Charting the Course of AI Risk

Practical Considerations for Financial Services Leaders

A publication of the FS-ISAC Artificial Intelligence Risk Working Group Future Use Team

The Future State of Generative AI

Contents

Contributors

- Co-Chairs: Chris Budd, Deutsche Bank, and Pam Simpson, TD Bank
- Benjamin Dynkin
- Donavon Swinney, WECU
- Dr. Carrie E. Gates, FS-ISAC
- Michael Silverman, FS-ISAC

Executive Summary

Artificial Intelligence (AI) simultaneously inspires optimism and panic. Its latest iteration — generative AI, which creates wholly new text, images, audio, and more — brings the financial sector to a new inflection point. While AI clearly offers opportunities that the sector can embrace, organizations have just begun to set their course into the future of AI.

To help chart what lies ahead for AI's use and consequences, FS-ISAC's Artificial Intelligence Risk Working Group assembled the Future Use team of financial sector experts in AI, risk, technology, and cybersecurity. In this publication, the Future Use group puts AI, especially generative AI (GenAI), in a broad financial services framework to help institutions:

- Predict the problems and identify the uncertainties involved in the phases of AI solution implementations
- Define the challenges and understand the nuances of AI in financial services
- Predict the short-, medium-, and long-term impacts of GenAI tools in realistic scenarios

We conclude this work with several questions to help your firm understand, then manage, your own AI risks.

There is not enough data to make long-term evidenced-based predictions on whether AI implementation will exponentially provide value or depreciate efficiency and effectiveness over time. Patents and publicly available academic research aren't sufficient to isolate trends. Few have wholly assessed the viability of moving human workloads onto machines and converting AI's promise into material business solutions. Some firms purport to be aggressively pushing toward adoption; others are running into significant concerns. Clearly, the sector needs a practical, actionable approach to navigating this new landscape.

Our paper seeks to offer an "all-hazards" approach – guidance on a broad spectrum of concerns, with practical guardrails and risk concepts in a useful framework to help the sector protect the advantages of AI. In so doing, financial firms can maximize AI's business value and minimize its potential disadvantages at the same time.

Defining AI for Financial Services Use Cases

Projecting risks to the financial services sector requires investigating how AI might transform business operations. Financial firms are likelier to make more accurate predictions by exploring their own needs for bringing AI into production in incremental steps. The first of these steps should be a concrete definition of AI.

The term "AI" itself stretches from the technically deterministic to science fiction. AI can mean any number of things, from machine learning to deep learning, from natural language processing to robotics to artificial general intelligence.

It's useful to focus specifically on generative AI, which creates new text, images, videos, and audio (or a combination of them, often referred to as multi-modal models) that didn't exist before.[i] The most pertinent GenAI models for financial services are probably large language models (LLMs), as these hold the most promise for transformation in the financial sector due to their capabilities for knowledge management.

Popular LLMs

Examples of popular LLMs include Microsoft's Co-Pilot, Open AI's Chat GPT, Google's Gemini, and the open-source framework Llama.

With this in mind, each organization should define AI for itself against its own objectives for the technology. We suggest that your definition:

- Incorporates industry standards (NIST, ISO, etc.)
- Is not so broad that it's not useful
- Is not so specific that it dates as conditions change

Understanding AI Capabilities vs. AI Expectations

Though they present potential risk vectors, LLMs demonstrate remarkable capabilities in a wide array of domains and allow organizations of any size to automate many knowledge-based tasks.[ii] The more advanced language models consistently outperform the average human's abilities in reasoning, reading comprehension, mathematical problem-solving, and computer programming.[iii, iv]

Many LLMs approach or surpass 90th-percentile human results on standardized tests such as the LSAT, the multistate bar exam, and the GRE.[v]

These technologies face some constraints. For example, in their current form, AI systems can't integrate knowledge independently as humans can and they lack autonomous agency – e.g., without human prompts, AI engines can only wait. Nonetheless, GenAI use cases in financial services are already well publicized, such as:

- Summarizing publicly available prospectus materials
- Customer support
- Cost-effective and fast coding for business applications

However, in these nascent stages, as with any new technology that disrupts business operations (like the commercialized Internet did in the 1990s), AI's inherent risks cannot be wholly understood.

Projecting Future Use: Defining the Problems

Predictions become less accurate as the time horizon grows more distant. This principle applies to AI usage in financial services firms – the further ahead you look at the future of AI, the less certain you can be about the impact and value of your AI implementations.

