MONITAUR

WHITEPAPER

Risks of AI and ML for Insurance: What Insurers Need to Know

Contact:
info@monitaur.ai
www.monitaur.ai

Table of Contents

Artificial intelligence and machine learning (AI and ML) are transforming the insurance industry. Many companies are already using it to assess underwriting risk, determine pricing, and evaluate claims. But for all of the promises that come along with AI and ML, there are potential pitfalls that professionals must also be aware of.

As data scientists and other technicians forge ahead with these new technologies, compliance officers, risk management, and actuarial professionals must come up to speed on the fundamentals of AI and ML and fully participate in key decisions about how these technologies will be deployed and managed. If the proper guardrails and governance are not put into place early, insurers could face legal, regulatory, reputational, operational, and strategic consequences down the road. Given the heightened scrutiny surrounding AI and ML from regulators and the public, those risks may be coming much sooner than many people realize.

Most professionals are well aware of the need to develop a baseline understanding of AI and ML, but given the nature of the technology, the pace of innovation, and the relative novelty of AI and ML applications, an enormous gap still exists at most companies between the technology experts implementing these technologies and the business leaders who are ultimately responsible for controlling risk, managing compliance, and steering their organizations toward successful business outcomes.

Insurance professionals must bridge that gap, first by understanding how AI and ML can lead to new types of risks, then by putting programs in place to manage and mitigate those risks. A rapidly evolving regulatory environment will impact compliance programs as these new technologies continue to gain widespread adoption. Companies must protect themselves from compliance actions, reputational harm, and legal risk by ensuring that 1) AI/ML models are implemented within broader governance frameworks, and 2) there are checks and balances built into the entire development and implementation process.

Insurers may face significant strategic, financial and operational risks as well. If technologists and analysts fail to account for potential shifts in the underlying data that feeds AI/ML prediction models, they could generate decisions with potentially catastrophic impact. Insurance professionals must adopt a “trust but verify” approach to the decisions generated by their AI/ML algorithms. Otherwise, they are exposed to potential risks around pricing, claims processing, and other business processes to which AI and ML are applied. When taken at scale, poor predictive models could have a significant impact on the bottom line and reputation of the company.

Traditional areas of risk — compliance, operational, strategic, reputational, and legal — all have the potential to be impacted by AI/ML. Regulators, legislators, and the general public are all increasingly focused on AI and ML applications, and that situation is evolving relatively quickly. Expect compliance risk to increase over the coming two to three years, as the recent focus on data governance and digital privacy expands to predictive analytics and the potentially powerful impact it can have on people's lives. The nascent field of machine learning assurance (MLA) has emerged to address these risks and ensure that AI and ML are deployed successfully, in a manner that controls risk and maximizes the overall business value of these technologies.

IN THE KNOW: Insurance Use AI Cases

AI helps insurers make faster and more accurate decisions in these areas, often by providing decision support to front-line service professionals.

- Underwriting
- Claims & Fraud Detection
- Pricing & Product
- Sales & Marketing
- Risk Management

A Primer in AI/ML

Artificial intelligence and machine learning (AI and ML) are transforming the insurance industry. Many companies are already using it to assess underwriting risk, determine pricing, and evaluate claims. But for all of the promises that come along with AI and ML, there are potential pitfalls that professionals must also be aware of.

We often hear the terms "artificial intelligence" and "machine learning" used interchangeably. In fact, the two are related but are not directly synonymous. Artificial intelligence refers to a broad category of technologies aimed at simulating the capabilities of human thought.

Machine learning is a subset of AI that is aimed at solving very specific problems by enabling machines to learn from existing datasets and make predictions, without requiring explicit programming instructions. Unlike futuristic "artificial general intelligence", which aims to mimic human problem-solving capabilities, machine learning can be designed to perform only the very specific functions for which it is trained. Machine learning identifies correlations and makes predictions based on patterns that might not otherwise have been noted by a human observer. ML's strength rests in its ability to consume vast amounts of data, search for correlations, and apply its findings in a predictive capacity.

