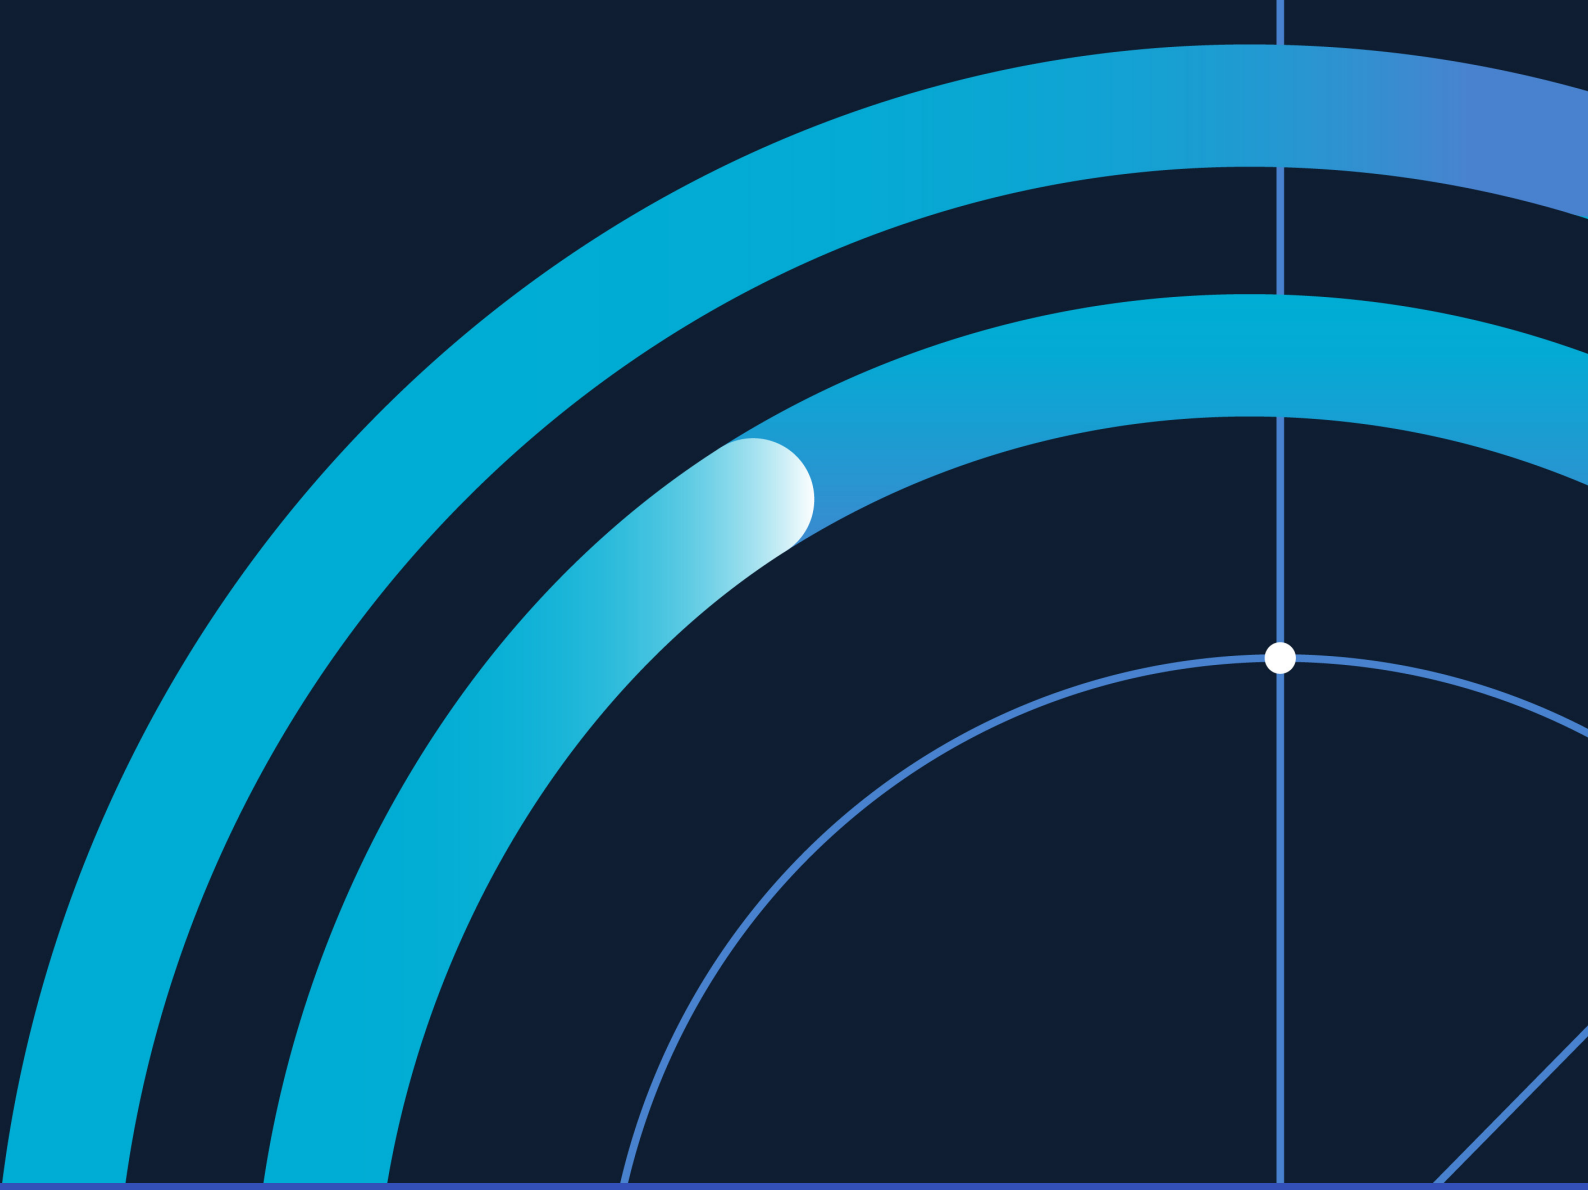

Whitepaper

Mind the Gap

A US-focused analysis of AI liability risks and the implications for insurance

Munich RE

HSB

Contents

1 Introduction 3

2 AI lawsuits: What can they tell us? 5

3 Damage caused by AI – potential financial losses, potential challenges, potential coverage 10
- 3.1 What losses can AI mistakes cause? 10
- 3.2 Who should be liable for the losses caused by AI mistakes? 12
- 3.3 What insurance policies could be implicated? 14

4 The need for a thorough analysis – spotty coverage of several relevant AI risks 16
- 4.1 How should AI exposure be assessed? 17
- 4.2 What damage seems covered by existing policies, and where do we see coverage gaps? 18
- 4.2.1 Damage most likely covered under existing policies 19
- 4.2.2 Damage with uncertain coverage under existing policies 20
- 4.2.3 Damage most likely underinsured under existing policies 20

5 Outlook for AI insurance 22

6 References 23

1 Introduction

Artificial Intelligence (AI) has captured the world's attention. AI is the broad term encompassing software systems that mimic human intelligence in decision-making. AI includes many different types of technologies including computer vision, natural language processing, and text and image generation (GenAI). While AI models have been used across industries for decades, the launch of ChatGPT, a GenAI model, in November 2022 [1] was an instant success, reaching 100 million users in two months [2]. ChatGPT caused a rapid intensification in interest in AI for many different facets of life. As of 2023, investment and research into AI is only expanding [3], with startups raising US$ 42.5bn across 2,500 equity rounds, the US seeing AI funding jump by 14% year on year in 2023 [4], and in academia the total number of AI publications doubling since 2010 [5]. While it is still early in the year, these trends appear to be continuing in 2024.

As increasingly more companies explore AI use cases in their business, the awareness of the risk of AI "going wrong" and causing damage is also increasing. From journalists cataloguing AI errors, like Aaron Drapkin in his blog *"AI Gone Wrong: An Updated List of AI Errors, Mistakes and Failures"* [6] to universities like George Washington University (GWU) collecting AI-related lawsuits in a public database [7], many people are keeping a close eye on AI models making mistakes and their repercussions to try and gauge the impact of AI errors.

Regulators have also not stayed quiet. SEC[1] Chair Gary Gensler recently commented on CNBC on algorithmic trading: *"what is the responsibility of someone using an AI model and the AI model hallucinates? [...] If a company is using Artificial Intelligence in a material way [...] and in that, that program has a tendency to hallucinate, they have to consider those risks"* [8]. Similarly, EEOC[2] Chair Charlotte Burrows pledged to *"vigorously use our collective authorities to protect individuals' rights regardless of whether legal violations occur through traditional means or advanced technologies"* [9] – with as a result that *"if the vendor is incorrect about its own assessment and the tool does result in either disparate impact or discrimination or disparate treatment discrimination, the employer could still be held liable"* [10]. These are not empty threats either – a tutoring company recently had to settle with the regulator for US$ 365,000 after using an AI model to hire tutors that was later found to be discriminating against older applicants [11]. Both statements seem to point in the same direction: not only AI providers but also intermediaries and the users of AI systems are responsible for the mistakes that their AI models make.

Despite the buzz around AI making mistakes, the question as to how existing insurance policies protect insureds if AI goes wrong is still unanswered [12]. As insurance follows liability, insurers as well as insureds are facing a multitude of questions.

The Insure AI team at Munich Re has been insuring the performance risks of AI models since 2018. As a leading provider of reinsurance, primary insurance and insurance-related risk solutions, Munich Re employs many experts in emerging technologies and risk management. In this white paper, we aim to raise some of the most important questions about the new risks posed by AI applications and provide an understanding of what to consider when answering those questions for both insurers and insureds. This paper seeks to explain:

[1] U.S. Securities and Exchange Commission (SEC): independent agency of the United States federal government. The primary purpose of the SEC is to enforce the law against market manipulation.
[2] Equal Employment Opportunity Commission (EEOC): federal agency established to administer and enforce civil laws against workplace discrimination.

