Whitepaper

From Governance to Insurance

Frontline Perspectives on Mitigating Corporate AI Risk

NOT IF, BUT HOW

Munich RE

About this Report

Munich Re has been writing AI insurance for more than five years. During that time, we've reviewed a diverse range of AI models and use cases, and developed a solid, proven, and extensible risk assessment methodology.

But with the advent of GenAI more than a year ago, the corporate appetite for AI risk is evolving rapidly and profoundly, and so are the tools and policies for managing this risk. To learn how risk is being managed in this fast-changing environment, we reached out to a broad cross-section of professionals working on the front-lines of corporate AI development and deployment.

In all, over a three-month period from October 6, 2023, to December 5, 2023, 32 AI governance stakeholders from across Europe, North America, and Asia took our calls, shared their insights, and thankfully did not complain when the calls ran over the allotted 30 minutes.

This report distills their experience and insights.

Key takeaways

While the topic of AI hallucinations grabs the headlines, corporate concerns about AI risk more broadly focus on bias and discrimination, IP infringement, PII compromise, cyber vulnerability, probabilistic error, and drift.

Companies are currently limiting AI risk exposure by constraining use case selection while waiting for current AI legal cases to work their way through the courts.

AI risk is viewed as an extension of technology risks that corporations have long faced and is assumed to be covered by traditional all-risk insurance. However, as more and more AI gets adopted, coverage gaps may emerge.

No AI model can be error-free and such mistakes can occur without negligence.

Even with a well-designed model, the frequency of errors may be rare but the severity of a single error may be dramatic.

Table of Contents

Introduction 4
Fig 1: Who's Worrying About AI: Interview Participants by Organization Type and Geographic Region 4

To Mitigate Risk, Start by Playing it Safe 5
Fig 2: What We're Building: AI Use Cases by Type 5
Fig 3. The KPIs Used to Measure AI Use Cases 6

Data: First Risk Among Equals 6
Fig 4: What We Talk About When We Talk About AI Risk 7

AI Innovation and Risk Mitigation 7
Fig 5: The EU AI Risk Framework 8

Residual Risk and Insurance 8
Fig 6: Established Insurance Policies and AI Coverage 9
Fig 7: Risks that Affirmative AI Covers Should Cover 9

Minding the Gaps 10
Fig 8: Assessing AI Performance: What Data, Which Type of Model, and How is it Monitored 11

“Open the pod bay doors, HAL.” “I am sorry, Dave. I'm afraid I can't do that.”
—Stanley Kubrick’s “2001: A Space Odyssey“

Introduction

The horror of the rogue machine has haunted humankind since the early days of the industrial revolution, but the theme of technology running amok may go back beyond the myth of Prometheus and the gift of fire. There seems to be something both frightening and abhorrent about a tool designed to empower human endeavor turning against its master and causing loss and harm, whether it’s Shelley’s Frankenstein or an autonomous vehicle.

From a risk governance and mitigation perspective, the perennial question is whether the harms wreaked by novel inventions differ in kind or intensity from the well-known and well-understood risks posed by established technologies and processes: for there is no machine without a failure rate and no process immune to error.

Fortunately, the discipline of quality by design in the new product development process is well established. In the case of the software release life cycle, the closest precursor to AI, the established governance framework and structure from requirements gathering to design, data sourcing and implementation through testing, release, monitoring, and maintenance can be applied to AI models and use cases without wholesale modification. In other words, tools and practices to mitigate the risks associated with AI—e.g., bias and discrimination, IP infringement, PII compromise, cyber vulnerability, probabilistic error, and drift—are well-defined.

As an insurance partner to companies that seek to build and deploy AI models, the task we set for ourselves was first to understand where the current gaps, if any, in process and technology lie within the AI governance framework. Then to learn what residual risks remain after all best practices have been applied, and how best to deal with any economic, reputational, and/or litigation repercussions that that may ensue from erroneous AI performance.