At the same time, machine learning has some serious limitations. First, most algorithms are not rapidly adaptive to new conditions. In other words, it makes predictions based on past patterns, without regard to any fundamental changes that could potentially impact its predictive accuracy. A machine learning algorithm trained to forecast traffic volume on major highways, for example, might be rendered inaccurate or obsolete when a new condition emerges that could radically alter people's driving habits. When something fundamental changes, such as widespread shutdowns arising from the COVID-19 pandemic, the ML model loses its ability t o make accurate predictions about traffic, unless or until it is trained with a sufficiently large sample of new data. This has implications for insurers, as we will discuss later.

A second major limitation inherent to machine learning is its focus on discovering correlation. In some cases, skilled data scientists can design models that are closer to causal in nature, but in general, ML doesn't care why two data points may be connected. It only cares that there is a correlation between them. This can lead to unintended consequences, especially with regard to so-called proxy discrimination, which can unintentionally drive disparate outcomes for different demographic groups based on factors that are only indirectly connected to race, religion, gender or other protected attributes.

We will discuss these limitations in further detail, particularly as we delve into the various risks that arise from the adoption of AI/ML technology. First, though, we'll explore the concept of machine learning models, focusing particularly on those models that are broadly applied within the insurance industry.

IN THE KNOW: Types of AI

ARTIFICIAL INTELLIGENCE

A broad group of systems that attempt to mimic intelligence as displayed by humans

MACHINE LEARNING

A model-based system that can learn without explicit programming directions

DEEP LEARNING

A type of ML model that uses multiple layers to create its own new features from raw data

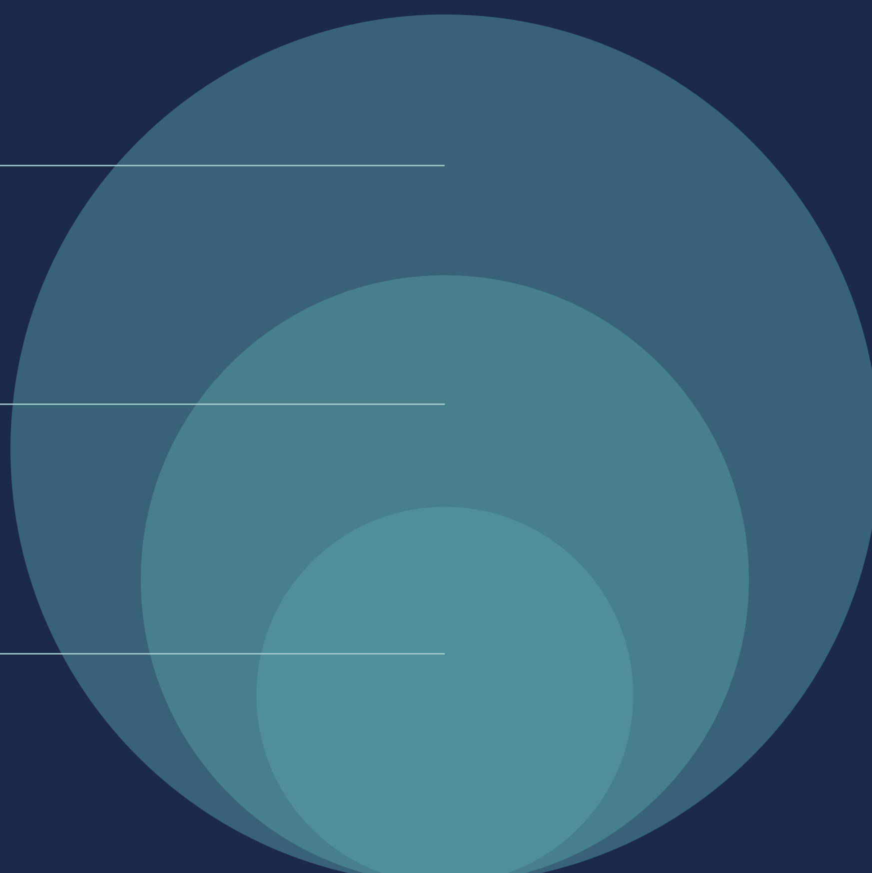

An Introduction to Predictive Models

There are three fundamental approaches to building predictive models:

Descriptive models are generally deployed with an eye toward the long-term goal of understanding causality or strong correlative connections among two or more variables. Statistical modeling and time series analysis generally fall into this category. A statistical analysis that discovers a strong correlation between smoking and premature death, likewise, is a descriptive model insofar as it establishes (or at least strongly implies) a causal relationship. Descriptive models have been used in the insurance industry for many years and are not necessarily limited to an AI/ML context. Nevertheless, they are coming under increased scrutiny as regulators realize their inherent limitations and potential for unintentional bias.

Predictive inference models are optimized to make predictions based primarily or solely on correlation, without less regard for causation. This approach can be useful when decision-makers are less concerned with the "why" of a prediction but merely need to understand likely outcomes and the appropriate actions to be taken in response to those predictions. For instance, when an insurer uses machine learning algorithms to predict the likelihood of a claim, they are generally more concerned with an increase or decrease in that likelihood than the specific factors that generated the change. In other words, they simply care about the accuracy of the prediction so that they can take effective business action. The majority of ML models fall into this category because business leaders see clear advantages in using data science tools to consume vast amounts of data and optimize their models for predictive performance.

Systems models address the phenomena being studied as a whole by attempting to determine with differential equations how various factors operate in conjunction with

IN THE KNOW: Machine Learning Ingredients

Developing a machine learning model requires the following:

- **Data**: Structured ("Tabular"), Image, or Unstructured data
- **Data Labels**: The "truth" that enables an algorithm to interpret data (e.g. cat vs dog, good risk vs bad risk)
- **Data Experts**: Individuals who intimately understand the specific datasets that the model will encounter. Data Scientists & Analysts.
- **Compute**: A computation environment, via cloud service or on-premise, with sufficient resources to process data and build ML models.
- **Tools & Software**: A solutions "stack" to process data, run experiments, and build ML models
- **Software Experts**: Data & ML Engineers, Data & ML Infrastructure Engineers, IT/DevOps

one another to produce a given result. Systems models are generally far more complex than descriptive and predictive models, and they usually require a multistep approach in order to reach a conclusion. This makes them far more complex and difficult to implement. They require significantly greater forethought because data scientists must properly identify key variables and relationships up front. Because of the high barrier to entry and other limitations, systems modeling is applied far less frequently than the descriptive and predictive inference approaches.

Business leaders in the insurance industry will be primarily concerned with the first two categories of models in this list, that is, descriptive (statistical and time-series) models and predictive inference models. Let's take a closer look at some of the risks arising from the use of these models.

Limitations and Pitfalls of AI/ML

Much of the potential concern about AI/ML applications in the insurance industry stems from the fact that predictive inference models are focused on identifying correlations in the datasets, which the models then employ in making predictions. Such correlations may reflect past discrimination, so there is a potential that, without oversight, AI/ML models will actually perpetuate past discrimination. The increase in the impact with AI/ML results from scale. Discrimination can occur without AI/ML, but the scale is much, much smaller and more difficult to detect.

Consider what happens, for example, when a life insurance company deploys an AI/ML model to assess life expectancy. The model takes into account the neighborhood in which each applicant lives and discovers that in communities disproportionately populated by people of color, life expectancies are lower than in wealthier suburban areas with predominantly white populations.

