the claims process would likely benefit large numbers of customers as well as the company itself, the displacement of human agents is a potential challenge. One future development that shows promise is in the use of AI for customer engagement through chatbots and digital avatars. The current technology has the potential to facilitate a radical innovation of service provision. As ever, the development of regulatory regimes and the ability to build environments that support safe and effective AI in insurance—and that balance the ethical considerations necessary to ensure the continued trust and overall stability of the financial intermediary—are of the utmost importance. Regulatory bodies can play a key part here by collaborating and encouraging dialogue between all relevant stakeholders. Ongoing research developments to further the efficiency of predictive analytics, including research into explainable AI and further-reaching technological change, are also key areas to reduce the potential for successful fraud. It is likely that the long-term future for predictive modeling, given the forthcoming changes in technology, is in a machine-learning framework. We endorse the view of overshoot as a serious obstacle facing the deployment of AI for verifications and argue that the ultimate limitation to AI fraud-checking is the social duty of care to those implicated and trusting in these systems.

5.1. Ethical and Privacy Concerns

Continued technological advancements in the deployment of AI and machine learning have rightly drawn increased attention to the ethical, legal, and privacy considerations of using data to drive claims handling and, more importantly, the potential

repercussions if data is used to undermine the industry's reliance on a foundation of good faith and trust. At the forefront of these considerations are data quality and integrity, bias, and unfair discrimination. The potential for predictive data, shared by the policyholder in good faith, to be misused or breached raises contractual and ethical concerns. There are valid concerns of data misuse in predicting good faith. There is a concern that "someone's DNA data might be used to inform and predict not only disease, but that you're a bad driver." In terms of algorithm protections, there is also an issue of discrimination in leveraging AI for claims. Organizations must always be transparent in requirements and decision models to ensure trust between parties. There have been spirited debates in the machine learning and automotive world regarding how compliance and regulatory requirements can stifle innovation. Unfortunately, the argument against a required regulatory regime rings tone-deaf to the wrongness of misusing a client's own data. We believe there not only should be new industry practices that all organizations utilizing AI in claims should follow, but ethical best practices that we can borrow from to ensure a commitment to privacy and fair use of data. A set of bias and fair treatment to ethical AI deployment can then be designed around these values, ensuring actual prevention of privacy and discrimination abuse.

5.2. Integration with Existing Systems

Many insurers looking to introduce advanced technologies such as machine learning and AI to their claims operations face a separate structural challenge: their existing legacy systems. The insurance industry continues to rely on traditional heavy, complex, and individualized enterprise systems built with layers of custom development that have been in place for dozens of years. The challenges associated with integrating both existing systems and the new AI systems include operating them simultaneously with no service disruption, user acceptance and buy-in, data management and quality, data duplication, and agriculture. Experience shows many companies struggle to get this right while at the same time many succeed. By starting with a blank canvas, some have used automated machine learning models to automatically generate approximately 10 million non-linear predictive models to best predict which legacy claims system would successfully update. This insight can be fed into an organization's EA strategy.

Another example is the Romans who started with local office rollouts integrating AI and legacy systems on an office-by-office basis over time. This was accompanied by

substantial training and change management across the business to get participation and engagement at each location. Officially, over 4,000 claims handlers in 60+ countries now use the system and claimants upload 75,000 pieces of information each month. Integrating AI with existing systems properly will help an insurance organization to increase the efficiency of claims handlers, allowing them to handle a higher volume of low-cost, low-complexity claims. By providing customers with a quicker claims experience, it should also increase NPS and business retention levels. However, insurance organizations need to develop the capacity, skills, and appetite to integrate with these systems so they can start to operationalize them and gain value from them. Technical integration is just part of the deployment and adoption challenge too. They will need to bring in stakeholders at an early stage to agree on the requirements and assist in an agile or lean UX process to design and validate the desired solutions for them.