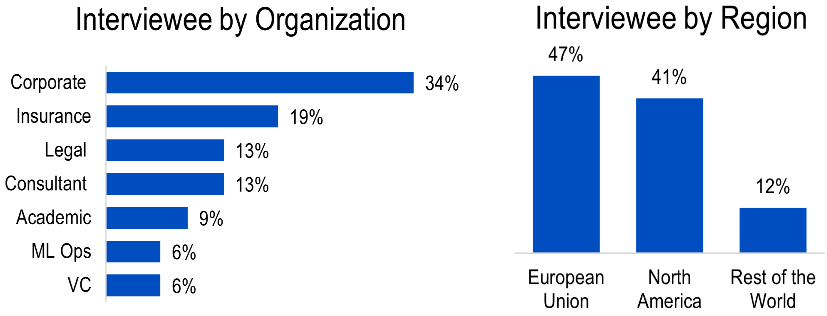

Fig. 1: Who’s Worrying About AI: Interview Participants by Organization Type and Geographic Region

To tackle these questions, we casted a wide net and interviewed a broad range of governance stakeholders. We started with our peers at brokerages, carriers, and wholesalers to the risk managers they serve, then over to the engineers and product developers making AI design decisions, upstairs to the CFO, COO and CISO offices, then broadened the discussion to inside and outside counsel, with interesting conversations with consultants, policy analysts, and a number of innovative start-ups in the ML Ops space.

Here's what they told us.

To Mitigate Risk, Start by Playing it Safe

If as Alexander Pope advises, "Fools rush in where angels fear to tread," then despite whatever frenzy has whipped up in the media over AI, there are few fools to be found in the realm of AI risk governance. As one of our Big Tech respondents put it, "Reality is different than the survey data which supposedly comes from the C-Suite. They say they want to achieve competitive advantage using Gen AI; but really, they want to gain experience today that will be useful later—without blowing anything up."

One basic way to de-risk this experiential phase of AI deployment is to avoid high-risk use cases that could lead to significant losses or reputational damage. Instead, a company can greenlight proven knowledge base and document retrieval use cases where legacy automation technologies are already deployed, so the error and risk factors are well understood. In the words of an Fortune 500 product development lead, "The safest course is to apply AI to improving established automation use cases—for example in the call center, leveraging data to improve response time and accuracy is a win. All the information is already there, but nobody can access it all efficiently. A marginal gain of even 1% to 2% is probably worth the investment, and the risk doesn't necessarily increase given the legacy customer experience was not great to begin with."

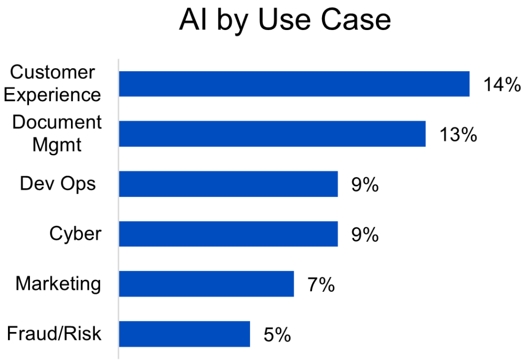

Fig. 2: What We're Building: AI Use Cases by Type

For higher value use cases, corporates lean towards a conservative approach to AI adoption, enlisting experts to spot-check the AI's work and ensure that it is functioning correctly. As the head of one European Large Language Model (LLM) vendor put it, "Every use case involves certain risks, whether it's automating the legal due diligence contract review or customer service assistance. There is always the possibility that the AI may underperform, not find certain things, or have an errant opinion. So, we advocate for keeping a human-in-the-loop (HITL) and try to design conservatively."

The issue with reliance on HITL as a risk mitigation measure is that dampens the productivity and scalability of an investment in a given AI use case. Even in these early days, the metrics of corporate investment in AI trend towards the quantitative, and in particular, process optimization/efficiency. An InsureTech executive noted, "Equal or greater accuracy in less time is the measure of any AI automation project". This thought was then echoed by another Fortune 500 product development lead: "Specifically in software development, using Gen AI is about responsiveness—accelerating time to market over the product life cycle."

This can lead companies to the paradox of evaluating AI use cases based on efficiency, which the insertion of the human in the loop impedes. A US IP attorney points to the case of the autonomous vehicle vendor, Cruise: "They were checking AI decisions every 5 or 6 minutes, which is not sustainable, given that companies are looking to automate the processes that will yield the most cost-savings." Acknowledging that the HITL approach is not cost-effectively scalable over the long-term, for now, the humans in the AI's loop would seem to have a measure of employment security as corporations scale the AI learning curve.