- **What can lawsuits alleging AI mistakes tell us about liability?**
  In Section 2, this document analyzes a database by researchers from GWU of 137 AI lawsuits filed and categorizes them by algorithm type, the reason for the lawsuit, and the industry. The analysis shows that AI lawsuits have steadily increased and that, over time, the reasons for lawsuits diversify. It also shows that many industries have been impacted so far and that the lawsuits have hit AI models across the board.
- **What potential damage caused by AI could trigger lawsuits? Why could it be difficult to attribute liability for the damage? What insurance policies could be triggered by what risks?**
  In Section 3, this document lists the diverse nature of the potential damage that AI can cause, explores the challenges with attributing legal liability, and shows that a wide variety of insurance policies could be triggered or lead to more severe losses than anticipated if AI were to make a mistake – what we will call "silent AI risk".
- **How can insurers and brokers determine whether traditional insurance policies cover AI risks? Which risks are covered in existing insurance policies and where might the coverage be spotty?**
  In Section 4, after giving a general guide where the insurance contract could (lacking explicit terms) include – or exclude AI risks, the analysis will touch upon the areas that are currently covered, such as physical damage and bodily harm caused by AI, as well as the areas that are less covered, especially AI decisions resulting in discrimination and pure economic losses.
- **How might the AI insurance market evolve in the future?**
  Drawing from the cyber insurance industry, Section 5 will outline the different steps that led to today's cyber insurance market and – because of the parallels in the adoption of the internet and AI – contrast them with what is happening (and is predicted to happen) in the AI insurance market to try and forecast the future of AI insurance.

Due to the bulk of litigation taking place in the US, and the authors' business focus, this paper has a US-centric perspective and will not comment on other jurisdictions. With the current discussions around the EU AI Act, it will be interesting in the future to contrast the conclusions of this paper with the European rules to come.

2 AI lawsuits: What can they tell us?

AI malfunctioning or "going rogue" has been a long-time theme in science fiction, but the first settlements and lawsuits for malfunctioning AI have already been filed and paid. Due to the probabilistic nature of the underlying algorithms and despite continuous improvements, AI will continue to make mistakes. This means that, despite implementing all technical steps to create an optimal model, the risk of AI making a wrong decision cannot be eliminated. These mistakes have the potential to become increasingly devastating as AI use continues to increase in vital business operations. An online AI Incident Database has tracked over 3,000 reports of harm caused by AI [13], showing that AI-related harm is occurring, whether or not a lawsuit is filed.

In a global survey conducted by McKinsey [14] that assessed the risks companies worry about for their implemented AI, the most-cited risks were inaccuracy, cybersecurity, and intellectual-property infringement. The results of the survey are shown in Figure 1.

Figure 1 – McKinsey & Company survey on perceived risks to Generative AI

Inaccuracy, cybersecurity, and intellectual-property infringement are the most-cited risks of generative AI adoption.

Generative AI–related risks that organizations consider relevant and are working to mitigate, % of respondents[1]

[1]Asked only of respondents whose organizations have adopted AI in at least 1 function. For both risks considered relevant and risks mitigated, n = 913.
Source: McKinsey Global Survey on AI, 1,684 participants at all levels of the organization, April 11–21, 2023

Source: The state of AI in 2023: Generative AI's breakout year, McKinsey Global Institute, August 2023.

Furthermore, AI-related lawsuits are on the rise. Researchers at GWU have constructed a searchable "AI Litigation Database" [7] where AI and AI-adjacent lawsuits can be searched according to the name of the algorithm, the cause of action or even the application area.

As of February 2024, 137 lawsuits are listed, going as far back as 2004, which is in stark contrast to a quick Google search about the "first AI lawsuit", which points towards a 2022 settlement by the EEOC. This database is maintained by volunteers and provides a picture of the state of AI litigation. The number of lawsuits may therefore not be exhaustive, but can provide a representation of the landscape of AI litigation.

The cumulative numbers in Figure 2 show a gradual increase in AI lawsuits over the past 20 years – besides indicating an increase in the reasons[3] for filing the lawsuits. As an increasing number of people and businesses start using, creating and selling AI applications, it is not surprising to see an increase in the number of lawsuits. As AI models change both in their application and in their capabilities, the reasons for filing lawsuits, the alleged harms, become more versatile. Lawsuits for due process violations, privacy violations and misrepresentations of AI use were some of the first lawsuits filed. When AI models were being increasingly used in making housing decisions where they delivered biased outputs, there was a rise in lawsuits filed for discrimination in 2018 and 2019. The rise of privacy lawsuits around 2019 can be attributed to a series of class-action lawsuits against Clearview for their scraping of images of faces from across the web. The rise of bodily injury lawsuits in 2020 stems mainly from Tesla's autopilot making mistakes. Finally, the jumps in IP infringement and defamation lawsuits are due to the rise in popularity of GenAI with the introduction of ChatGPT. As the use cases for AI continue to proliferate through the economy, we expect the number of lawsuits to increase.

[3] The classification of these lawsuits was performed by the authors. **Due process Admissibility / Civil and Constitutional Right Infringements due to AI:** Cases under this heading included areas where AI was being used in a way unacceptable to the injured party's most basic human rights. An example was the denial of parole based on a risk assessment conducted by AI (Rodriguez v. Massachusetts Parole Board), but also areas where AI was being used in courts. **Privacy violations:** Cases under this heading included areas where AI or the training of AI breached consumer privacy. Examples included several cases of facial recognition technology (Burke v. Clearview AI, Inc., FTC v. Rite Aid Corp, et al.); **IP Infringement:** Cases under this heading included areas where plaintiffs accused companies of using their intellectual property to train algorithms. Example cases include Getty Images (US), Inc. v. Stability AI, Inc. where a company sued an image creator for using its images in training their models; **Other:** Cases under this heading include areas not otherwise classified; **Discrimination:** Cases under this heading represent lawsuits alleging discrimination by AI systems. Example cases include lawsuits where an algorithm was discriminating based on race, gender, etc. (National Fair Housing Alliance v. Facebook); **Economic and Financial losses e.g. business interruption:** This category includes inaccuracies, business interruptions, hallucinations, and underperformance resulting in economic and financial losses. An example case is C.S. et al. v. Saiki, where an algorithm made decisions that reduced benefits for Oregon residents; **Bodily injury / Property Damage / Death:** This category includes lawsuits alleging bodily injury or death. Examples include lawsuits against companies electing to use an AI (incorrectly) which later led to bodily injury (Murphy v. Essilorluxottica); **Defamation:** Lawsuits in this category allege that generative AI is defaming them by hallucinating (Walters v. OpenAI, L.L.C.)

Figure 2 – Cumulative number of AI-related lawsuits mostly in the US according to [10]. X-axis starts in 2012 for clarity.