Even without any explicit reference to race, the ML model discovers a correlation that exists at least in part because of racial disparities driven by past discrimination. Economically disadvantaged communities typically suffer from limited access to healthy food, higher crime rates, and higher rates of drug and alcohol addiction. The result is that if an applicant lives in the "wrong" neighborhood, they may be assessed as a higher risk than someone who lives in the wealthy suburbs.

A predictive inference model is not concerned with causation; it is simply trained to find correlation.

Even when the ML model is programmed to explicitly exclude race as a factor in its decisions, it can nevertheless make decisions that lead to a disparate impact on applicants of different racial and ethnic backgrounds. This is known as proxy discrimination, and in many cases, it can be far more subtle than the example outlined above.

In 2019, Apple launched a branded credit card in partnership with Goldman Sachs. Before long, users noticed that women were generally being offered lower preapproved credit lines than men. They took to social media to decry what they saw as blatant sexism on Apple's part. Even the company's co-founder Steve Wozniak joined in on the discussion, pressing the company to respond to the allegations.

Unfortunately, neither Apple nor Goldman Sachs were able to immediately explain how their ML models worked. Their systems were not auditable, so they couldn't explain the differences in pre-approved credit levels except to say that gender was never explicitly part of their predictive model. Neither organization had a plan in place to respond to charges of proxy discrimination, so the result was a public relations disaster. The reputational damage to Apple and Goldman is evidence that companies deploying AI need to get ahead of such risks by having governance structures in place in advance.

As noted previously, there is a second major deficiency inherent in predictive inference models, namely that they are incapable of adapting to new information unless or until they are properly acclimated to the "new reality" by training on updated data. Earlier, we gave the example of traffic patterns before and after the arrival of the COVID-19 pandemic. Let's consider a different example this time, one which is perhaps better aligned with the insurance industry.

Imagine that an insurer wishes to assess the likelihood that an applicant will require long-term in-home care. They train their ML models based on historical data and begin making predictions based on that information. Imagine that a breakthrough treatment is subsequently discovered (for example, a cure for Alzheimer's disease) that leads to a 20% decrease in required in-home care services. The existing ML model is unaware of this development; it cannot adapt to the new reality unless it is trained on new data. For the insurer, this leads to overpriced policies and diminished competitiveness.

The lesson here is that AI/ML requires a structured process of planning, approval, auditing and continuous monitoring by a cross-organizational group of people.. Business leaders cannot afford to take a "set it and forget it" approach. In the auto insurance space alone, we can imagine multiple examples in which changing circumstances must be taken into consideration to ensure that machine learning is effective.

When mandatory seat belt laws went into effect in many locations and when airbags became standard on new vehicles, incidences of death or serious injury in motor vehicles diminished markedly. Current production vehicles incorporate automatic braking, heads-up displays, and other mechanisms that warn drivers of impending hazards or intervene directly to prevent accidents. What happens when multiple factors like this, each having some influence on outcomes, act in combination? Machine learning can only tell us that if it is trained on high-quality data that is routinely audited and updated to reflect the current reality.

The key takeaway here is that without oversight, AI/ML can generate the wrong results, leading to bad business decisions and competitive disadvantage. An intentional approach to AI is absolutely essential.

IN THE KNOW: Understandable vs Explainable Models

A dizzying array of model types are available to data science experts today. A rule of thumb: model complexity is strongly correlated with the difficulty for humans to understand how decisions are made by an AI/ML system.