Every institution faces this paradigm, but it can shake leaders' confidence in their AI implementations. To develop realistic predictions about the outcomes of AI solutions, we recommend predicting your future use in a framework of the following five phases:

- Phase 1: Certainty
  The firm assesses the known AI value versus the cost of implementing the AI solution. The firm begins with this basic business approach, which requires understanding the business, use cases, and people, processes, and technology resources available. That background enables it to set budgets and select appropriate solutions.

- Phase 2. Some uncertainty
  Use cases are prioritized, including automation of repetitive tasks. Those tasks are compared to the value of the new opportunities identified above. This can only be accomplished if the business understands the value of AI (Phase 1) and is sufficiently confident it can adequately apply that knowledge to prioritize use cases.

- Phase 3. Increased uncertainty
  Ad hoc implementation issues begin to arise. The institution understood the value of AI and what it wants to implement but doesn't have an adequate concept of the risks and threats. The firm is not alone: The sector as a whole is still identifying the

technology and cybersecurity risks as they apply to confidentiality, integrity, and availability relating to AI.

- **Phase 4. Increased uncertainty**
  Use cases do not operate as expected. As business opportunities evolve and even grow, technology risks may compound with useability. The AI model may not be able to perform as intended. If the institution faces a shortage of AI or machine learning engineers, it may experience significant risks if the tool cannot meet objectives.

- **Phase 5. Most uncertainty**
  The sustainability and repeatability of processes is not understood. Due to the fast-changing nature of AI, governance and processes may be upheaved as AI transforms along with groundbreaking new technologies like quantum computing.

Loss of Certainty in Gen AI Implementations over Time

Like any other business-facing risk, certainty fades over time. Monitoring this process throughout an AI implementation's lifecycle helps teams make more informed decisions, manage uncertainties more effectively, and ultimately deliver a successful product.

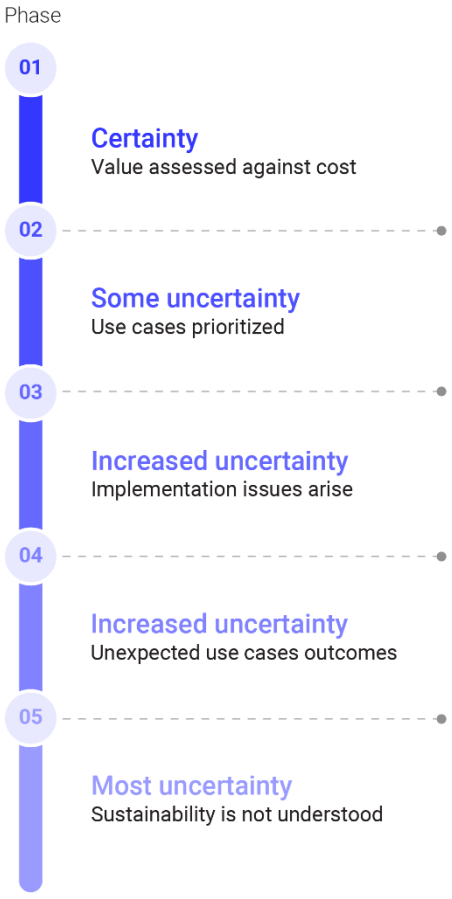

Uncertainties and Potential Challenges of GenAI

Your current state, as well as your future one, likely includes the following uncertainties and challenges — as other organizations have encountered – presented by GenAI.

Shortage of staff skilled in AI: Internally-deployed LLMs introduce less risk but a shortage of skilled AI technicians opens firms to operational risks, including misconfiguration, lack of maintenance, inadequate support, etc.

As financial services firms adopt more AI-driven tools, cybersecurity practitioners will need to understand the risks and threats of those tools. As with any new technology, training and professional development will be imperative to defend against emerging risks.

As use cases mount, the ability to use GenAI productively in proprietary deployment will require:

- Building capabilities and skill sets
- Conducting AI training so employees know how to ingest data into a training dataset safely
- Identifying and deploying the duties necessary to move AI models from a developmental environment to a production environment

Creative destruction: Innovation dictates that new, more efficient, more effective, more powerful forms of technology cannot functionally coexist with previous generations of technological systems. As organizations implement AI, they must understand what can and cannot be automated, and they must communicate how the workforce will change as a result of task automation.

This is key to reducing tensions between knowledge workers and the business line deploying the GenAI tools. Ultimately, this is also the difference between creative destruction and business disruptions.