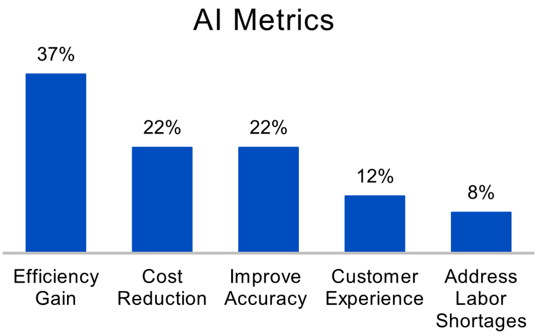

Fig. 3: The KPIs Used to Measure AI Use Cases

Data: First Risk Among Equals

Another paradox in the development of the AI market currently is that the potential risks of deployment may be easier to calculate than the potential rewards. "How are AI risks any different from what we already face?" an Operations VP asks. "These risks are all existing problems, maybe just exacerbated by AI. If you run a business, then you already have an established risk appetite. 'How do you manage risk?' is the question."

In fact, running our interview notes through a word cloud generator, the list of worries probably looks a little different than that for a conventional product launch view, perhaps with the exception of "hallucination." It is also perhaps no coincidence that the same frequency analysis reveals "data," whether company privileged or protected third party, ranks right alongside the headline-grabbing topic of hallucinations. "Only developers think about hallucination," a corporate governance lead notes, "But it's the data privacy issues that are happening in non-regulated, non-policed GenAI environments that keep us up at night."

Not surprisingly, the perceived intensity and urgency of a given risk type varies according to the use case. As one European COO told us, "Risk completely depends on the use case. If we are dealing with people, the issue is bias. For autonomous driving, on the other hand, prevention of cyber attacks is extremely important, where data protection is most important in using LLM models for decisioning." In addition, the use of LLMs brings with it a heightened concern about unintended consequences. "You might inadvertently violate an NDA merely by using an open model, since the lineage of data used in open models can't be guaranteed," an IP attorney points out.

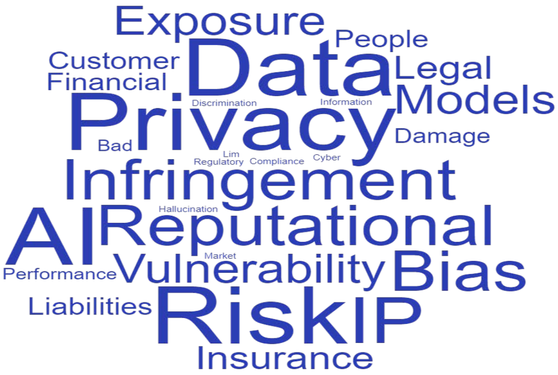

Fig. 4: What We Talk About When We Talk About AI Risk

In fact, the corporations we interviewed by and large are avoiding open AI models altogether and implementing strict output filtering and leakage monitoring processes. One development operations engineer commented, “Whenever we deploy an application in the organization it is within a secure network,” but here again the conservative approach creates a paradox. As a CISO noted, “To realize any improvement in data quality and delivery, you need to open up data to APIs. This is inherently high-risk and contradicts standing CISO guidance. But AI insists you open up the silos.”

AI Innovation and Risk Mitigation

Whether intentional or not, the corporate tactic of limiting risk exposure by constraining use case selection aligns with the EU AI risk framework, which classifies AI uses according to potential impact on human rights, health, and safety. Not surprising, the provisional EU AI Act subjects AI use cases in the high-risk category to a higher level of compliance scrutiny surrounding data governance, model accuracy, cybersecurity, documentation and transparency, and the overall risk management process in place.

For the majority of experts we reviewed, this is viewed as a welcomed development. As one European chief technology officer put it, “There are probably too many best practices regarding organizational AI governance, development and deployment, and benchmarking, along with any number of tools in the software stack.” A US insurance brokerage reported, “We already have some US corporate clients who conform their policies to EU privacy standards. Companies like the assurance of standards.”

Given the emerging patchwork of regional and national policies, companies are already looking forward to the advent of ISO standardization. “We expect ISO/IEC JTC1/SC 4 to be as impactful for AI as ISO/IEC 27001 was for cyber security, as it will be the first global standard that an organization can be certified against. To say you complied with that standard will be really useful for defining best practices,” the CTO told us.