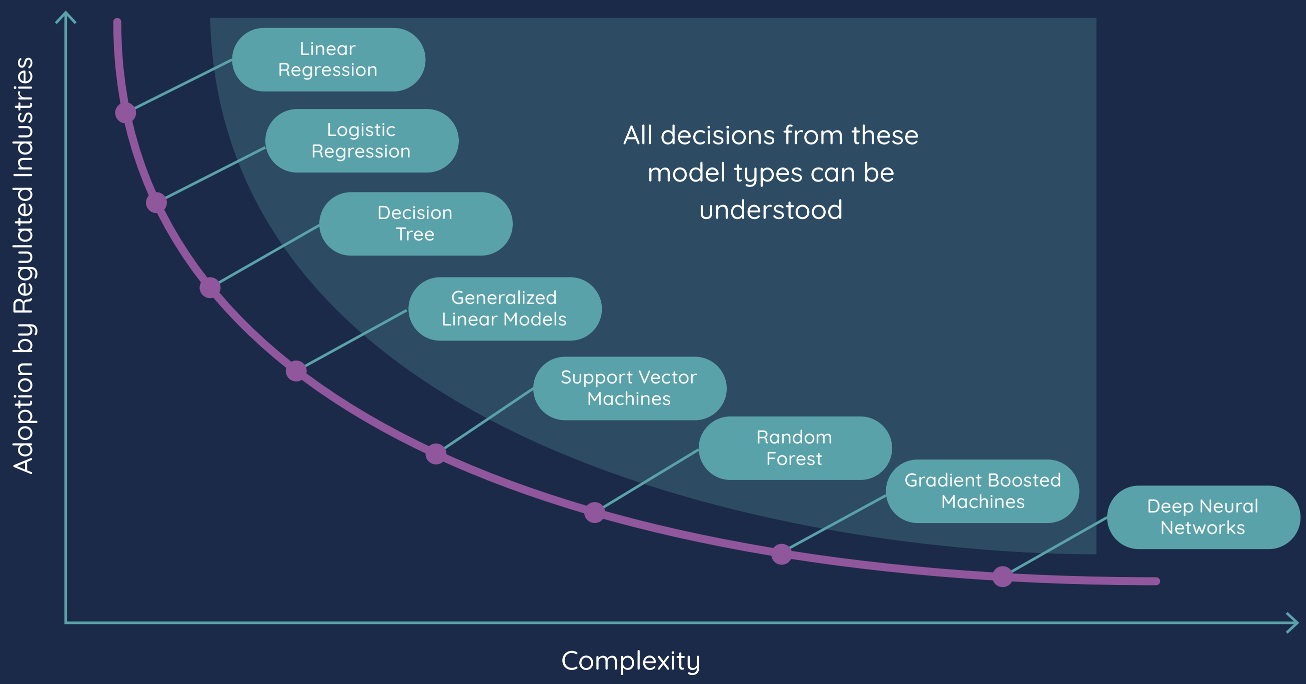

The vast majority of models deployed in highly regulated environments offer some ability for humans to comprehend how they were made through a combination of transactional auditing, model reperformance, and technology tools.

In many cases, technical teams can use careful practices to design “interpretable” models, also known as “white box” models. Emerging research into Explainable AI (XAI) holds promise for “black box” models; however, it is best to think of XAI tools as providing probable approximations of how models make decisions to expert data scientists, rather than exact explanations that are intelligible to non-technical reviewers.

Organizational Readiness for AI/ML

As with any new technology that holds great promise, a number of leading insurers have forged ahead with AI/ML initiatives, understanding that there are clear first-mover advantages to be gained. As adoption picks up steam and business practices evolve to incorporate those gains, though, it's worth pausing to consider questions of organizational readiness with respect to AI/ML. Does your company have the right people in place? Are executives aware of what it will take to make AI/ML initiatives successful? Are your technology professionals well-versed in the rapidly evolving domain of compliance as it pertains to AI and ML?

There are a range of questions that touch on the technical domain as well: Which specific business problems lend themselves to AI/ML applications? Does the company have the necessary technical resources and infrastructure in place to implement and manage AI/ML initiatives? Is the right data available to feed machine learning models, and is it of sufficient quality?

Following the deployment of AI/ML programs, insurers must ask whether their machine learning models are performing as expected and must see to it that results are monitored for bias, drift, outliers and anomalies. AI/ML systems must be auditable and controllable so that previous decisions or recommendations can be assessed after the fact, with a clear understanding as to how AI/ML models arrived at their conclusions.