Mistakes tolerance: There may be instances of overlap between GenAI and human workers, specifically where the task requires some automation but it can't fully replace a human. Machine error (hallucinations) and human mistakes are a possible result of that overlap. Therefore, organizations will need to define a threshold or tolerance for mistakes. In establishing such parameters, the organization would do well to consider certain factors, including the tolerance threshold for different kinds of mistakes and the impact on job roles that incorporate AI.

Over-confidence in GenAI outputs: Data analysis relies on information integrity, but AI outputs can contain hallucinations and errors. Nonetheless, financial services employees – even those entrusted to apply critical thinking – may unknowingly rely on information lacking necessary integrity. Blind acceptance of AI outputs, especially those from the black boxes of public LLMs,[vi] is unwise and AI may streamline processes to a degree that review becomes minimal, introducing the risk of acting on incorrect information.[vii] GenAI usage for financial services needs to meet the highest standards of veracity, and employees must know how to test outputs for accuracy.[viii] A relationship between AI usage, oversight, and review will be essential to healthy productivity and growth, especially in mission-critical scenarios.

- What percentage of incorrect GenAI and human responses is acceptable for the given task?
  - Do employees know which errors could incur business disruptions, opportunity costs, and other revenue losses?
  - Do employees know which errors could cause regulatory non-compliance?
  - How will errors be detected in operations that prioritize speed over accuracy?
- What are the defined roles between human workers and AI?
  - Will job roles be re-defined to include AI response fact-checking rather than more creative work?
  - How will worker apathy be addressed?
  - How will skills training optimize productivity and reduce risks?

Over-reliance on GenAI outputs: In crisis or incident situations, where cyber teams have little time to collect or act on information, GenAI can be leveraged to increase efficiency.[ix] But systems can cease to operate, and lack of alternative structures can hobble responses.

Secure by Design

Secure by Design principles centralize customer security as a core business requirement, not as a technical feature or add-on.

Ongoing due diligence alongside the development and implementation of AI will need to be built into employee skillsets that adhere to a "Secure by Design" framework. As organizations seek to achieve their targeted state of AI implementation, human elements must remain able to achieve business objectives in the event of AI failure. As new roles for AI emerge within the business – such as a Chief Artificial Intelligence Officer – firms may benefit from a mix of employees whose prior business experience makes them knowledgeable about older or existing processes and new hires who have AI or GenAI-specific experience to innovate more established processes.

Rules and regulations: Regional, national, and international legislation and frameworks exist (and more will come) to guide the implementation of AI.

The regulatory landscape that defines how GenAI can and should be used is still developing. Traditionally, regulations are put in place to stimulate innovation, however, current proposed AI regulations indicate governments may seek to limit AI innovation and adoption.

Noteworthy Legislation

- The Council of Europe Framework Convention on Artificial Intelligence is the first international legally binding treaty on AI.[x]
- The EU Artificial Intelligence Act classifies AI risk and denotes accountabilities for it. [xi]
- The Principles to Promote Fairness, Ethics, Accountability and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics was drafted by the Monetary Authority of Singapore (MAS) in partnership with banks and the technology sector.[xii]
- The Social Principles of Human-Centric AI and AI Guidelines for Business are Japan's risk-based, soft-law approach to AI regulation.[xiii]

Nonetheless, and despite the slower adoption of AI by EU companies compared to US firms, the EU AI Act may produce 'the Brussels effect,' impacting regional, national, and international AI regulatory initiatives.[xiv]

The Brussels Effect

The European Union's unique ability to shape the world's business environment due to its impact on global markets despite its limited legal jurisdiction. Global uptake of GDPR privacy standards is one example of the Brussels Effect.

Upcoming EU and US regulations and legal trends may impact considerations beyond the here and now. For example:

- The EU AI Act's 24-month implementation timeline encourages self-regulation through the early adoption of its provisions, best practices, voluntary compliance commitments, and public pledges (e.g. via the AI Pact). The most direct influence on financial institution deployments worldwide would be the development of the EU AI Act standards complementing the existing ISO/IEC 23894 and 42001.[xv]
- The Organization for Economic Cooperation and Development (OECD) principles were adopted in 2019 and updated in May 2024.[xvi] They rely on a two-tiered approach for implementing trustworthy AI, guided by value-based principles to ensure AI remains a positive force rooted in inclusivity, sustainability, transparency, and accountability. The framework also provides recommendations for policymakers, which represents a pursuit towards a common good use of AI. While these principles may not be applicable to financial institutions in total, understanding the trustworthy AI principles within a larger context can help create a framework aligned to overarching societal goals.