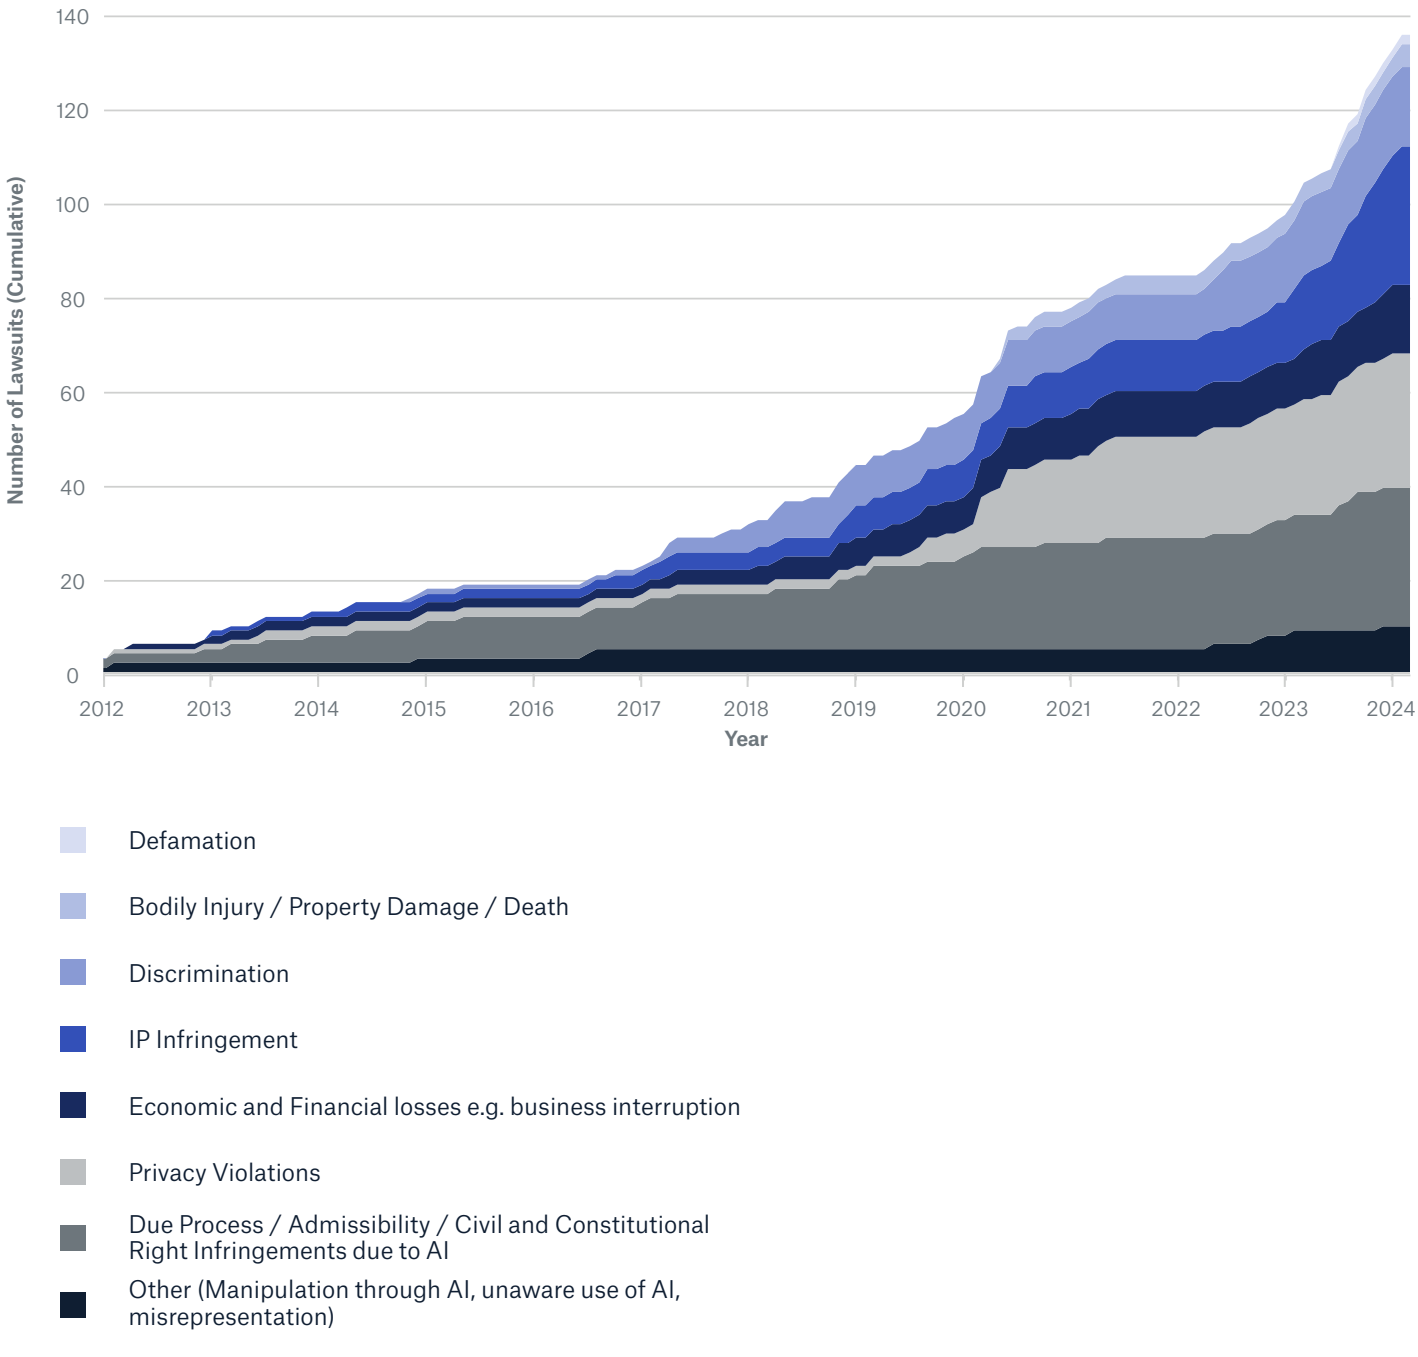

The data shows that, over the past decade, what started as due process lawsuits has given way to a rise in claims alleging discrimination and IP infringement.

Figure 3 shows the lawsuits listed by algorithm, where that information was available. Facial recognition and self-driving vehicles are high-risk use cases for AI, and significant care should be taken when such technologies are deployed. The EU AI Act has named these two use cases as "unacceptable risk" (meaning that in future these use cases cannot be deployed in the EU) and "high risk"[4] (which are use cases subject to assessment by third parties prior to being placed on the market), respectively [15]. The risk is reflected in the number of lawsuits against Clearview (facial recognition) and Tesla's Autopilot (self-driving help). However, GenAI is also notably represented in the lawsuits, with many major AI model providers being impacted, such as OpenAI, Microsoft, Stability AI and Meta. These lawsuits are colored grey in Figure 3. Finally, algorithms used in law enforcement, like MIDAS and COMPAS, are also being sued.

Figure 3 – Lawsuits by algorithm. Grey bars indicate GenAI models

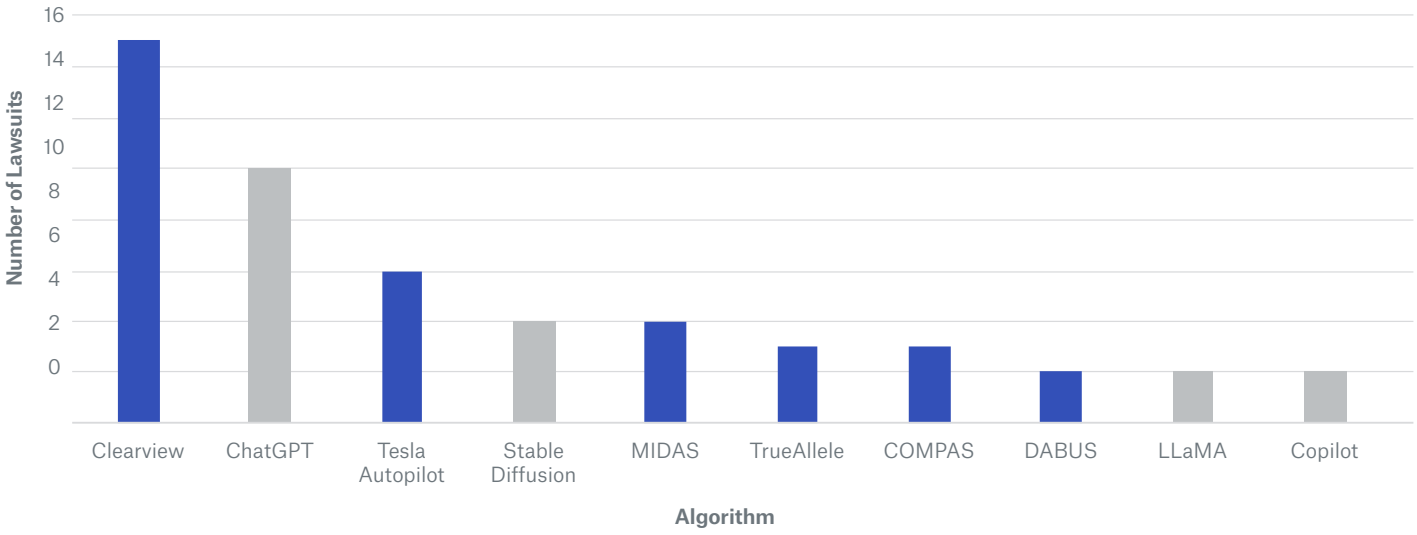

[4] Facial recognition is a high-risk application when used for remote real-time biometric identification.

These algorithms are being applied in many different domains of life. AI has a broad reach across the economy. The data in Figure 4 shows that the lawsuits span a wide range of industries.

Figure 4[5] – Number of AI lawsuits by industry[6]

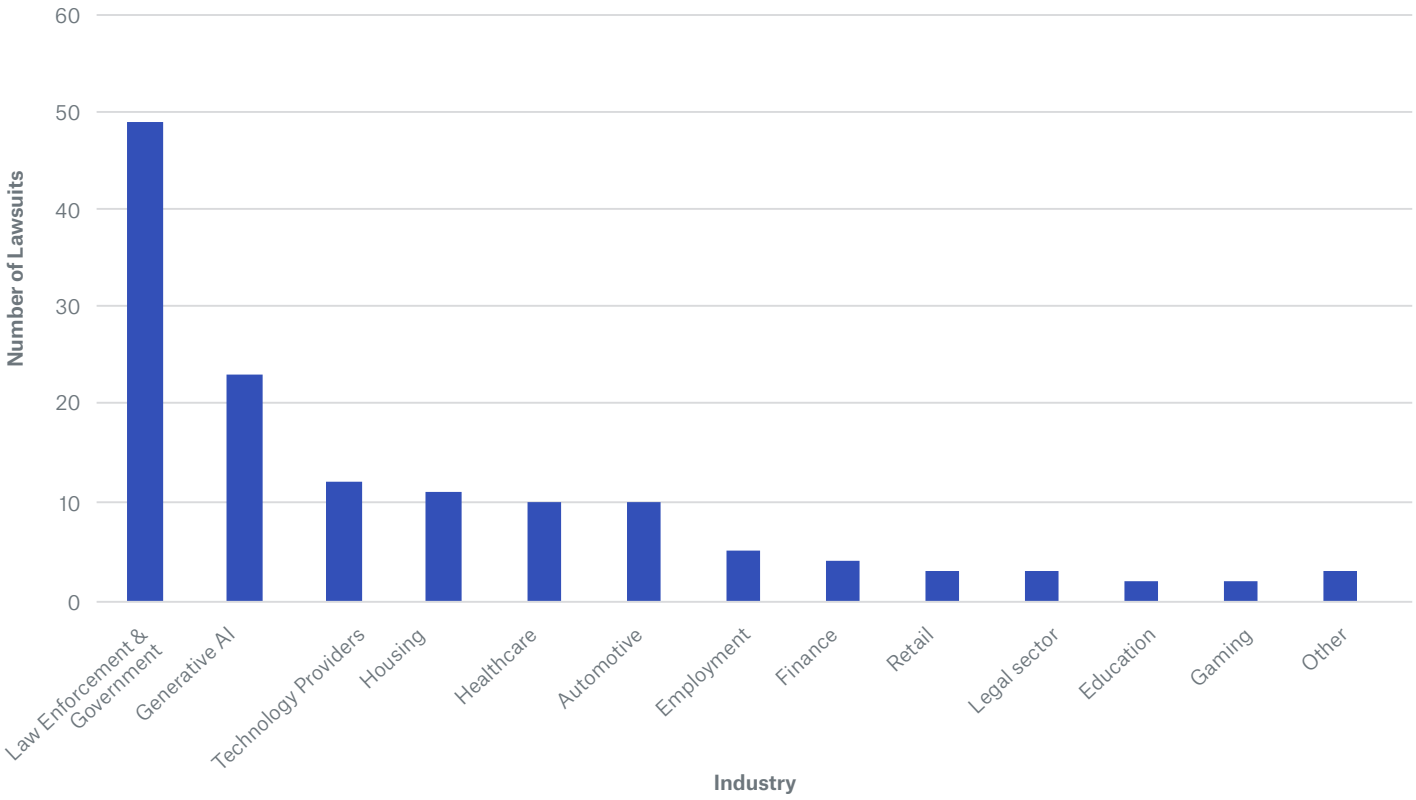

Analyzing the AI Litigation Database illuminates a few key points:

- Despite the first AI-related lawsuit being filed in 2004, an uptick in AI lawsuits in the past six years is visible
- New avenues to sue for AI-related losses (read: reasons for AI lawsuits) have been steadily increasing over the years
- It is not one "faulty model" that is being sued – all models can make mistakes or cause different types of harm; risks cannot be mitigated by "avoiding the bad model"
- As AI enjoys widespread use in many industries, most industries are represented in the lawsuit database – the potential for harm caused by AI is not confined to one industry

We have shown that the number of AI-related and AI-adjacent lawsuits is steadily rising. As most of these lawsuits still have to be concluded, we are currently unable to comment on damages quantum, length of trial, jurisdiction clustering and other factors that influence the impact of lawsuits for the parties. However, taking the points above, we can infer that AI lawsuits should start to become a real concern for AI providers, AI adopters and AI users.