Insurers must also guard these systems against both internal and external threats, ensuring that reliable controls are in place to prevent malicious behaviors that could put the company at risk. Without a structured governance framework in place, AI/ML systems stand a chance of being compromised, with potentially disastrous results.

AI and ML are already providing tremendous value for insurers, but unless leaders take the initiative to create a governance framework around these technologies that includes risk mitigation strategies, those AI/ML programs may expose the companies using them to a range of hazards, including reputational, legal, compliance, and strategic and operational risks.

Categories of AI/ML Risk

Broadly speaking, five categories of risk related to AI/ML exist that insurers should be concerned with: reputational, legal, strategic/financial, operational and compliance/regulatory. Let's look at each of these in turn.

Reputational risk

Reputational risk arises from the potential negative publicity surrounding problems such as proxy discrimination that may occur. As the Apple Card debacle illustrates, the predictive models employed by most machine learning systems are prone to introducing bias. Recently, an early adopter of AI in the insurance space suffered backlash from consumers when its technology was criticized due to its potential for treating people of color differently from white policyholders.

Proxy discrimination should be prevented whenever possible through strong governance, but when bias occurs despite a company's best efforts, business leaders must be prepared with an explanation. The technical field of "explainable AI" is gaining increased attention among data scientists in response to this risk. However, most tools for explainability are only suitable for expert technical users and create a new dimension of risk since different model types and datasets demand different explainability methodologies. As insurers roll out AI/ML, they must proactively prevent bias in their algorithms and must also be prepared to fully explain their automated AI-driven decisions.

Key questions:

1. In what unexpected ways might ML/AI model decisions impact our customers, whether directly or indirectly?
2. How are you determining if model features have the potential for proxy discrimination against protected classes?
3. What changes have model risk teams needed to make to account for the evolving nature of ML/AI models?

Legal risk

Legal risk is looming for virtually any company using AI/ML to make important decisions that affect people's lives. Although there is little legal precedent with respect to discrimination resulting from AI/ML, companies should take a more proactive stance toward governing their AI to eliminate bias. They should also prepare to defend their decisions regarding data selection, data quality, and auditing procedures that ensure bias is not present in machine-driven decisions. Class-action suits and other litigation are almost certain to arise in the coming years as AI/ML adoption increases and awareness of the risks grows.

Key questions:

1. How are we monitoring developing legislation and new court rulings that relate to AI/ML systems?
2. How would we obtain evidence about specific AI/ML transactions for our legal defense if a class-action lawsuit were filed against the company?
3. How would we prove accountability and responsible use of technology in a court of law?

Strategic and financial risk

Strategic and financial risk will increase as companies rely on AI/ML to support more of the day-to-day decisions that drive their business models. As insurers automate more of their core decision processes, including underwriting and pricing, claims assessment, and fraud detection, they risk being wrong about the fundamentals that drive their business success (or failure). Perhaps just as importantly, they risk being wrong at scale.

Currently, the diversity of human actors participating in core business processes serves as a buffer against bad decisions. This doesn't mean bad decisions are never made; they are, but as human judgment assumes a diminished role in these processes and as AI/ML take on a larger role, errors may be replicated at scale. This has powerful strategic and financial implications.

Key questions:

1. How are we preventing AI/ML models from impacting our revenue streams or financial solvency?
2. What is the business problem an AI/ML model was designed to solve, and what other non-AI/ML solutions were considered?
3. What opportunities might competitors realize by using more advanced models?

Operational risk

Operational risk must also be considered, as new technologies often suffer from drawbacks and limitations that were not initially seen or which may have been discounted amid the early-stage enthusiasm that often accompanies innovative programs. If AI/ML technology is not adequately secured — or if steps are not taken to make sure systems are robust and scalable — insurers could face significant roadblocks as they attempt to operationalize it. Cross-functional misalignment and decision-making silos also have the potential to derail nascent AI/ML initiatives.