Thaler V. Perlmutter

In August 2023, the US District Court for the District of Columbia decided in *Thaler v. Perlmutter* that AI-generated work can't be copyrighted as copyright law only protects human-generated artwork.

Copyright issues: Challenges to the legitimacy of AI copyright protections will continue with the increased adoption of AI technologies (i.e. Thaler V. Perlmutter[xvii]), especially where organizations seek to copyright work generated from AI models.[xviii] The concept of "intellectual property (IP) laundering" may emerge for institutions where IP is unintentionally used within organizational outputs and then embedded in proprietary processes or documentation. Firms should be prepared for the legal and moral implications of IP laundering, with controls in place to map data lineage and data supply chain through GenAI models to outputs from the GenAI models.

Organizations should consider the jurisdictional laws and ongoing legal cases when conceiving ideas or documentation that LLMs or GenAI will help to create.

With that in mind, the following questions should be considered:

- As AI models routinely produce results that incorporate parts of, or are similar to, copyrighted material, would organizations be subject to copyright lawsuits?
- When is content generated by GenAI no longer considered previously copyrighted content?[xxii]
- Unless it creates a proprietary AI, how can an organization state with certainty its AI models were not trained with copyrighted data?
- Does using models that contain copyrighted materials constitute a violation of copyright — even if the copyrighted data is not accessed by employees or distributed publicly in any capacity?

Cost-prohibitive energy requirements: AI technologies require significant and ever-increasing amounts of power – the infamous nuclear power plant, Three Mile Island, has reopened to power a Microsoft data center.[xxiii] Considering the exponential growth and deployment of AI technologies, current energy sources may prove insufficient for the foreseeable future unless significant investments are made in AI infrastructure.

↑ 50%-100%
GPU power consumption
driven by just three chip makers: AMD, Intel, and Nvidia[xix]

AI service providers, especially cloud service providers, should consider capital investments into private power generation as public grid operators may be constrained at peak times. The costs will be high: On 21 May 2024, Ben Fowke, Interim Chief Executive Officer and President, American Electric Power Company, testified before the Senate Committee on Energy and Natural Resources that power generation construction costs will run businesses hundreds of billions of dollars.[xxiv]

↑ ~30%
Microsoft's CO_2 emissions
since 2020 due to data center expansion[xx]

↑ ~50%
Google's GHG emissions
between 2019 and 2023, largely due to data center energy use[xxi]

Environmental impact: Goldman Sachs research shows that in its current state, one Chat GPT query "needs nearly 10 times as much electricity to process as a Google search," and by 2030, carbon dioxide emissions may double due to increased GenAI usage.[xxv]

Most of the power consumed by AI models goes into training the models. A third of that consumption is attributed to the inference capabilities these LLMs provide.[xxvi] Institutions

committed to environmental sustainability should consider investing in more efficient technologies, as well as leveraging alternative power solutions for data centers.[xxvii]

Short-, Medium-, and Long-Term Considerations

With those broad considerations in mind, we can now examine potential short-, medium-, and long-term risk scenarios and considerations.

Please note that we look at these as broader milestones – not fixed timespans (i.e. three months, six months, one year) – along your organization's course to AI integration *and* alignment with the evolving AI landscape for the broader financial sector.

Short term: When AI achieves the first real, demonstrable impacts — either positive or negative — that your organization's planning for now

Medium term: When your current AI initiatives have gained traction and infrastructure and regulatory trends are taking even higher priority.

Long term: When AI is no longer novel to your organization or the sector. It's just another fixture of the business world, like the internet is now.

Short-Term Considerations: Augmented Workforce and AI Deployment

The short-term implications are primarily concerned with how AI will transform the 'way we work.' The FS-ISAC AI Risk Working Group projects that there will not be a significant shift in the labor force in the short term. Rather, we foresee greater interest in testing and learning how the labor force can be augmented with GenAI. These decisions will lay the foundation for future developments.

Questions specific to the short-term outlook that your organization can use to spark important dialogue:

- Which GenAI use cases are likely to be deployed first?
- Which use cases will have the greatest impact on the labor force?
- How can you assess labor force impacts from GenAI solutions?
- Is FS-ISAC's projection that we will likely not see a substantial impact on the labor force applicable to your organization?

Key Short-Term Considerations

Labor: Increased efficiency and productivity, notably in cybersecurity functions.

- GenAI could augment the workforce and enable cybersecurity functions to become more efficient, diminishing the impact of labor shortfalls in monitoring, detecting, and preventing cyber attacks.

Deployment: Financial firms will have to determine and articulate the economic case for GenAI deployments.