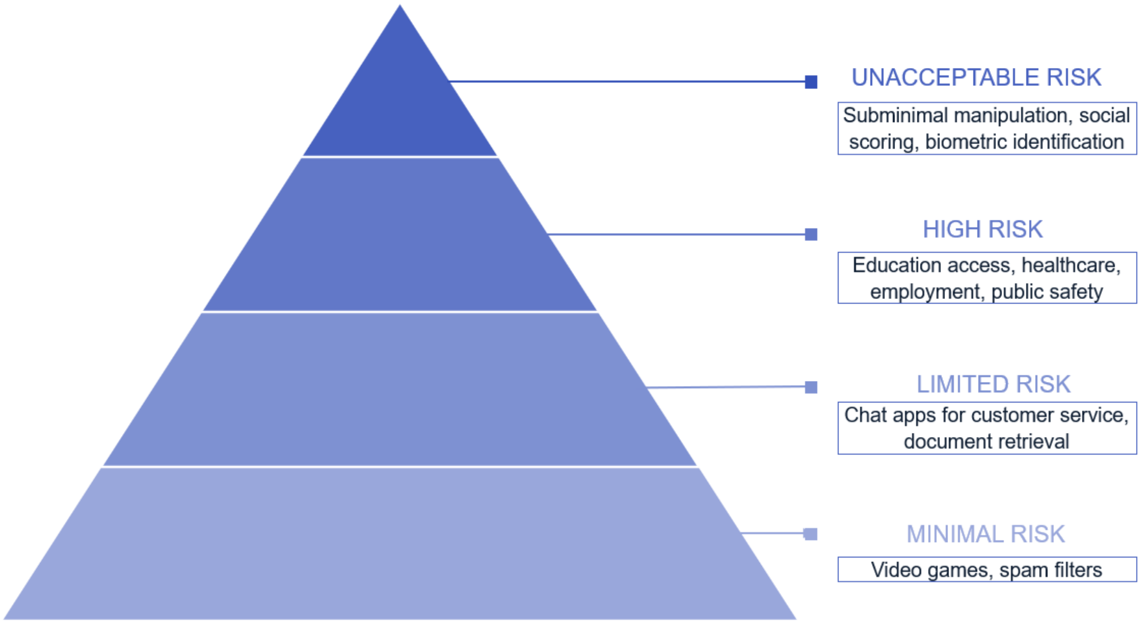
UNACCEPTABLE RISK
Subminimal manipulation, social scoring, biometric identification
HIGH RISK
Education access, healthcare, employment, public safety
LIMITED RISK
Chat apps for customer service, document retrieval
MINIMAL RISK
Video games, spam filters

Fig. 5: The EU AI Risk Framework

At the same time it is recognized that as mature and robust as a given governance framework may be, risk prevention and avoidance measures can never be infallible. A Fortune 500 Software Development Engineer warns, “We are dealing with probabilistic models. There are internal tools and processes to act as a safety net, but they are not perfect.”

Risk transfer takes over where governance and management reach their limits. As a US IP attorney put it, “As a lawyer, you are looking for how a risk can be shifted off a corporation. First off, we want the client to shift AI risk over to the technology vendor or onto end users, and then we want to make sure that they have the available resources to indemnify themselves for any liabilities they are held responsible for.”

To this end, the copyright infringement indemnification announcements from Microsoft, Google, and Amazon have clearly resonated within the corporate sphere, despite concerns about limitations, such as Google’s reported exclusion of models that have been “fine-tuned” by customers using internal data. That notwithstanding, IP is only part of the residual risk picture and it is unclear whether future regulatory policy or case law will allow AI application providers to transfer IP risk completely over to foundation model providers, leading our IP attorney to conclude “Whatever risk you can’t shift, that’s when you look to insurance to distribute it.”

Residual Risk and Insurance

For many, AI risk is viewed as an extension of the technology risks that corporations have long faced. Therefore, it is reasonable to proceed with the assumption that it will be covered within the phalanx of professional liability/indemnity, product liability and tech E&O, and cyber coverages that corporation currently maintain, especially for corporations limiting their AI deployments to low-risk use cases. “Right now we are not making important decisions with AI with significant financial exposure”, a corporate AI innovation director explains. “But with big damages on the line, that’s when you would look to specialized insurance.” A similar sentiment was shared by European scale-up founder: “There’s clearly a need for an insurance that would cover an economic loss caused by AI.”

