AI SIMPLIFIED

INSIGHTS YOU CAN USE

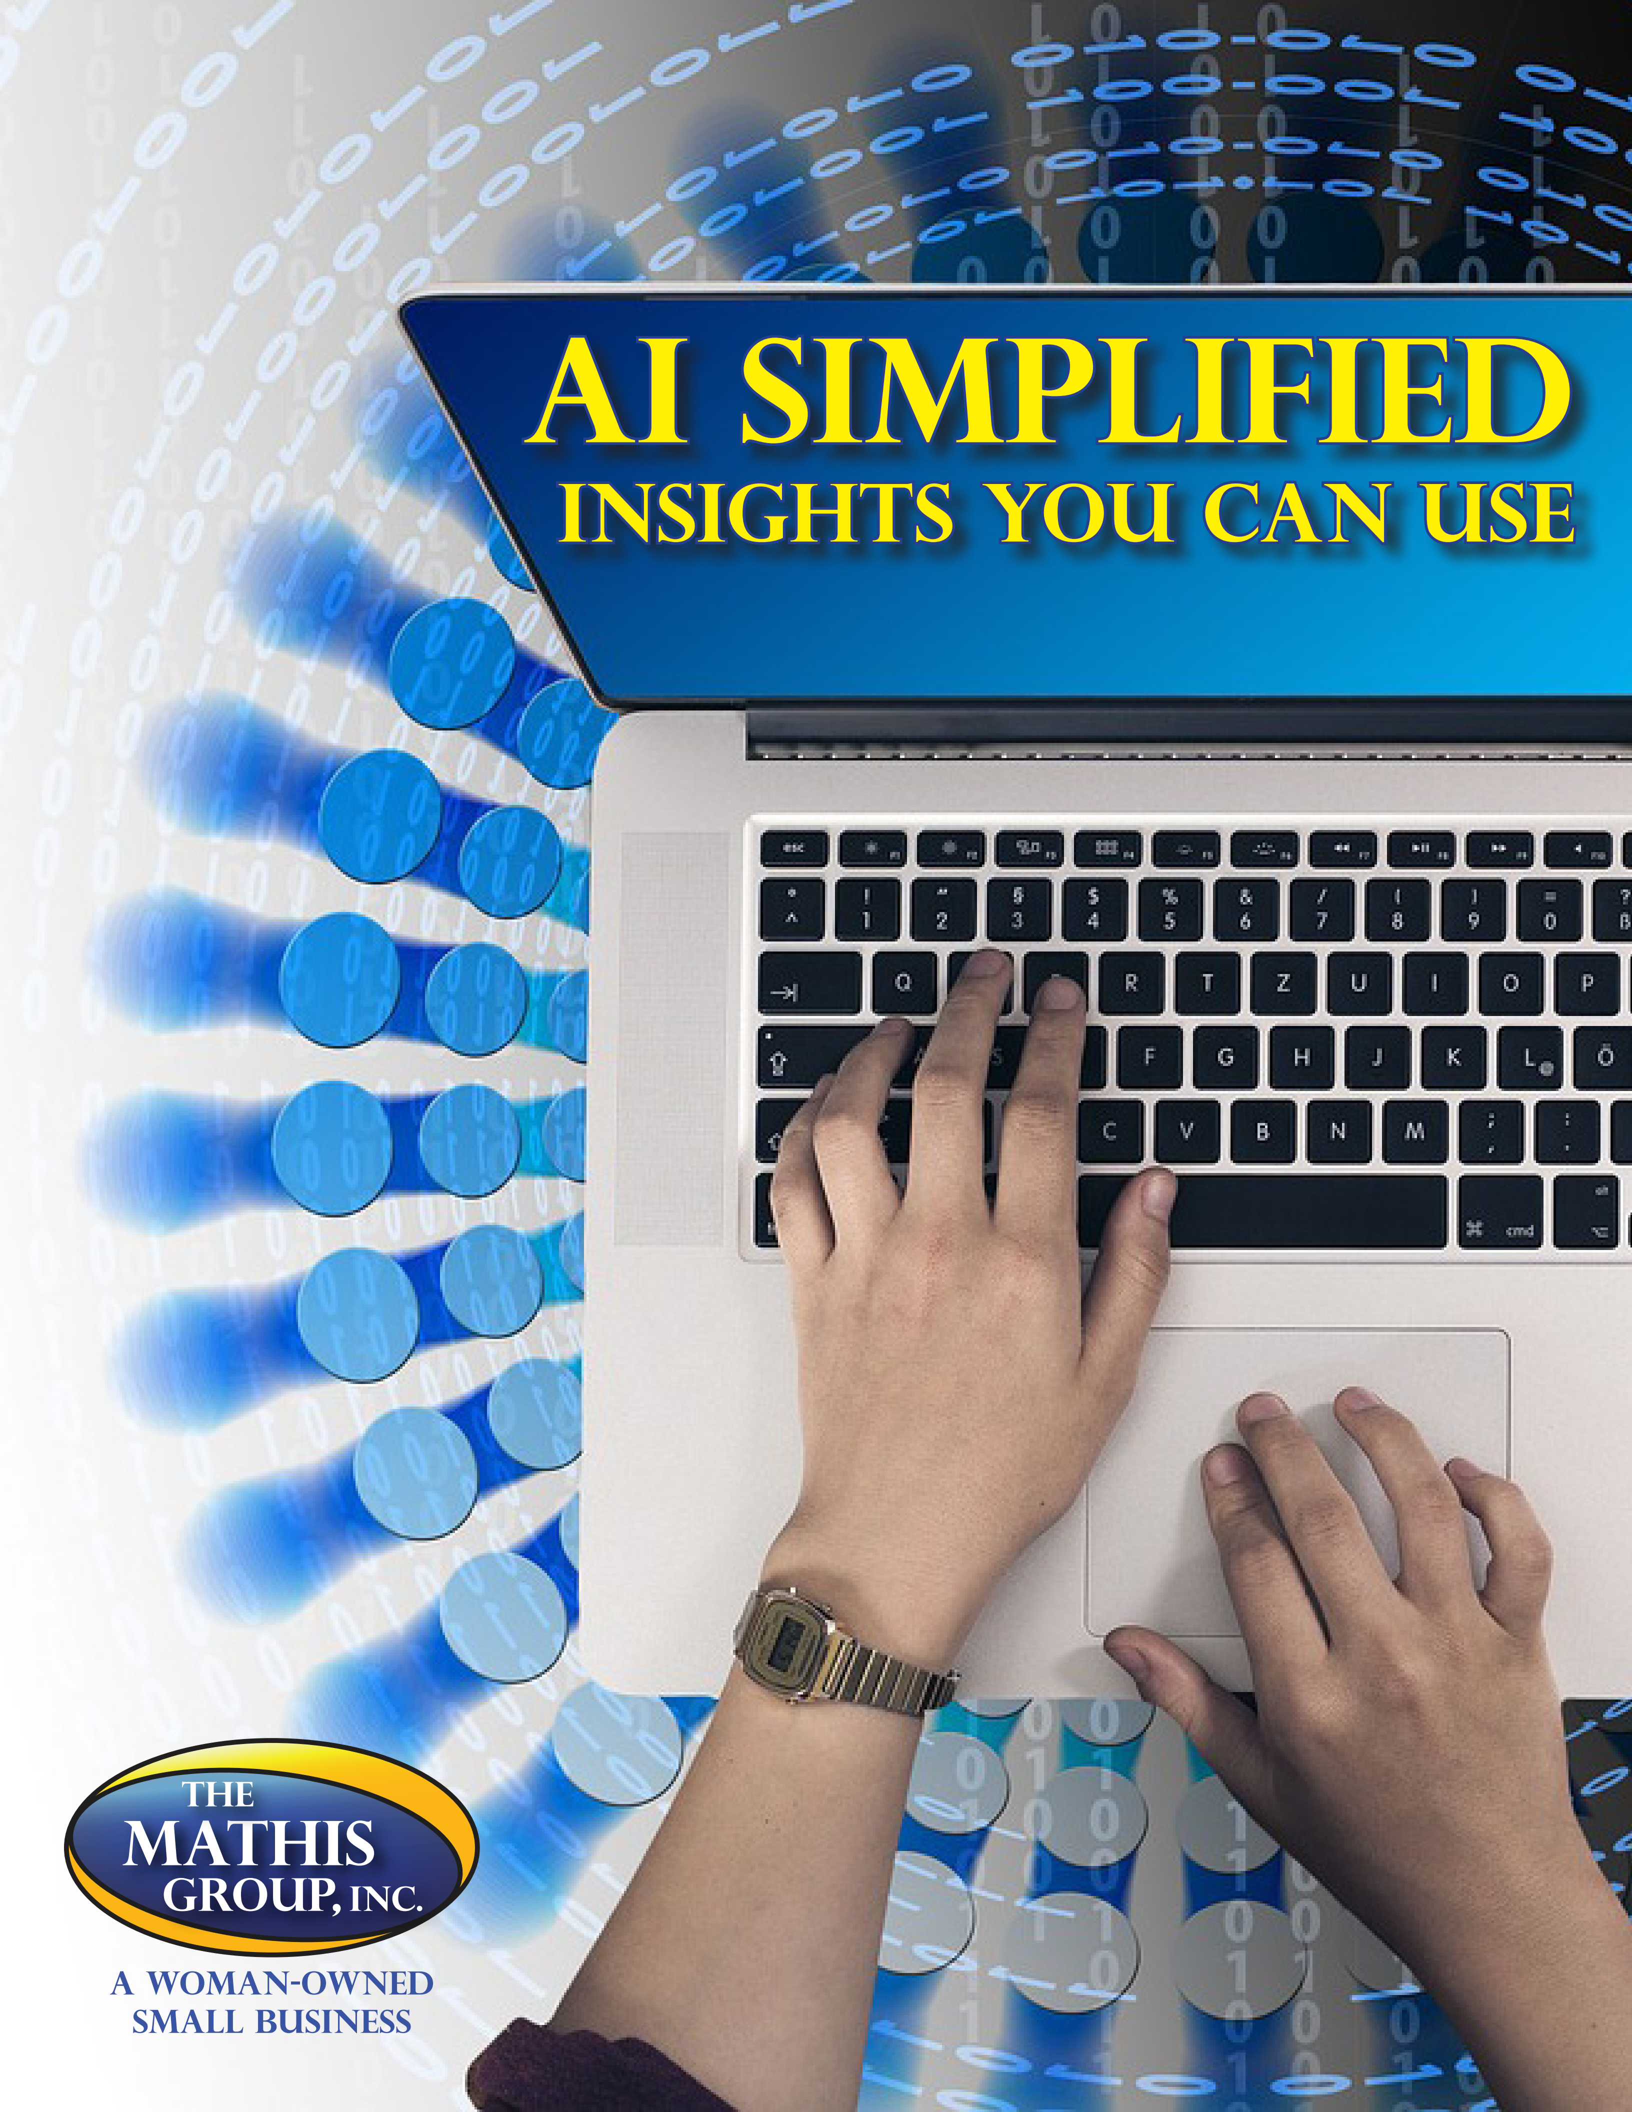

THE MATHIS GROUP, INC.

A WOMAN-OWNED SMALL BUSINESS

FEBRUARY 2026
AI FOR RISK FORECASTING: PREDICT ISSUES BEFORE THEY DERAIL THE PROJECT

Artificial intelligence is reshaping the risk management landscape by leveraging predictive capabilities to forecast, monitor, and anticipate responses to risks and to identify solutions to mitigate them. Risk forecasting in project management has focused most of its research and writing on historical risk and on how it was adjusted to achieve success. Using historical information will help in some projects, such as software projects. However, implementing AI is more of a data project that does not follow the same processes, restrictions, and constraints as other projects. This newsletter will examine risk and use AI to assist with forecasting and management, and to predict how AI will work in a given situation.

AI transforms the culture of project management by providing consistent, evidence-based forecasts and serving as an assistant to anyone assessing current risks and planning how to address them. In this way, AI becomes not just a tool for forecasting but a catalyst for disciplined execution and long-term project success.

HOW DO PREDICTIONS HELP RISK FORECASTING?

Risk predictions and forecasting will change how one conducts risk analysis and develops project mitigation and contingency plans. The ability of AI to anticipate which risks will occur, or to give project managers and teams advance notice that a risk is happening sooner than initially expected, enables