- Stakeholder priorities will vary – increasing output volumes, increasing the integrity of the work, cutting costs, etc. – which will add complexity to business decisions for deployment.
- Financial institutions' risk-averse nature and the sector's lack of universally accepted AI risk and control frameworks may hamper GenAI adoption.
- Vendors' increased adoption of GenAI may indirectly impact firms that did not deploy their own GenAI or establish guiding frameworks in their organizations.

Medium-Term Considerations: Labor, Data, Cybersecurity, and Legal Challenges

Incremental innovations that will enable more GenAI autonomy, delivered in alignment with other research and development efforts currently underway, are expected for the medium term. The IBM AI Roadmap predicts paradigm shifts for AI year over year, which will increase efficiency and broaden use cases.[xxviii] As GenAI technologies advance in the medium term, we expect broader questions about demand saturation.

The following questions can help your financial institution examine your medium-term outlook:

- What kinds of new workforce opportunities might arise from an AI-augmented future in your firm?
- What are the kinds of AI innovations that you should consider in this time frame?
- How might different regional environments give rise to significant differences in regulatory requirements that you'll have to navigate?
- Do projections, especially regarding regulatory fragmentation and potential demand fluctuation, apply to your organization?

Key Medium-Term Considerations

Labor: The first shifts toward the full replacement of human labor with AI in limited circumstances.

- This will likely set the stage for larger-scale labor changes that will be more fully realized in the long term.
- Intern programs and apprenticeships may evolve as technology-related programs turn their focus to AI oversight tasks like reviewing code, training data prep, and reviewing AI output.

Agentic AI: The current limitations of infrastructure and training model sets will probably increase agentic AI technologies in financial services.[xxix,xxx]

New processes: GenAI could become a familiar feature, complete with risk and control frameworks, enabling the ubiquitous use of AI for new processes.

Model collapse: As more and more data created by GenAI sources is re-incorporated into training data, data diversity becomes a significant problem.

- Homogenous datasets can perpetuate bias and cause LLMs to reinforce their own outputs. Over time, the model becomes less able to generalize to existing and potential issues and less effective in business applications.
- Confidence in outputs will be reduced as models collapse.

More sophisticated attacks: Cybercriminals will use AI to strengthen and improve attacks.

- If AI becomes able to invent new cyber threats, as opposed to the current ability to augment or accelerate such threats, then whole new cybersecurity roles may arise in the sector.

Legal challenges: Issues around the unlicensed use of scraped internet data will have implications for the further use of curated data scraping.

- Issues such as those raised in *OpenAI vs the New York Times*[xxxi] will raise questions about the way models are built and used. The trial, regardless of its verdict, is likely to be the first of its kind as other legal challenges may arise in the medium term.

Regulations: The regulatory environment, especially globally, will remain fragmented.

- Requirements will differ, perhaps even conflict, between countries and between US states.
- Ever-changing political environments will likely exacerbate such changes.
- Financial institutions will have to continually review their governance and compliance related to AI; organizations may need to ensure that data can be segmented into various learning models to ensure achievable separation, removal, or anonymization as per varying regulatory requirements.

Fewer false positives in security alerts: AI will better isolate genuine threats.

- Employees will have less response training, so hands-on experience with security platforms will mostly occur in times of crisis when an organization can least afford an inexperienced response.

Long-Term Considerations: Code, Competition, and Complexity Issues

It is difficult to make accurate predictions of the long-term implications of GenAI. Funding and interest in AI research may be considerably reduced or the technology may experience exponential growth. It's also possible that a new technology may emerge that renders AI archaic, or AI's value may reach a point of diminishing returns. Consider the possibility of an "AI Winter" when planning for long-term investments in permanent infrastructure and a continuation of challenges.

Below are questions that can help your institution frame its long-term outlook:

- What kind of AI innovations should you focus on?
- What risk appetite do you want for your future of AI, and what strategy should you embrace to get within that appetite?
- How should you handle uncertainty versus prescriptive insights?
- How might regional regulatory environments give rise to significant differences in this space over time?
- Is FS-ISAC's projection, especially regarding the AI feasibility for smaller and larger organizations, reasonable for your organization?

Key Long-Term Considerations

Labor: Threats could materialize if too few employees are able or willing to handle manual workloads.

- Continued innovation could eliminate jobs, reduce working hours, or create new labor categories.

Security: GenAI products will have unique security concerns, amplifying entanglement issues.

- Security, interoperability, and other developmental factors will create complexities similar to those introduced by the Internet of Things (IoT).