Key questions:

1. How are we evaluating the security and reliability of our AI/ML systems?
2. What have we done to test the scalability of the technological infrastructure that supports our systems?
3. How well do the organization's technical competencies and expertise map to our AI/ML project's needs?

Compliance/Regulatory risk

Compliance/Regulatory risk should be a growing concern for insurers as their AI/ML initiatives move into mainstream use, driving decisions that impact people's lives in important ways. In the short term, federal and state agencies are showing increased interest in the potential implications of AI/ML.

The Federal Trade Commission, state insurance commissioners, and overseas regulators have all expressed concerns about this technology and are seeking to better understand what needs to be done to protect the rights of the people who live under their jurisdiction. Europe's General Data Protection Regulation (GDPR), California's Consumer Privacy Act (CCPA), and similar laws and regulations around the world are continuing to evolve as litigation makes its way through the courts.

In the longer term, we can expect regulations to be defined at a more granular level, with the appropriate enforcement measures to follow. The National Association of Insurance Commissioners (NAIC) and others are already signaling their intentions to scrutinize AI/ML applications within their purview. In 2020, NAIC released its guiding principles on artificial intelligence (based on principles published by the OECD) and in 2021, created a Big Data and Artificial Intelligence Working Group. The Federal Trade Commission (FTC) has also advised companies across industries that existing laws are sufficient to cover many of the dangers posed by AI. The regulatory environment is evolving rapidly.

Key questions:

1. What industry and commercial regulations from bodies like the NAIC, state departments of insurance, the FTC, and digital privacy laws affect our business today?
2. To what degree have we mapped regulatory requirements to mitigating controls and documentary processes we have in place?
3. How often do we evaluate whether our models are subject to specific regulations?

Best Practices for Managing AI Risk

Proactive management is essential for controlling AI risk. For insurers, that means developing cross-functional awareness of the potential exposure created by these technologies, especially with respect to compliance, audit, and internal controls.

Incorporate AI/ML into your existing governance structure. Insurers must create accountability and alignment across their organizations for ownership, management, and problem mitigation. This starts with building a baseline competency in AI/ML throughout the company, with strong cross-functional communication and collaboration at the center. Compliance officers, data scientists, and line-of-business owners must work together to ensure that AI/ML is implemented and managed with an eye to the five key areas of risk outlined above.

Stay up to date on statutory, legal, and regulatory developments related to AI/ML, both within the insurance industry and horizontally. Foster relationships with regulators to ensure that your work efforts to mitigate risk are clearly understood, and work proactively to ensure regulatory developments recognize the machine learning assurance and governance accountability and alignment efforts.

Proactively implement a machine learning assurance program. Companies must be prepared to show and explain their work, that is, to demonstrate how their AI/ML models are designed, how the organization went about selecting data and ensuring its quality, and what steps they took to proactively eliminate potential sources of bias. They must demonstrate that they have well-considered controls in place throughout the model lifecycle to protect against internal and external threats, to audit AI/ML systems for potential violations, and to take prompt corrective action when necessary.

Machine Learning Assurance (MLA) provides a controls-based framework and supporting technology to establish confidence and verifiability that AI/ML systems are functioning as expected. MLA safeguards the transparency, fairness, compliance, safety, and optimal operation of AI/ML systems.

Trust and Confidence for ML and AI

Monitaur is the premiere provider of AI Governance and ML Assurance software solutions for insurance companies around the world. We believe that, with the right approach and effective oversight, AI/ML have tremendous power to improve people's lives. Founded by a team of deep domain experts, Monitaur's software helps insurers deliver human audit and verification of models and their decisions, align cross-functional teams around assurance of models, and protect themselves from unanticipated impacts and reputational harm.

If your organization has already deployed AI/ML or other advanced models – or if you're considering it in the near future – we would love to talk to you about how Monitaur can help. **Contact us at info@monitaur.ai today.**