The next Section will therefore examine the potential damage that malfunctioning AI could cause, point out the challenges regarding establishing liability for the damage and, finally, list some of the insurance coverages that could potentially cover AI mistakes.

[5] The 'Other' category includes insurance, art, customer services and activism.
[6] Law enforcement and government includes the judiciary, border control, police, and general government.

Analyzing AI related lawsuits shows that AI mistakes can cause a wide variety of losses:

- Property damage, bodily injury or even death
- Fines and penalties imposed by a regulator or authoritative body
- Privacy violations
- Data leaks
- IP infringement
- Pure financial losses
- Defamatory comments, images, or references generated by an AI model
- Discrimination

3 Damage caused by AI – potential financial losses, potential challenges, potential coverage

Like for many other business risks, there are mitigation strategies that companies can adopt to minimize their exposure to costly or vexatious lawsuits, caused by AI malfunctioning. Some of these strategies have been outlined in Munich Re's previous white paper [16]. Technical mitigations reduce the risks of AI malfunctioning, but cannot fully eliminate them.

Below, we will outline the potential damage that AI systems could cause, why attributing that damage to involved parties can be tricky, and which insurance policies could be triggered.

3.1 What losses can AI mistakes cause?

- **Property damage, bodily injury or even death** could be the result of AI malfunctioning within physical products. Examples are self-driving cars injuring pedestrians, as happened when a driverless taxi in San Francisco ran over a woman in October 2023 [17], cleaning robots causing property damage, or an automated plant malfunctioning and injuring employees, like in South Korea in 2023 where an industrial robot crushed a man after misidentifying him as a box [18].
- **Fines and penalties imposed by a regulator or authoritative body** where the use of an AI model violates due process rights or other constitutional rights, or where AI models discriminate against protected groups and a regulatory authority imposes fines to compensate the victims or deter the wrongdoers. Many US agencies, including the FTC[7], EEOC, CFPB[8], DOJ[9] and SEC, have issued statements and strategies describing their no-tolerance stance on AI-related wrongdoings [9]. At the heart of the statement is the clear message that "the AI did it" will not be considered a valid defense for model users and developers to disclaim liability for injury caused to third parties.
- **Privacy violations** might be alleged if an AI model unlawfully collects confidential, personal or identifying information, or if the scraping (a technique where a computer program extracts data from human-readable output coming from another program) used in the training of the models unlawfully captures personally identifiable information. Exemplary lawsuits include the alleged unlawful surveillance activities by Clearview AI in California (Renderos v. Clearview AI, Inc.) [19], the unlawful collection of customers' voiceprints without their consent in a fast-food restaurant drive-thru (Carpenter v. McDonald's Corporation) [20], and the training of GenAI models on private information from internet users (P.M. et al. v. OpenAI LP) [21].
- **Data leaks** often lead to huge damage, as they can result in the dissemination of trade secrets and lead to financial losses, public embarrassment or loss of shareholder or customer confidence. Leaks can occur as a result of using AI. To date, no lawsuits have been filed for data leaks. This is not to say that they haven't happened yet. Most well-known is probably the incident at Samsung in 2023, where Samsung employees inadvertently leaked sensitive company

[7] Federal Trade Commission (FTC): an independent agency of the United States government whose principal mission is the enforcement of civil antitrust law and the promotion of consumer protection.
[8] Consumer Financial Protection Bureau (CFPB): an independent agency of the United States government responsible for consumer protection in the financial sector.
[9] United States Department of Justice (DOJ): federal executive department of the United States government tasked with the enforcement of federal law and administration of justice in the United States.

data, including the source code of software responsible for measuring semiconductor equipment [22]. Other risks include the possibility of leaking training data, as experienced by researchers successful in tricking an Nvidia large language model into releasing personally identifiable information of over 71,000 employees [23] via an adversarial model prompt.

– **IP infringement** related to AI is an area that is dominated by GenAI lawsuits. Lawsuits and other disputes pertaining to alleged IP infringements, insofar as they are caused or triggered by the use of GenAI, could come about from several different directions:

  - Firstly, the user of the GenAI model could be sued for using generated images that are substantially similar to existing copyrighted images. GenAI platforms, aware of this risk and the impact it could have on their business, are already offering protection against this risk for a subgroup of their clients [24, 25, 26].
  - Secondly, the owners of copyrighted images or protected text could sue GenAI platforms for unauthorized use of their images or text as training material, as alleged by, amongst others, Getty Images (Getty Images (US) Inc. v. Stability AI [27]), the New York Times (The New York Times Company v. Microsoft) [28] and the Authors Guild (Authors Guild v. OpenAI, Inc.) [29], which is the bulk of copyright cases against AI models at the moment.
  - Finally, there have been lawsuits to determine whether AI-generated images could themselves have copyright protection, with different outcomes depending on the jurisdiction [30, 31].

– **Pure financial losses** are financial losses suffered that are not directly attributable to physical injury to persons or property. They could happen, for example, where an AI model inaccurately creates commercial projections that are relied upon. One example is Zillow's recently terminated "iBuying" home-flipping program, where its AI models inaccurately forecasted home prices that Zillow relied upon, causing losses for Zillow, as well as the subsequent shutdown of the business in 2021, which resulted in 2,000 employees losing their employment [32]. Similarly, pure financial losses can occur where a trading algorithm autonomously makes trades based on accidentally wrongly published figures, like in the case where Lyft's blunder in a published profitability metric caused a stock price surge – and subsequently an abrupt stock price correction – of over 60% after hours, caused by the trading algorithms [33]. Business interruptions are also examples of pure financial losses, e.g. in manufacturing plants using AI. As the reliability of AI improves, we expect to see more and more examples of business interruption caused by malfunctioning AI. Economic and financial losses can come from a variety of different issues caused by AI:

  - **Underperformance and inaccuracies** can lead to damage and lawsuits in a variety of different areas. If error-prone AI models are used in regulated areas like housing (C.S. et al. v. Saiki: 04/01/2017) and healthcare (Barrows et al. v. Humana, Inc. 12/12/2023), regulators may fine the responsible companies for allowing "bad models" to make decisions in sensitive areas. Those fines result in financial losses for the company. Furthermore, misrepresentations and overpromising of models' performance can lead to lawsuits brought by investors in the company, as regulated by the SEC in the case of BlueCrest Capital Management Ltd and others (Tewson v. DoNotPay, Inkie Lee v. Tesla, Inc.).

- **Hallucinations**[10] are a subgroup of inaccuracies, usually found in GenAI models. Chatbots are especially prone to creating hallucinations, which can be a problem when they provide information that is relied upon. Air Canada was ordered by a Canadian tribunal to reimburse a customer for incorrect information its chatbot provided that resulted in the customer being charged a higher fare[11] [34].

– **Defamatory comments, images, or references generated by an AI model:** As GenAI creates new content, although accurate in most cases, it can generate content that mistakenly links a person or company to a situation by misinterpreting the connections between those two references. That is what a radio host alleged in his lawsuit when ChatGPT produced a false text stating that he embezzled money from a gun-rights organization (Walters v. OpenAI LLC) [35]. This lawsuit is currently the only one of its kind, but AI generating defamatory statements is not, as shown by a professor in Stanford who was falsely labelled a "terrorist" [36] by ChatGPT.

– **Discrimination,** more specifically disparate impact discrimination, might be alleged if a biased model makes discriminatory decisions against a protected group. The EEOC settled its first AI discriminatory hiring suit in September 2023, ordering a tutoring provider using AI in hiring decisions to pay USD 365,000 [37].

Analyzing existing law shows the challenges of assigning liability for AI making mistakes, due to the lack of transparency in models because:

– the autonomy of the models and the lack of foreseeability of AI mistakes
– the "black box nature" of algorithms making it impossible to understand what "caused the mistake to happen"
– the openness and contributory nature of the AI building process

3.2 Who should be liable for the losses caused by AI mistakes?

"A robot may not injure a human being …" begins the directive laid out in the sci-fi visionary's work "I, Robot" [38]. As the potential losses discussed above show, Isaac Asimov's first law of robotics is receding in the rearview mirror [39]. Since AI does indeed make mistakes, the question is: who is responsible for them, and who should pay for the costs that arise? Despite decades of academic discussion on the topic, a clear attribution of responsibility has yet to emerge.

There is general agreement that "where the hand of human involvement in machine decision-making is so evident, there is no need to re-examine liability rules" [40]. This means that where the error or omission of a person caused an AI to malfunction, that person bears the responsibility (read: is liable) for the resulting damage caused.

Legal uncertainty starts where the human involvement is less clear-cut. When AI is autonomous – meaning that it could operate in ways that are unforeseeable by the original programmers – difficult questions about liability arise [41]. Some scholars, like Madelaine Clare Elish, suggest that the moral blame does not go to the algorithm itself but to the "nearest human" responsible [42]. However, finding the "nearest human" is not straightforward and causes significant challenges to the current legal system, as discussed below.