Overreliance on Code Generation: AI-generated code opens the door to superfluous tools.

- While code generation simplifies software development capabilities, there is potential for the market to become saturated with superfluous tools that don't meet critical SDLC requirements designed by untrained individuals.

Higher barriers to entry for small financial firms.

- Larger, better-resourced organizations will consolidate their AI capabilities making it difficult for smaller in-house AI platforms to compete.

Questions to Prepare You Today for GenAI Future Use

Rather than recommending concrete actions to prepare your firm for future use cases, we suggest that financial services leaders ask the following questions of their institutions right away. The answers may well reveal the best next steps for your organization to undertake and help you get things started.

1. Does your organization have an "AI Champion," either as a role or group (such as a steering committee), to track your AI resources — initiatives still in development, platforms already established, etc.?

Such an entity could ensure that various AI initiatives align with each other, the organization's risk appetite, and your broader business objectives. As a corollary, knowing how to track AI initiatives can also mitigate the risks of 'shadow AI' — i.e., AI developments that, like Shadow IT, have no approval or oversight.

Ideally, AI Champions are poised to consider the operational risks surrounding AI usage: increased energy needs, additional regulatory requirements, and all the other considerations detailed throughout this publication.

Shadow IT: IT systems deployed without the IT department's knowledge — SaaS apps and services are common examples — that evade the IT department's restrictions. Though the intent is usually benign, shadow IT poses regulatory and security risks in financial services, especially regarding data loss, leakage, and use.

2. What mechanisms does your organization have for identifying and mitigating Shadow AI? Could your existing Shadow IT mitigations be expanded to seek out unapproved AI platforms already operating on existing infrastructure?

It's not unknown for business units to rely on software and other deployments that the IT department doesn't oversee – or even know about – in significant revenue streams.

Even without an AI Champion, your organization needs to mitigate shadow AI just as it needs to mitigate shadow IT. Best practice is to assess any newly discovered AI usage to discover vulnerabilities or legal exposures.

3. Has your organization already begun considering the risks surrounding AI? Does it assess new technology deployments in general – for example, are energy use, regulatory alignments, ethics, or other aspects covered by this paper being considered?

Many regulations in our sector call for assessing cyber-relevant risks of new deployments, and your organization's risks may span beyond confidentiality, integrity, or availability. Your enterprise risk management group (or equivalent) may already have a robust register of things to assess; a conversation with them and your AI subject matter experts would show diligence and help you thoroughly assess your AI deployments.

It might also help identify where the use of AI might intersect with a model risk management framework in institutions using one.

4. Are AI risks embedded in your risk and control frameworks?

Embedding AI risks into your control sets is a necessary practice. Your deployments or general use of AI probably won't incur risks that weren't previously identified in the register; perhaps no new control investments are needed either. But it's still healthy to assess your organization's AI-relevant risks, and then to update the risks and control frameworks as needed – or simply formally document that your current framework is sufficient for AI already.

5. How does your organization train employees to interact with and support AI?

While your organization may have robust cybersecurity to thwart attacks from the outside, your organization has insiders who could do significant damage by mistake.

Training employees can reduce the risks of misuse, and it can also reduce data leakage concerns. Remember that any employee could be tempted to put sensitive, proprietary data on a publicly available AI platform, even if just to run additional analytics on it.

6. Does your organization look for diversity in its training data?

Diversity could help mitigate the risk of your models training on their own output or a set of models being trained on each other's output. Insufficient training data diversity could turn outputs into echo chambers that feed directly into critical business decisions.

7. Does your organization deploy GenAI to produce new content? Are your models trained on data that includes copyrighted material?

If the answer is "yes" then there's a risk that your organization has committed IP laundering, as detailed previously in this publication.

8. Is your organization planning for the impacts of AI, including the impacts described in this publication, on your workforce?

Some roles will be made redundant, and other roles will become vital. As GenAI technology develops and business requirements change, employee skill sets and training must keep pace. We strongly recommend considering this now as your organization prepares for its most optimized future.

Conclusion

We agree with those who suggest that the future is perpetually in motion. However, we believe this paper provides an "all-hazards" approach for the future of AI and a framework to help you think through the unknowable changes that will come in the future. However, it's worth noting that those changes fit a perspective with which you're probably familiar: risk. You have likely experienced meetings in which cybersecurity teams urged caution while business managers lobbied for rapid deployment to maximize competitive advantage. Looking at it that way, charting your course for AI will take you through familiar territory.