The challenges will also need to be addressed in terms of ensuring the right level of data quality and correct data. Organizations must also make sure that their foundational capabilities and business operations are suitably prepared to stand up against such technological advancements. They must be 'AI-ready' by increasing an understanding of AI and what can be achieved, ensuring relevant and clean data are captured and easily accessible, having the right infrastructure and technology in place, including data management tools and BI, and having the right talent, with the skills and knowledge in AI. Applying AI to the business as a whole, a strategic holistic view needs to be considered, so these have created an AI strategy making it a part of their business process, the wider business and operations strategy, a product, and a complementary tool. The use of this leading thinking will ensure this strategy is in line with the organization's innovation agenda, making it an agile and adaptable process. Finally, insurance firms will have ethical concerns around the process; for example, they will need to articulate and understand the processes used to make sure they comply with data usage legislation and regulations, ensuring they are transparent, provide up-market, explainability, and trust in the process of utilizing AI.

5.3. Potential for Further Innovation

While AI and machine learning have started to be integrated into back-office functions in insurance organizations, the potential for further innovation should not be

underestimated. There are currently several technology advances emerging that could complement existing systems and improve operational efficiency. Blockchain is often hailed as an answer to processing inefficiencies, as its decentralized and immutable nature can ensure all parties involved in the claims process have real-time access to a single source of truth. Hyper-personalization of claims processing, driven by AI, enables the insurer to understand the customer on a much deeper level. Linking added security to rapidly changing customer experiences is real-time analytics—driven once again by AI and machine learning. Being able to respond to customer demands and new technologies within milliseconds positions insurers in a prime place to prevent churn, attract new customers, and enhance customer relationships. Today, predictive analytics are being used to predict which claims are likely to escalate, and close these out as quickly as possible to keep the claim cost down. In the future, this could be customized to understand a customer’s specific needs, because making a customer wait until the weekend to settle a small at-fault claim because of policy terms won’t prevent churn. AI will make the intelligent decision on the best outcome and product to offer and inform the customer at the time of loss. Finally, there is the cultural element of innovation in insurance. It’s not just about the technology. Insurance organizations may invest in the latest skills and technology, but they must also be willing to continually evolve their customer proposition and approach. Change needs to be about a continual shift, not just happening in one massive overhaul every decade. Evolution will also ensure that organizations can keep pace with the latest technology and continually improve processes as customer expectations move forward. Future innovation may also be constrained by the growing time it takes to deploy new technology. The technology of the future may be as valuable in terms of innovation as the technology of today is, but how easy is it to replace one with another when the life expectancy of a technology becomes increasingly hard to predict? The more difficult a technology is to replace, the more risk is associated with making an investment in it. Finally, the other major risk is the pace of change of consumer needs and perceptions. Keeping up with this could be the biggest innovation challenge for insurers. Therefore, an innovation ecosystem and the ability to innovate and implement at speed will be a critical part of the future insurance landscape. This ecosystem will center on developing an innovation that is specific to claims using the latest data and technology. It will involve mutual collaboration between figures in the insurance and insurtech industries, claims

organizations, and major figures from other industries through partnerships and alliances. The future of AI in claims is an exciting prospect that gives insurance a goldmine of possibility.

6. Conclusion

This paper highlighted how AI and machine learning can improve operational efficiency, increase accuracy, and client satisfaction, as well as reduce claim handling. Within the case studies, three forms of AI were identified: document classifiers, which use computer vision to identify and extract the content of documents in any format or process particular forms and emails, as well as natural language processing that can turn free-format notes into data. A second form involves using analysis of historic and near-time portfolios to build up a range of analytic predictive models that can highlight cases, claims, or activities being either attempted or underway that meet known fraud indicator suspicions. The third form is offering a rules-based engine that can provide automated decision-making, triggering different case management pathways according to a predefined rule set. Hence, current applications for AI in claims are largely focused on improved customer engagement and communication, reaching cost, process, and efficiency goals through claims decision-making, resource deployment, and fraud detection and management.

As with any AI implementation, there are challenges associated with its application—most notably in the areas of ethical considerations, system integration, and innovation potentials. The drive to be more innovative, to provide unique capabilities with effective solutions, is within tantalizing reach. Markets are becoming increasingly aware of the potential of these techniques. There are practical applications already that can demonstrate proof of the art of the possible. We are keen to work with the marketplace to make this happen—hence the thought leadership paper—to stimulate dialogue and address these challenges today, getting a step ahead of the game to potentially develop these capabilities for disruptive differentiation tomorrow. In summary, AI has the potential to be a transformative technology in the claims space and to change the way things are done now.