Common-law tort doctrines generally require a breach of a duty of care owed, and a harm being caused due to that breach. Existing tort law might be insufficient to assign liability due to the human-centric responsibilities they assess.

[10] „AI hallucination" is a phenomenon wherein a large language model – often a GenAI chatbot or a computer vision tool – perceives patterns or objects that are nonexistent or imperceptible to human observers, creating outputs that are nonsensical or altogether inaccurate.

[11] The chatbot gave the passenger the advice that they could apply for a discount in retrospect (which was wrong) while displaying a link to the official policy. The courts in Canada held that the customer would be awarded the discount despite the link to the correct information, because Air Canada was not able to demonstrate why the customer should not have been able to trust what the chatbot was saying, thus not needing to click the link.

As uncertainty in the law affects effective risk management, the concepts that will be most challenged by AI will be outlined briefly. Especially questions around the standard of care required, how AI particularities affect causation, who will have the burden of proof to establish the tort and, finally, the question of strict liability, will be the more challenging ones for courts and regulators to answer.

- **Standard of care in tort:**
  More than half a century ago, when discussing the efforts that a driver had to make to avoid an accident, a Louisiana court ruled that "A human being, no matter how efficient, is not a mechanical robot and does not possess the ability of a radar machine to discover danger before it becomes manifest. Some allowances, however slight, must be made for human frailties and for reaction [...]." [43]. Continuing this thought logically would point towards a more stringent duty of care owed to third parties by AI algorithms than humans, as David Vladeck correctly points out [40] – a question that will surely be discussed by the courts at some point in the future.
- **Causation considerations:**
  Questions around which effect actually caused the damage – an important question to answer when understanding who is liable – as well as whether it could have been avoided because it was foreseeable, are obscured by the autonomous decisions AI makes, the opaqueness of AI decisions (also often referred to as the "black box") and the openness of AI systems:
  - **Autonomy and foreseeability requirements:**
    AI is taking over more and more complicated tasks without active human control or supervision[12]. As AI is able to analyze more possible scenarios of certain outcomes than humans, sometimes even find optimal ones, to outside human observers the system's actions can be unexpected. If the AI system acts without the action being foreseeable, it would be unfair to create a rule to, for example, hold the designers liable for the harms the systems caused – as is the norm for manufacturers. As such, the actions of a learning AI system could be viewed as a superseding cause, similar to an employee acting outside the scope of employment [41], breaking the chain of causation necessary to establish liability.
  - **Opaqueness and black-box system:**
    All algorithms are complex and require specialized knowledge to be understood. Sometimes, particular features or weights can be so deeply buried in the algorithm that they are hard to identify and understand – even for their creators. Today's LLMs sometimes contain billions of parameters, which makes understanding how a parameter correlates to an output impossible. The opaque nature of "black box" algorithms and deep learning applications results in no one knowing exactly what caused the AI to act a certain way. This becomes especially true when the AI system is further modified through updates or self-learning [44]. If the expert, the creator, is unsure of how and why a decision was made by the algorithm, how should this be determined in a courtroom?
  - **Openness of AI models:**
    Complications might also arise out of the fact that algorithms are often made up of different contributions that are not necessarily classified as coordinated input (think open-source contributions). Furthermore, as it can be incredibly expensive to develop and build a fully-functional LLM from scratch, companies will license a publicly available model, a foundation model as the basis of further fine-tuned models that fit a specific business need. The outcomes of the model will then depend on a variety of conditions from different systems and data, making it very difficult to point to the parties responsible for specific errors the model makes.

[12] E.g. https://waymo.com/waymo-one-san-francisco/.

– **Burden of proof:**
Because of the way that causation is obscured by autonomy, opaqueness and the openness of AI models, proving liability is a very difficult task. In order to win a case in court, it will therefore be vital to establish which party has the "burden of proof" – meaning the duty to show that the elements necessary to establish liability were present – and the point at which that burden of proof will be shifted to the other side to prove that liability cannot indeed be established.

As seen, a traditional application of tort law is difficult when considering the new attributes and touchpoints of AI models. While this is just one example of the many possible causes of action an injured party may follow when seeking damages for losses caused by an AI model, it demonstrates that assigning liability to any one party in the ultimate value chain of AI is not straightforward or certain. Given the human-centric approach to negligence in tort, legal scholar Zeynep Tufekci suggested assigning human-level agency to AI models – an approach that would allow us to claim "the algorithm did it" – against the wishes of the regulators mentioned above [45]. Other authors have suggested that, rather than attempting to fit into the confines of the law, a specific duty of care doctrine should be introduced for algorithms [46] or a "reasonable algorithm" standard created that would be applied to self-learning systems [47].

In the EU, as suggested by the EU Proposal for a New Product Liability Directive [48], AI might soon fall under the no-fault product liability regime. By introducing this regime, the European Union is expecting legal certainty, enhanced consumer trust in AI and assistance in consumers' liability claims (through burden of proof shifts and information requirements) for damage caused by AI-enabled products and services [49]. Similar treatment in the US would also allow AI to fall under product law and its corresponding stricter liability regime – with different outcomes depending on whether the "defect" was classed as a manufacturing or design defect [50].

3.3 What insurance policies could be implicated?

While the above considerations are exciting for legal scholars, the resulting legal uncertainty is less attractive for risk managers and small businesses. Until express guidelines from lawmakers are introduced, courts must continue to interpret the law as it appears before them. As a consequence of this landscape, we expect many exemplary (and sometimes frivolous) lawsuits, as parties try to establish the boundaries of conduct and responsibility when it comes to providing or using an AI model.

The involvement of multiple parties in AI-related lawsuits – be it the many parties affected, or the experts needed to establish fault and causality – could escalate legal costs substantially. The more complicated the discovery process is, the more time, reports and expert judgements and witnesses will be needed – another factor driving up litigation costs. Lawyers aiming to sue successfully might "throw claims at the wall to see what sticks" as per the saying, which might prove fatal for SMEs who have to front preparation, defense and settlement costs against these lawsuits. Given the unpredictability of the assignment of liability and the quantum of damages that might be awarded against whoever is found liable, businesses might benefit from seeking to offset these costs through a third party. This practice can improve the predictability of cash outflows and balance possible volatility arising from high unexpected costs.

One such method of offsetting liability and its associated costs is by obtaining an insurance policy to indemnify a company for legal fees and damages it is required to pay[13].

As of now, the insurance market has not started to expressly exclude AI-caused losses in policies or jurisdictions. Conversely, affirmative coverage for AI risks is still sparse (but existent – see our previous white paper [16]). Traditional insurance policies that may respond to a loss caused by an AI model decision or mistake are listed below:

- If the training or use of AI results in privacy breaches or digital threats, **cyber liability policies** could be triggered.
- If the use of AI results in algorithmic bias and system failures, **errors and omissions policies** could be triggered.
- If the deployment (or even lack of deployment) of AI in a company leads to mismanagement in company processes, **directors' and officers' policies** could be triggered.
- If services or products using AI result in injury to others, **commercial general liability** policies could be triggered.
- If an AI-controlled industrial robot injures an employee, **workers' compensation policies** could be triggered.
- If the use of AI results in physical damage, resulting business income and extra expenses could also be covered if the "all-risks" **property policy** is triggered.
- If the use of GenAI leads to pictures that are too similar to copyrighted materials, or if the training data used falls under IP protection, **intellectual property policies** could be triggered.
- If third parties are harmed due to defective AI or AI malfunctions, **product liability policies** could be triggered.
- If AI makes biased choices in employment and is discriminatory in determining applications, promotions or wrongful determinations, **employment practices liability policies** could be triggered.
- If people sue against AI-generated content, e.g. for making defamatory statements, **media liability policies** could be triggered.

Furthermore, industry-specific coverages like medical malpractice insurance could be triggered if an AI tool used in the medical field leads to a wrong diagnosis.

All of the coverages mentioned above are scenarios where insurance coverage may respond. However, none of the coverages mentioned above were originally created to cover the particular losses that AI mistakes could lead to, leading to so-called "silent AI"[14] as these policies still might respond to losses caused by AI, representing an unexpected risk to insurers' portfolios.

Insurers may seek to re-price these policies, amend the wordings or other conditions, or expressly exclude AI-related events from future contracts once companies attempt to claim.

[13] It should be noted, that AI can (and frequently does) reduce the frequency and severity of losses for in existing policies. AI-based sensors and cameras can be e.g. used in sensitive equipment, alarming the owners when repairs are needed - reducing their breakdown and replacement

[14] the expression "silent AI" is a nod towards the term „silent cyber", a frequently discussed topic in the cyber insurance industry

4 The need for a thorough analysis – spotty coverage of several relevant AI risks

As a provider of targeted AI risk policies, we, the Insure AI team, are often asked the question "isn't this risk already covered by one of my existing policies?". So far, there have been very few coverage analyses publicized on the market that detail the coverage of AI perils by traditional policies. A few international broking houses [51, 52] and a primary insurance company [44] analyzed the coverages, but reached contradicting conclusions.

As wordings and conditions differ across markets, general statements made about the extent of coverage should not be taken as absolute fact. In any event, comments about the extent of coverage available to businesses by their insurers is always subject to the decision of the respective insurance companies' claims departments (or the courts presiding over insurance disputes). However, due to the significance of the topic, the difficulty of providing overarching and universally correct answers in this ever-changing field should not deter us from attempting to shed some light on the issue.

As stated in section 3.2, AI-related lawsuits are projected to include a multitude of parties and experts, likely becoming very expensive. This risk, coupled with the potential of silent AI, as many traditional insurance policies being impacted by AI-related insurance claims, leads to a need for insurers to be aware of which of the traditional policies in their portfolio might be exposed to risks and damages resulting from malfunctioning AI. As many policies could be affected, unexpected limit stacking could occur, further exposing insurance companies to unexpected claims.