We encourage your firm to use the contexts and questions above to generate meaningful assessments in the short term and next steps towards achievable objectives that are more clearly understood – from the server room to the board room – with deliverables that inhibit risk and promote useful outcomes. Your resulting deliverables may include:

- A formal framework and guidelines for AI usage
- Prioritized use cases
- Impacts mapped back to the people, processes, and technology that can mitigate risks
- Assessments against AI infrastructure and tools
- Decisions regarding roles and responsibilities and how the workforce will interact with newly implemented GenAI automations
- Clear mission statements about increased efficiency, reduced costs, increased business, etc.
- A risk register that is better adapted for the 21st-century
- Achievable longer-term strategies to determine and then realize your competitive advantage in this space

We hope this publication causes you to bring your experiences and feedback to future engagements with FS-ISAC's AI Risk Working Group. We can never expect to be finished looking at the future, and we hope to incorporate your concerns and perceptions in our ongoing work.

Appendix: Copyright Issues

Challenges to the legitimacy of AI copyright protections will continue with the increased adoption of AI technologies (i.e. Thaler V. Perlmutter), especially where organizations seek to copyright work generated from AI models.[xxxii]

The US Copyright Office deems "that copyright can protect only material that is the product of human creativity. Most fundamentally, the term "author," which is used in both the Constitution and the Copyright Act, excludes non-humans." [xxxiii]

However, the same document allows for technological assistance (the machine as a tool) to create copyrighted material if the material is created by the author's own mental conception.[xxxiv] Jurisdictional laws and ongoing legal cases will likely apply to ideas or documentation that LLMs or GenAI will help to create.

IP laundering may emerge when IP is unintentionally used within organizational outputs and embedded in proprietary processes or documentation. Be prepared to deal with the legal and moral implications of IP laundering, with controls in place to map data lineage and data supply chain through GenAI models to outputs from them.

Many examples of copyright infringement in the curation of marketing materials are directed at AI developers. For instance, artists may sue AI developers who use their art in training data (that is later reproduced by users of the model) without explicit permission.[xxxv] US law allows original content creators to sue users or organizations, especially if the art is recognizable in the artist's creative expression and is being used by an organization for profit.[xxxvi]

References and Resources

[i] Though we refer to these models as "Generative AI," there are significant epistemic questions as to whether these models actually "generate" new content or are simply very complex "stochastic parrots" that regurgitate human generated content in nicely packaged content. See e.g. Bender et. Al: *Bender, E.M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21). Association for Computing Machinery.* Retrieved from http://dl.acm.org/doi/10.1145/3442188.3445922

[ii] Forbes Technology Council. (2025). Foundation models and LLMs: 19 real-world practical use cases. *Forbes*. Retrieved from https://www.forbes.com/councils/forbestechcouncil/2025/02/05/foundation-models-and-llms-19-real-world-practical-use-cases/

[iii] Anthropic. (2023). *The Claude 3 Model Family: Opus, Sonnet, Haiku* [Model Card]. Anthropic. https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf

[iv] OpenAI. (n.d.). Learning to reason with LLMs. Retrieved from https://openai.com/index/learning-to-reason-with-llms/

[v] OpenAI. (n.d.). Learning to reason with LLMs. Retrieved from https://openai.com/index/learning-to-reason-with-llms/

[vi] Harvard Data Science Review. (n.d.). *AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap*. Retrieved from https://hdsr.mitpress.mit.edu/pub/aelql9qy/release/2

[vii] Matuchniak, T. (2010). *The "Good," the "Bad," and the "Ugly" of AI as a risk communications tool*. Retrieved from https://www.emerald.com/insight/content/doi/10.1108/S2040-726220220000025008/full/html#:~:text=Matuchniak%2C%202010

[viii] MIT Sloan School of Management. (n.d.). *Machine learning explained*. from https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained

[ix] IEEE CAI. (2024). *LLM-Assisted Crisis Management: Building Advanced LLM Platforms for Effective Emergency Response and Public Collaboration*. Retrieved from https://ieeecai.org/2024/wp-content/pdfs/540900a862/540900a862.pdf

[x] Council of Europe. (n.d.). The Framework Convention on Artificial Intelligence. Retrieved from https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence

[x] European Commission. (n.d.). *Artificial Intelligence Act*.