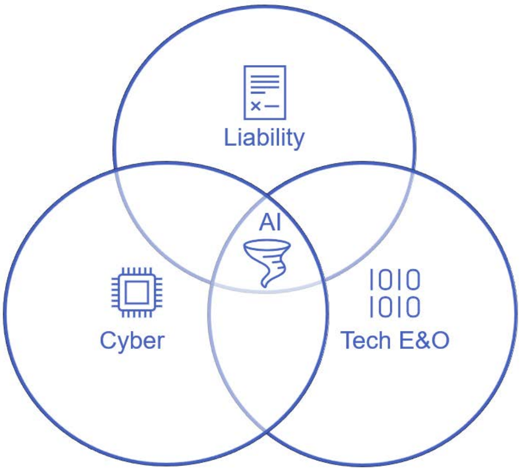

Fig. 6: Established Insurance Policies and AI Coverage

Even in this cautious environment, there are AI models in use, such as Fraud Detection and Prevention, that are running at a sufficient transactional scale to require the type of Own Damages policy envisioned by these experts. The chief risk officer of a large commercial bank explains, "We're considering how insurance could make fraud losses a constant cost for the bank, by absorbing the volatilities" of model performance. A big tech ERM executive concurs, "We look to insurance to take any revenue volatility out of AI use."

Beyond the scenario of own economic losses brought on by AI underperformance, there is also a concern about increased liability claims, in the event of IP infringement, privacy breaches, or systemic discrimination. But as a policy analyst points out, "We have a good idea of the price of human errors. But the market has yet to price the cost of AI errors."

The pace at which this develops will be determined largely by the courts, a US IP attorney predicts: "Highly visible AI claims are already in the courts. How these work out will determine how urgent the need for AI Liability insurance is." An executive at a specialty insurance wholesaler sums it up by stating, "No one wants to sign up to be the first to pay a big liability insurance premium when there are no losses. But as soon as there is an AI claims scare, that's when companies will look twice at AI Liability."

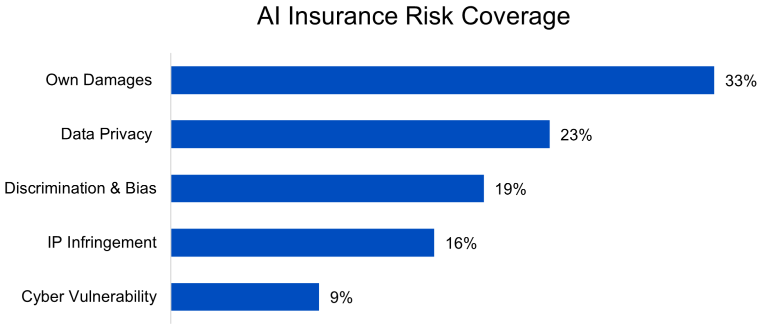

Fig. 7: Risks that Affirmative AI Covers Should Cover

For now, the indemnity and liability policies currently in place remain mostly silent on AI, neither explicitly including nor excluding AI risks. That will change according to the pace of corporate AI adoption, according to a prominent data privacy and cyber security attorney–"As more and more AI gets adopted, that's when conventional all-risk insurance will start getting poked full of holes. That's when the demand for an affirmative cover would materialize." In the meantime, he advises his clients to, "Look for the gaps between lines of insurance and try to make sure that you are adequately protected."

“You know of course though he's right about the [HAL] 9000 series having a perfect operational record. They do.”

“Unfortunately, that sounds a little like famous last words, [Dave].”

Minding the Gaps

Insurance companies tend to be seen as reactive, responding to market needs, rather than anticipating them. Nevertheless, the industry, and Munich Re in particular, have been keeping pace with the rise of the industrial, underwriting the foundational technologies with great promise but unproven track records. It started with the steam engine and the telegraph, then on to green energy and nanotechnology. AI is just the latest in a long line of technologies that Munich Re has taken the lead in insuring. In fact, we have been writing AI coverage for more than five years now.