For a related reason, it is also important for insureds to consider how extensively the policies they currently have protect them against AI-related risks they might be exposed to. The potential for hefty regulatory fees, lengthy legal proceedings and the implication of (even remote) parties in lawsuits are risks that insureds are increasingly becoming aware of. Generally speaking, large companies will most likely have many different insurance policies acting like an umbrella. They should focus on whether their limits are adequate – and potentially search for very specific coverage gaps. Smaller and mid-sized companies, who often do not have the widespread insurance coverage of bigger companies, could be far more exposed to AI risks than they expect[15].

The purpose of this analysis is,

1) to increase the awareness of potential AI-related losses being covered by traditional policies that were not specifically designed (and priced) to cover AI risk "silent AI" [53].
2) to point to spotty areas in the current coverage where risks that are exacerbated or introduced by AI are not risks that the insured is usually (adequately) covered against in their existing policies.

[15] For example, a local restaurant using GenAI to create marketing flyers will probably be exposed to the risk of IP infringement lawsuits, as they don't have Media Liability insurance policies and the limits in their other policies covering IP infringement can be very low.

After giving broad guidance on how to assess the coverage of AI in traditional policies – including the challenges we see – we will provide our view on some of the perils that we see as currently covered, and which we see as underinsured.

Our opinion is just one of many relevant and sometimes conflicting views on this topic, and we would recommend businesses and insurers to discuss the scope of coverage for each contract they are party to.

4.1 How should AI exposure be assessed?

When trying to assess the coverage of AI risks in traditional policies, there are three "places" in the insurance agreement that are particularly relevant: 1) the trigger 2) the covered event 3) the exclusions.

– **Trigger:**
Coverage triggers are defined in an insurance policy as an event that will lead to a payout. For example, for general liability coverages, this trigger is a legal claim brought against the insured; for cyber coverages, this trigger can be the occurrence of a specific (cyber) event.
With regards to AI, the occurrence trigger could present a point of ambiguity, as it could reference to the point at which the model was first introduced, or the time when the damage was first realized.

– **Insuring agreement:**
Insurance coverage is the amount of risk or liability that is covered for an individual or entity by way of insurance services. It is the part where the insurer agrees to take certain actions such as paying losses for covered perils, providing certain services, or agreeing to defend the insured in a lawsuit.

It is important to make a distinction between two types of policies: Open peril and named peril policies. Named peril policies cover damage arising out of events specifically listed in the policy. Open peril insurance policies generally cover all events, once the insurance is triggered, that are not specifically excluded in the policies. As AI is not currently an explicitly excluded peril, all open peril insurances –general liability policies, for example – cover AI damages if they fall under the insuring agreement.

The most obvious policies that could be impacted by malfunctioning AI are general liability policies and cyber/tech E&O policies, both of which could be triggered by a claim being made.

- The insuring agreement of general liability policies covers damages arising from a legal liability that resulted in physical harm to a third party (bodily injury, property damage, etc.). General liability coverages often exclude pure financial losses and business interruption costs.
- Pure economic losses could be covered by cyber and tech E&O policies instead. There is a requirement of "failure of a technology product or service" in order for the insuring agreement to cover the damage. With systematic AI failures, this requirement could be fulfilled, as it could be said that the product or service is not fit for the agreed purpose. However, when it comes to one-off hallucinations or short-term underperformances, these can result in serious economic losses without potentially fulfilling the requirement of "failure" as required in the case of cyber coverage.

Product liability policies could also be impacted and could offer protection against AI failures. However, for AI to fall within the insuring agreement, AI would need to be considered a product (as opposed to a service), an important question that has not yet been decided upon.

When analyzing whether one is sufficiently covered for AI mistakes caused by AI used or deployed, one must understand which damage could occur when using AI, how likely and how severe that damage might be, and whether the specific damage is covered under the Insured's existing policies.

– **Exclusions:**
Exclusions take coverage away from the insuring agreement. After finding the policy triggered and the insurance agreement covering the damage, it is therefore important to consider whether any exclusions apply.

Damages arising from AI could be excluded from coverage, such as when:

- They arise out of lawsuits alleging discrimination, which may broadly include compensatory damages, costs required to re-train or fix a model, lost revenue or profit or possible shareholder action. This is a very relevant risk, as there are multiple lawsuits, published cases and scholarly articles that allege biased models leading to the discrimination of protected groups.
- They are punitive damages awarded by courts, which could become more relevant as the awareness of AI perils increases. This is a particular concern for businesses in the US, where courts award higher punitive damages than in the EU. However, there are already areas in which malfunctioning AI could lead to punitive damages, especially in the context of GDPR rulings.
- They are regulatory settlements, like the ones that have already been reached with the EEOC and the FTC.
- They are caused by a reliance on technical instruments, e.g. in the field of medical malpractice, since only errors by FDA-approved technical instruments are covered, and AI may or may not be FDA-approved.
- Other exclusions that could impact coverage are "intentional act" exclusions, war and terrorism exclusions, losses excluded (like pure economic losses, business interruption losses, losses stemming from product recall, etc.) and many more "specific exclusions".

4.2 What damage seems covered by existing policies, and where do we see coverage gaps?

Understanding potential coverage gaps and silent AI exposure is a task that involves an understanding of insurance, liability workings and the particularities of AI. As many of the lawsuits are still in the infancy stages, another important part of determining AI risk coverage is creativity. Only by imagining what mistakes AI could make in different scenarios, what harm those mistakes could cause, and what policies might respond, can the coverage of AI risks be mapped. The imagined scenarios will be signaled by icons. This exercise, by its nature, will not provide us with a comprehensive picture, but will instead highlight some of the more obvious risks and corresponding coverage considerations.

4.2.1 Damage most likely covered under existing policies

If **AI causes physical harm to a person or property,** most underwriters will consider this a loss that would be covered by a property insurance or general liability insurance product. For a single AI-driven event, the limits of multiple policies could be stacked.

Firstly, let us examine the scenario in which a company installs AI-based water leakage sensors. A pipe bursts and the AI fails to detect it, which leads to damage of the insured's property (building and contents). The AI in this case was intended to reduce the risk of water leakage in the building. In the first instance, the liability for a failure of the AI in this scenario would fall onto the property facility manager and be covered under the company's property policy. Subsequently, the property facility manager might seek compensation from AI providers, therefore the AI provider might also be held liable for the failure of their system, in that it failed to perform as designed.

Secondly, a car manufacturer uses AI to align wheels (or to perform another step of the manufacturing process) and, due to systemic underperformance of the AI, a batch of cars is produced with faulty axes. This is only realized after they are sold and there is a recall. The risk of a recall might be excluded by many product liability policies, but will usually be covered by a general liability policy. Again, the AI provider might also be held liable for the failure of the product.

Thirdly, let us imagine a company employing AI for chemical manufacturing process control. The AI is intended to optimize a parameter of a process but underperforms badly, exceeding safety limits. The AI causes a runaway chain reaction in the reaction process. The reaction causes a fire, damaging property and injuring several employees and bystanders. Such losses are likely covered by property policies (fire damage to the property), general liability policies (bodily injury to the bystanders) and workers' compensation policies (injury to the employees).

Usually, in first instance, the AI users would be liable for the damage caused by AI. However, they could seek compensation from the AI providers. There is a standard professional liability exclusion in many general liability policies that excludes coverage for third-party claims that allege bodily injury or property damaging arising out of the selling, licensing or furnishing of computer software [54]. For AI providers, this might constitute a potential coverage gap for AI products that cause bodily injury or property damage claims.

If AI mistakes can be classed as a **"failure of a technology product to perform as intended"**, third-party claims and potentially also first-party losses could be covered under cyber insurance and tech E&O insurance, including resulting regulatory liability and media liability. For an AI mistake to be classed as a failure to perform as intended, the mistakes made must be so severe and/or frequent as to render the use of AI pointless. An AI mistake that could be classed as a failure to perform could be, for example, self-driving cars not being able to determine their lanes when it rains. At the other end of the spectrum, damage resulting from a GenAI-powered chatbot that hallucinates every now and then, but otherwise generates reliable results, would probably not be considered a failure of technology and would therefore not be covered under cyber and tech E&O policies. Cyber coverages typically exclude coverage for contractual liabilities, while tech E&O policies typically exclude coverage for liability arising out of bodily injury or property damage.

If there are significant AI model or implementation errors, misrepresentations about the AI to shareholders, or potentially even the lack of AI integration (very unlikely in our current legal environment), **cause lawsuits against a company's directors or officers,** costs associated with such lawsuits are likely to fall within the scope of D&O insurance.

4.2.2 Damage with uncertain coverage under existing policies

When examining the area of AI-driven **intellectual-property infringements** in its many different forms, the coverage becomes less clear-cut. Consider the case of a small plumbing business using a GenAI model to create "two Italian plumbers" as its business mascot, and the GenAI creates a spin-off Super Mario picture. The plumber uses the picture on his car and is sued by Nintendo. Several insurance policies could be implicated:

- Intellectual property policies cover litigation expenses associated with pursuing those who infringe a company's intellectual property rights, as well as litigation expenses associated with a company defending itself against a claim of infringement of another's intellectual property right.
- For tech E&O and cyber policies, while patent infringements are excluded, copyright infringements could be covered. However, economic losses (e.g. lost revenue) arising out of the infringement of copyright are not covered under these policies.
- General liability policies cover copyright infringements, as well as defamation, under their advertising injuries clauses if the injury was committed inadvertently.

Another example of IP infringement lawsuits related to AI are the current lawsuits against GenAI companies for IP infringement in training data. Losses resulting from these lawsuits would probably be covered by media liability insurance, as well as cyber insurance. It is not surprising that we are hearing the first voices in the market contemplating excluding AI in Cyber policies. Therefore, more tech companies are buying media liability policies, policies that were traditionally bought by publishing companies, to cover their extensive lawsuits for unauthorized IP infringement during training.