[xi] European Commission. (n.d.). *Artificial Intelligence Act*. Retrieved from https://artificialintelligenceact.eu/

[xii] Monetary Authority of Singapore. (2018). Principles to Promote Fairness, Ethics, Accountability and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector. Retrieved from https://t1.daumcdn.net/brunch/service/user/2rc/file/dFjsAgzWVzQk_tg8IL415h3l-28.pdf

xiii Center for Strategic and International Studies. (2023). *Japan's Approach to AI Regulation and Its Impact on the 2023 G7 Presidency*. Retrieved from https://www.csis.org/analysis/japans-approach-ai-regulation-and-its-impact-2023-g7-presidency
xiv European Parliament. (n.d.). *EU AI Act: first regulation on artificial intelligence*. Retrieved from EU AI Act: first regulation on artificial intelligence | Topics | European Parliament]
xv European Parliament. (n.d.). *EU AI Act: first regulation on artificial intelligence*. Retrieved from EU AI Act: first regulation on artificial intelligence | Topics | European Parliament]
xvi OECD. (n.d.). AI principles. Retrieved from https://oecd.ai/en/ai-principles
xvii United States Courts. (n.d.). *Case 1:22-cv-01564*. Retrieved from https://www.govinfo.gov/app/details/USCOURTS-dcd-1_22-cv-01564/USCOURTS-dcd-1_22-cv-01564-0

xix Kindig, B. (2024). AI power consumption rapidly becoming mission-critical. *Forbes*. Retrieved from https://www.forbes.com/sites/bethkindig/2024/06/20/ai-power-consumption-rapidly-becoming-mission-critical/
xx NPR. (2024). *Three Mile Island nuclear power plant and Microsoft AI*. Retrieved from https://www.npr.org/2024/09/20/nx-s1-5120581/three-mile-island-nuclear-power-plant-microsoft-ai
xxi Google. (2024). 2024 environmental report. *Google Blog*. Retrieved from https://blog.google/outreach-initiatives/sustainability/2024-environmental-report/
xxii The results of the pending New York Times v. OpenAI lawsuit may provide some clarity on this issue.

xxiv U.S. Senate Committee on Energy and Natural Resources. (n.d.). Retrieved from https://www.energy.senate.gov/services/files/7F2AC3C4-87CB-4562-99F8-5BB999FC6433
xxv Goldman Sachs. (n.d.). *AI poised to drive 160% increase in power demand*. Retrieved from https://www.goldmansachs.com/insights/articles/AI-poised-to-drive-160-increase-in-power-demand
xxvi U.S. Senate Committee on Energy and Natural Resources. (n.d.). *Document title*. Retrieved from https://www.energy.senate.gov/services/files/7F2AC3C4-87CB-4562-99F8-5BB999FC6433
xxvii World Economic Forum. (2024). *Generative AI and energy emissions*. Retrieved from https://www.weforum.org/agenda/2024/07/generative-ai-energy-emissions/
xxviii IBM. (2025). *AI Roadmap*. [PDF] Available at: https://www.ibm.com/roadmaps/ai.pdf
xxix IBM. (n.d.). *AI Roadmap*. Retrieved from https://www.ibm.com/roadmaps/ai/
xxx Humansdotai. (2024). A new AI era: Agentic AI. *Medium*. Retrieved from https://medium.com/humansdotai/a-new-ai-era-agentic-ai-2cfe4f0635ea
xxxi Harvard Law Review (2024) 'NYT v. OpenAI: The Times's About-Face', *Harvard Law Review Blog*, 4 April. Available at: https://harvardlawreview.org/blog/2024/04/nyt-v-openai-the-timess-about-face/
xxxii Center for Art Law. (2023). *Case summary and review: Thaler v. Perlmutter*. Retrieved from https://itsartlaw.org/2023/12/11/case-summary-and-review-thaler-v-perlmutter/
xxxiii U.S. Copyright Office. (n.d.). *AI policy guidance*. Retrieved from https://www.copyright.gov/ai/ai_policy_guidance.pdf
xxxiv U.S. Copyright Office. (n.d.). *AI policy guidance*. Retrieved from https://www.copyright.gov/ai/ai_policy_guidance.pdf

[xxxv] Voon, C. (2023) 'Class Action Lawsuit Filed Against AI Generators DeviantArt, Midjourney, and Stable Diffusion', *Artnet News*, 17 January. Available at: https://news.artnet.com/art-world/class-action-lawsuit-ai-generators-deviantart-midjourney-stable-diffusion-2246770
[xxxvi] Business Insider (2024) 'How Artificial Intelligence and Intellectual Property Will Affect Marketing', *Business Insider*, 12 June. Available at: https://www.businessinsider.com/how-artificial-intelligence-intellectual-property-will-affect-marketing-2024-6