In that time, we have uncovered a number of insurance gaps that companies would be well-advised to consider, starting with the fact that no AI model can be error-free. With a well-designed model, the frequency of errors can be rare, and yet, the severity of an individual error or an accumulation of errors may be dramatic. What is unique about AI is that such errors can occur when the model is working the way it was designed. Traditional insurers only pay for own financial damages when there is proof of negligence. Mistakes without negligence – that is novel to AI, but AI Own Damages coverages is expressly designed to cover it.

Now let’s look at discrimination: it is currently excluded from most standard policies, although specialty Employment Practices Liability Insurance do cover employee lawsuits regarding discrimination, sexual harassment, or wrongful termination. However, that will not cover most use cases that fall under the EU’s proposed High-Risk category, such as access to education and medical care. Yet, AI models can increase discrimination risk by systematizing patterns from limited training data without evidencing “discriminatory intent”.

The question remains, if the HAL 9000 were a real-world application of a machine learning model existing outside the domain of fiction, would Munich Re have insured it? To that very definite and specific question, the underwriter must answer, as always, “it depends”.

In this case the resolution of “it depends” would be achieved by a careful review of the model, the data and methodology used to train, analysis of the error rates attained with both testing and real-world data, and a review of the monitoring processes in place to detect and remediate drift.

Fig. 8: Assessing AI Performance:
What data, which type of model, and how is it monitored?

Engineering Pipeline

We would like to understand the data engineering pipeline that was constructed to design these inputs. We collect and consider any relevant scientific literature about use cases that have already used such inputs.

External Factors

We assess if and how external factors were taken into account when developing the AI model. If an AI solution is supposed to monitor the battery health of electric vehicles for example, outside temperature and driving style are external factors that need to be considered.

Data Sources

An AI model may use one or different data sources. We investigate why these particular sources were chosen, how large and how up to date they are, and why the developers consider them to be the best choice.

Performance Monitoring

How is the performance of the AI model monitored while being used by the end customer? We ask for performance data, i.e. for the KPI that is supposed to be guaranteed, to check how well the model has performed in the past.

Performance Management

We review if and how the company is prepared to handle performance drops on the part of their AI model. How frequently is the model re-trained? How is additional data generated in case re-training is necessary? What measures are taken to ensure that the new re-trained model is as accurate as the earlier one?

Quality Management

What practices are implemented to ensure the overall good quality of the software after deployment? We assess quality control measures like code reviews and merge decisions.

However, to return from the world of science fiction to the world at hand, what we do know from five years of experience, is that 90% of AI companies that undergo Munich Re's thorough technical due diligence process are offered an insurance contract.

The reality is even the HAL 9000 series can only have a perfect operational record for so long. That is true even if the best available design, engineering, and governance practices have been applied in its creation and maintenance, operational error is predictable, in fact more predictable than say natural, political, or economic catastrophes.

What is predictable is insurable. The good news for innovators and leaders in the AI space is that the mathematics of calculating and pricing risk is well understood, and Munich Re has substantial experience in insuring AI and the ongoing risk appetite. For those committed to responsible AI innovation, Munich Re is your risk transfer partner.

Contact

Ted Pine
Sr. Business Development Manager
at Munich Re

tpine@munichre.com

Learn more:

About Munich Re

Munich Re is one of the world's leading providers of reinsurance, primary insurance and insurance-related risk solutions. With more than 142 years of insurance expertise and a continuous will to innovate, the company is driving the transformation of its industry. Munich Re is developing new products for ever new types of risk, including a growing portfolio of tech insurance, including artificial intelligence solutions.

Münchener Rückversicherungs-Gesellschaft
Königinstrasse 107, 80802 München, Germany

Picture credits: Munich Re

Münchener Rückversicherungs-Gesellschaft (Munich Reinsurance Company) is a reinsurance company organised under the laws of Germany. In some countries, including in the United States, Munich Reinsurance Company holds the status of an unauthorised reinsurer. Policies are underwritten by Munich Reinsurance Company or its affiliated insurance and reinsurance subsidiaries. Certain coverages are not available in all jurisdictions.

Any description in this document is for general information purposes only and does not constitute an offer to sell or a solicitation of an offer to buy any product.