Private users of generative AI, however, usually have no extensive cyber or liability policies and are currently mostly unprotected against IP lawsuits. This is why several GenAI companies are including indemnities as part of their premium offerings that cover users in the case of a lawsuit for copyright infringement caused by an AI to fill this market need.

When looking at **fines and penalties from regulators,** the majority of US states preclude coverage as a matter of public policy [55], if the statutory fines are penal rather than remedial. Due to the difficulty of assessing whether a "fine or penalty" is punitive or compensatory in nature, case law across the jurisdictions in the USA tends to be fact-specific, due to the nature of the specific penalty or fine, the language of the policy and the state's public policy on punitive damages. As such, it is uncertain whether these potential (and not unlikely) losses are sufficiently covered in existing policies.

4.2.3 Damage most likely underinsured under existing policies

Finally, we have identified a few risks that we consider to be very relevant to the adoption of AI, but which in our opinion are currently underinsured.

Firstly, we take a step back from third party losses and look at the risk of first-party **pure economic losses** (aka losses not stemming from a physical harm) caused by AI, while very relevant given the many business applications of AI, are likely underinsured:

– A potential example could be: A bank uses AI to conduct property valuations during loan onboarding, instead of using (human) property valuation service providers. If the AI underperforms and produces a valuation for a property that is too high, besides the borrower subsequently becoming insolvent, resulting in foreclosure, a financial loss might arise for the bank. After liquidation of the property, the bank suffers a loss amounting to the difference between the actual value and the wrong valuation (because the AI assessed it too high). As the AI values many properties over time, there is an additional risk of accumulation due to a systemic error for many properties. Historically, when the bank contracted property valuation service providers and they were held to fail to deliver their services or do so negligently, the bank was protected from losses through the service provider's professional indemnity insurance. Now, this potentially unwanted and systemic risk is on the bank's balance sheet. As a pure economic loss, damage resulting from AI underperformance in this case is not insured under the bank's traditional coverages (e.g. professional indemnity, cyber, tech E&O, liability policies). As AI is used in many cases that could lead to economic losses if relied upon (trading algorithms, revenue predictions, etc.), we are of the opinion that this could leave an insurance gap in traditional insurance policies.

– Another scenario could be business interruption costs caused by AI: A car manufacturer implements AI to adjust axle. The AI underperforms and, as a result, cars with unadjusted axis are put on the market. Production is halted to prevent accidents, and retraining the updated AI will take weeks. Revenue from unsold cars is lost. The operation is interrupted.[16] Pure economic losses are not covered under traditional property insurance. Without a "breach" or an actual "system failure", there is no coverage under a cyber policy. Potentially (but highly case-dependent), there could be coverage under the model provider's tech E&O policy, but only if a legal liability can be established that was breached by the AI provider. We believe this is a new type of risk stemming from AI underperformance that is not sufficiently covered under available policies.

Finally, the Insure AI team at Munich Re considers that bias in AI model decisions which can lead to discriminatory outcomes against protected groups is not sufficiently covered under existing policies. **Losses caused by discrimination** are excluded under general liability, cyber and tech E&O policies. EPLI policies cover the damages caused by discrimination in the employment space. Some EPLI policies also cover discrimination against "third parties", which could provide far-reaching protection against alleged discrimination. However, as EPLI policies are not part of the standard insurance portfolio of most AI users or AI providers, we consider this risk to be currently underinsured.

Generally, our recommendation is for AI users, AI providers and AI integrators to examine and calculate the extent to which their business operations are exposed to risks of AI underperformance or loss-causing events, and to contact their insurer or broker to discuss, assess and potentially further protect themselves against these risks.

[16] For clarification: we do not explore BI following software failure and shutdown or BI following a security or privacy breach, as those are cyber policy risks arising from using IT systems. The AI-induced risk is either the so-called irreducible (random) error or a systematic change in error due to model drift, such that the system must be stopped and BI occurs until retraining is completed. This scenario is similar to unscheduled maintenance following a machinery malfunction, even though the AI is not malfunctioning in this case: The scenario is the realization of the natural error fluctuation and randomness inherent in every AI model.

5 Outlook for AI insurance

The widespread adoption of AI across nearly all corporate functions and all industries, as well as the consequential impact on insurance, has frequently been compared to the adoption of the internet in the 2000s. Much like the rise of cyber risks in the late 1990s prompted insurers to explore new territories, we believe that the surge in AI use will soon become a focal point for insurers and brokers.

When the first cyber policies written, they focused on specific loss scenarios, were tailor-made and had a strong technology focus. When we compare this with the AI insurance market, this seems to be the stage of AI insurance to date: In order to navigate the current complexities of AI adoption, specific risks are addressed through tailor-made policies, such as Munich Re's standalone aiSure™ products.

If we continue drawing this parallel into the future, the emergence of a standardized AI insurance market seems a likely consequence – comparable to the cyber insurance market. Our predictions:

As businesses grapple with the transformative potential of AI, insurers will start developing policies to manage AI-related liabilities. When losses from cyber incidents started spiking, risk managers, brokers, and insurers started thinking about cyber risks more systematically and strategically. Considering the recent uptick in IP-related lawsuits, lawsuits against healthcare companies using AI, and the increased interest shown by regulatory agencies in not tolerating discrimination by AI models, an increase in AI related losses seems to be on the horizon.

As insurers recognized the accumulation potential of cyber incidents in which cyber-attacks and cloud systems downtime could affect many insureds at the same time, as well as the "silent cyber" exposure of property and general liability insurance policies, widespread cyber exclusions in traditional policies were adopted. The adoption of exclusions for cyber incidents accelerated the purchase of cyber policies. Similar accumulation potential could be attributed to the few existing foundation models, which are often a basis for more specific AI services, and single models being relied upon by many different companies. Additionally, when looking at all the potential insurance policies that could be triggered by "AI going wrong", this whitepaper argues that "silent AI" exposure is found within many traditional policies.[17]

Regulatory landscapes, exemplified by e.g. the GDPR for cyber risks, play a pivotal role. Similarly, AI regulation will likely spur businesses to follow evolving guidelines and adopt responsible AI initiatives, paralleling the regulatory journey in cyber insurance. Once the regulatory cyber landscapes were more clearly defined, markets started navigating compliance phases and developing standardized processes aligning with regulatory norms. This shift simplified underwriting and marked a transition towards an informed, standardized market practice, echoing the journey of other established insurance sectors.

[17] It is of note that if AI risks are to be covered by existing traditional insurance policies, then AI risks should be appropriately reflected in the pricing of the existing policies.

The Insure AI team at Munich Re deems the possibility of the emergence of an AI insurance market, marked by standardized practices and structured pricing, as likely - akin to the evolution witnessed in cyber insurance. Similar to the evolution of the cyber insurance market, we believe that in the future various policies like those mentioned throughout this document, will start to exclude coverage for AI risks, as they are difficult to price with traditional means. Our recommendation to risk managers with AI exposure is to determine the whole context of their AI activities as it relates to creating new AI risks. As the AI regulatory environment continues to change and technology continues to rapidly progress, this will stay an interesting area to watch.

Do you have any questions? Would you like to discuss in more detail? Please reach out.

Contact

Dr. Peter Bärnreuther
Senior Underwriter for Artificial Intelligence Risks
PBaernreuther@munichre.com

Ted Pine
Business Development Manager Sr
tpine@munichre.com

More information:
https://www.munichre.com/insure-ai

6 References

[1] OpenAI, „Introducing ChatGPT," 30 November 2022. [Online]. Available: https://openai.com/blog/chatgpt.
[2] K. Hu, „ChatGPT sets record for fastest-growing user base - analyst note," Reuters, 2 February 2023. [Online]. Available: https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/.
[3] N. Maslej, L. Fattorini, E. Brynjolfsson, J. Etchemendy, K. Ligett, T. Lyons, J. Manyika, H. Ngo, J. C. Niebles, V. Parli, Y. Shoham, R. Wald, J. Clark and R. Perrault, „The AI Index 2023 Annual Report," Stanford University, Stanford, CA, 2023.
[4] CB Insights, „State of AI 2023 Report," CB Information Services, New York, NY, 2024.
[5] Stanford Universtity, „The AI Index Report: Measuring trends in Artificial Intelligence," Stanford Institute for Human-Centered Artificial intelligence, Stanford, CA, 2023.
[6] A. Drapkin, „AI Gone Wrong: An Updated List of AI Errors, Mistakes and Failures," tech.co, 19 February 2024. [Online].
[7] J. Fattah, X. (. Qu, A. Schiele, J. Chen, S. Huppert, Z. Kim, M. Brown, R. Brauneis, J. Wang and S. Zhao, „AI Litigation Database - Ethical Tech Initiative," Institute for Trustworthy AI in Law and Society at George Washington University, [Online]. Available: https://blogs.gwu.edu/law-eti/ai-litigation-database/.
[8] G. Gensler, Interviewee, CNBC's full interview with SEC Chair Gary Gensler. [Interview]. 14 February 2024.
[9] U.S. Equal Employment Opportunity Commission, „Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems," [Online]. Available: https://www.eeoc.gov/joint-statement-enforcement-efforts-against-discrimination-and-bias-automated-systems#:~:text=We%20also%20pledge%20to%20vigorously,traditional%20means%20or%20advanced%20technologies..
[10] U.S. Equal Employment Opportunity Commission, „Select Issues: Assessing ADverse Impact in Software, Algorithms, and Artificial Intelligence Used in Employment Selection Procedures Under Title VII of the Civil Rights Act of 1964," 18 May 2023. [Online].
[11] U.S. Equal Employment Opportunity Commission, „EEOC Sues iTutorGroup for Age Discrimination," U.S. Equal Employment Opportunity Commission, 05 05 2022. [Online]. Available: https://www.eeoc.gov/newsroom/eeoc-sues-itutorgroup-age-discrimination. [Accessed 04 03 2024].
[12] W. Wang and M. Infantino, „Algorithmic Torts: A Prospective Comparative Overview," Transnational Law & Contemporary Problems, Vol. 29, No. 1, Forthcoming, 2018.
[13] Responsible AI Collaborative, „Welcome to the Artificial Intelligence Incident Database," [Online]. Available: https://incidentdatabase.ai/.
[14] QuantumBlack AI by McKinsey, „The state of AI in 2023: Generative AI's breakout year," McKinsey, 2023.
[15] EU Parliament, „EU AI Act: first regulation on artificial intelligence," 19 December 2023. [Online]. Available: https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence#:~:text=Unacceptable%20risk%20AI%20systems%20are,encourage%20dangerous%20behaviour%20in%20children.
[16] Munich Re, „Insuring Generative AI: Risks and Mitigation Strategies," 2024.
[17] A. Rose, V. Miracle and J. Kopp, „A woman was found trapped under a driverless car. It wasn't the first car to hit her | CNN Business," CNN, 3 October 2023. [Online]. Available: https://www.cnn.com/2023/10/03/tech/driverless-car-pedestrian-injury/index.html.

[18] Associated Press, „Industrial robot crushes worker to death as he checks whether it was working properly - CBS News," CBS News, 9 November 2023. [Online]. Available: https://www.theguardian.com/technology/2023/nov/08/south-korean-man-killed-by-industrial-robot-in-distribution-centre.
[19] J. Bhuiyah, „Clearview AI uses your online photos to instantly ID you. That's a problem, lawsuit says," 9 March 2021. [Online]. Available: https://www.latimes.com/business/technology/story/2021-03-09/clearview-ai-lawsuit-privacy-violations.
[20] Segal McCambridge, „Technology & Cyber Risk Client Alert: Evolution of BIPA Voiceprint Claims in Illinois Persists - Segal McCambridge Singer & Mahoney," 24 February 2022. [Online]. Available: https://www.segalmccambridge.com/blog/technology-cyber-risk-client-alert-evolution-of-bipa-voiceprint-claims-in-illinois-persists/.
[21] I. Poritz, „OpenAI Hit With Class Action Over ‚Unprecedented' Web Scraping," Bloomberg Law, 28 June 2023. [Online]. Available: https://news.bloomberglaw.com/ip-law/openai-hit-with-class-action-over-unprecedented-web-scraping.
[22] S. Ray, „Samsung Bans ChatGPT Among Employees After Sensitive Code Leak," Forbes, 2 May 2023. [Online]. Available: https://www.forbes.com/sites/siladityaray/2023/05/02/samsung-bans-chatgpt-and-other-chatbots-for-employees-after-sensitive-code-leak/?sh=4a7a803c6078.
[23] C. Criddle and M. Srivastava, „Nvidia's AI software tricked into leaking data | Ars Technica," Financial Times, 9 June 2023. [Online]. Available: https://arstechnica.com/gadgets/2023/06/nvidias-ai-software-tricked-into-leaking-data/.
[24] MIcrosoft, „Introducing the Microsoft Copilot Copyright Commitment," 7 September 2023. [Online]. Available: https://www.microsoft.com/en-us/licensing/news/microsoft-copilot-copyright-commitment.
[25] P. Venables and N. Suggs, „Protecting customers with generative AI indemnification," google, 12 October 2023. [Online]. Available: https://cloud.google.com/blog/products/ai-machine-learning/protecting-customers-with-generative-ai-indemnification.
[26] OpenAI, „New models and developer products announced at DevDay," 6 November 2023. [Online]. Available: https://openai.com/blog/new-models-and-developer-products-announced-at-devday.
[27] BakerHostetler, „Getty Images v. Stability AI | BakerHostetler," [Online]. Available: https://www.bakerlaw.com/getty-images-v-stability-ai/.
[28] M. M. Grynbaum and R. Mac, „The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work," New York Times, 27 December 2023. [Online]. Available: https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html.
[29] The Authors Guild, „Authors Guild Supports Nonfiction Writers in Lawsuit Against OpenAI - The Authors Guild," 21 December 2023. [Online]. Available: https://authorsguild.org/news/ag-supports-nonfiction-writers-in-lawsuit-against-openai/.
[30] A. Guadamuz, „Chinese court declares that AI-generated image has copyright – TechnoLlama," TechnoLlama, 9 December 2023. [Online]. Available: https://www.technollama.co.uk/chinese-court-declares-that-ai-generated-image-has-copyright.
[31] Authors Alliance, „Copyright Protection in AI-Generated Works Update: Decision in Thaler v. Perlmutter | Authors Alliance," 24 August 2023. [Online]. Available: https://www.authorsalliance.org/2023/08/24/copyright-protection-in-ai-generated-works-update-decision-in-thaler-v-perlmutter/.
[32] M. Clark, „Zillow is moving out of the home-selling business - The Verge," The Verge, 2 November 2021. [Online]. Available: https://www.theverge.com/2021/11/2/22760080/zillow-offers-home-selling-ibuyer-wind-down-excess-inventory-losses-financial-results.

[33] J. Broughel, „What Lyft's Big Typo Teaches Investors About New Technologies," Forbes, 16 February 2024. [Online]. Available: https://www.forbes.com/sites/jamesbroughel/2024/02/16/what-lyfts-big-typo-teaches-investors-about-new-technologies/?sh=11ac779677d1.
[34] K. Melnick, „Canada Airline to pay customer after chatbot gave false information - The Washington Post," The Washington Post, 18 February 2024. [Online]. Available: https://www.washingtonpost.com/travel/2024/02/18/air-canada-airline-chatbot-ruling/.
[35] I. Poritz, „OpenAI Fails to Escape First Defamation Suit From Radio Host," Bloomberg Law, 16 January 2024. [Online]. Available: https://news.bloomberglaw.com/ip-law/openai-fails-to-escape-first-defamation-suit-from-radio-host.
[36] T. Hsu, „What can you do when AI lies about you? | The Seattle Times," The Seattle Times, 7 August 2023. [Online]. Available: https://www.seattletimes.com/business/what-can-you-do-when-ai-lies-about-you/.
[37] U.S. Equal Employment Opportunity Commission, „iTutorGroup to Pay $365,000 to Settle EEOC Discriminatory Hiring Suit," U.S. Equal Employment Opportunity Commission, Washington DC, 2023.
[38] I. Asimov, I, Robot, Garden City, NY: Doubleday, 1950.
[39] C. Kirby, „Iowa law magazine," Iowa College of Law, 02 11 2022. [Online]. Available: https://law.uiowa.edu/iowa-law-magazine/news/2022/11/when-algorithms-harm-us. [Accessed 27 02 2024].
[40] D. C. Vladeck, „Machines without Principals: Liability Rules and Artificial Intelligence," Washington Law Review, vol. 89, no. Symposium: Artificial Intelligence and the Law, pp. 117 - 150, 2014.
[41] M. U. Scherer, „Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies and Strategies," Harvard Journal of Law and Technology, vol. 29, no. 2, pp. 353 - 400, 2016.
[42] M. C. Elish, „Morale Crumple Zones: Cautionary Tales in Human-Robot Interaction," Engaging Science, Technology, and Society, p. 29, 2019.
[43] Arnold v Reuther 92 So. 2d 593 (La. Ct. App. 1957), 1957.
[44] Zurich Insurance Company Ltd., „Artificial intelligence gives rise to 'algorithmic liability' | Zurich Insurance," 29 July 2021. [Online]. Available: https://www.zurich.com/en/knowledge/topics/digital-data-and-cyber/artificial-intelligence-gives-rise-to-algorithmic-liability.
[45] Z. Tufekci, „Algorithmic Harms Beyond Facebook and Google: Emergent Challenges of Computational Agency," Colorado Technology Law Journal, vol. 13, pp. 203 - 218, 2015.
[46] A. Glaubnitz, „How should liability be attributed for harms caused by biases in Artificial Intelligence," Yale Jackson Institute for Global Affairs, New Haven, CT, 2021.
[47] K. A. Chagal-Feferkorn, „How Can I Tell if My Algorithm Was Reasonable?," Michigan Technology Law Review, vol. 27, no. 2, pp. 213-261, 2021.
[48] S. D. Luca, „European Parliament," 12 2023. [Online]. Available: https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2023)739341. [Accessed 28 02 2024].
[49] A. Gold, „Next up for AI in the EU: Liability," Avios Pro, 10 10 2023. [Online]. Available: https://www.axios.com/pro/tech-policy/2023/10/10/next-up-for-ai-in-the-eu-liability. [Accessed 12 03 2024].
[50] K. A. Chagal-Feferkorn, „Am I an Algorithm or a Product? When Products Liability should apply to Algorithmic Decision Makers," Stanford Law & Policy Review, vol. 30, no. 61, pp. 61-114, 2019.
[51] B. Dyson, „S&P Global Market Intelligence," S&P Global, 12 12 2023. [Online]. Available: https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/insurers-brace-for-claims-from-generative-ai-surge-79584195. [Accessed 05 03 2024].
[52] Boyum, Eric; Ratcliff, Rachel; Gonzales, Jesus, „Use AI and Insurance Insights to Make Better Decisions," AON plc., https://www.linkedin.com/posts/adam-furmansky-b508a5a_use-ai-and-insurance-insights-to-make-better-activity-7088249675392114688-i7D3/, 2023.

[53] S. Aloni, „Surge in AI Ushers A „New Silent Cyber" Risk | Cyber Insurance Academy," Cyber Insurance Academy, 24 September 2023. [Online]. Available: https://www.cyberinsuranceacademy.com/knowledge-hub/guide/surge-in-ai-ushers-a-new-silent-cyber-risk/.
[54] J. Meagher and A. Connelly, „Navigating the New Frontier: Insurance for Artificial Intelligence Risks," the National Law Review, vol. XIV, no. 65, March 2024.
[55] Marsh McLennon, „Insurability of fines and penalties | Marsh," 21 June 2022. [Online]. Available: https://www.marsh.com/tr/en/services/financial-professional-liability/insights/insurability-of-fines-and-penalties.html#map.
[56] European Parliament Think Tank, „New Product Liability Directive | Think Tank | European Parliament," 12 April 2023. [Online]. Available: https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2023)739341.

© 2024
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.