faster responses to the challenge. A rapid response reduces the risk that the issue will escalate and harm other areas of the project. For example, a risk log anticipates risks associated with deploying an AI system and changing the culture. AI can warn of a negative trend earlier or identify an opportunity in how the culture is changing faster than expected. In either case, the data and AI recommendations will influence the next steps.

The difference between today's risk analysis and AI-driven analysis lies in the level of predictability and forecasting. In risk analysis, decisions are made more quickly when time is limited, enabling risk responses to be implemented faster than AI, which examines historical data across the project and predicts future adjustments. Implementing these recommendations can keep humans firmly in charge of solution selection. Research shows that dynamic indicators, not static lists, are better at surfacing early warnings. In practice, teams maintain their register for accountability while allowing forecasts to prioritize which risks to review first. This combination boosts vigilance without creating unnecessary tasks (Vanhoucke, 2012; PMI, 2021).

Guardrails, such as ethics and governance, must be examined and created for each organization. The ability to use AI responsibly in project settings requires adherence to the rules and regulations that govern responsible conduct. Some of these guidelines will eventually be dictated by governance oversight, while each company will customize other policies. Acting responsibly builds confidence and support in the project world while giving team members hope for a new level of efficiency that has not been achieved or anticipated.

FORECASTING AND PREDICTING POTENTIAL RISK BEFORE IMPLEMENTING AN AI PROJECT

Requirements for artificial intelligence projects differ from those of a typical software project. Organizations that treat them as standard software projects face hidden risks. In this section, we identify seven significant risks that AI projects face. Some of these risks are similar to those in software projects, while others require a different level of analysis and prediction. Project teams must plan for and, where possible, mitigate these risks to reduce the likelihood of failure and increase the likelihood of success in implementing the AI system.

The following graphic demonstrates some risks associated with an AI project. Because the project is viewed as a data project rather than a software development project, the risk profile may differ from what one anticipates. This graphic highlights topics covered in this newsletter and may not apply to every artificial intelligence project. However, we have identified a reputational risk from the lack of explainability in AI, as it is a common challenge that most AI projects must mitigate to succeed. The lack of explainability will lead many to be cautious, as they may misunderstand how AI makes decisions and recommendations. The reputational risk of AI errors or misuse is a significant concern for any AI project in a high-risk area, often leading many to delay implementation out of fear of what could happen if it does not work correctly.

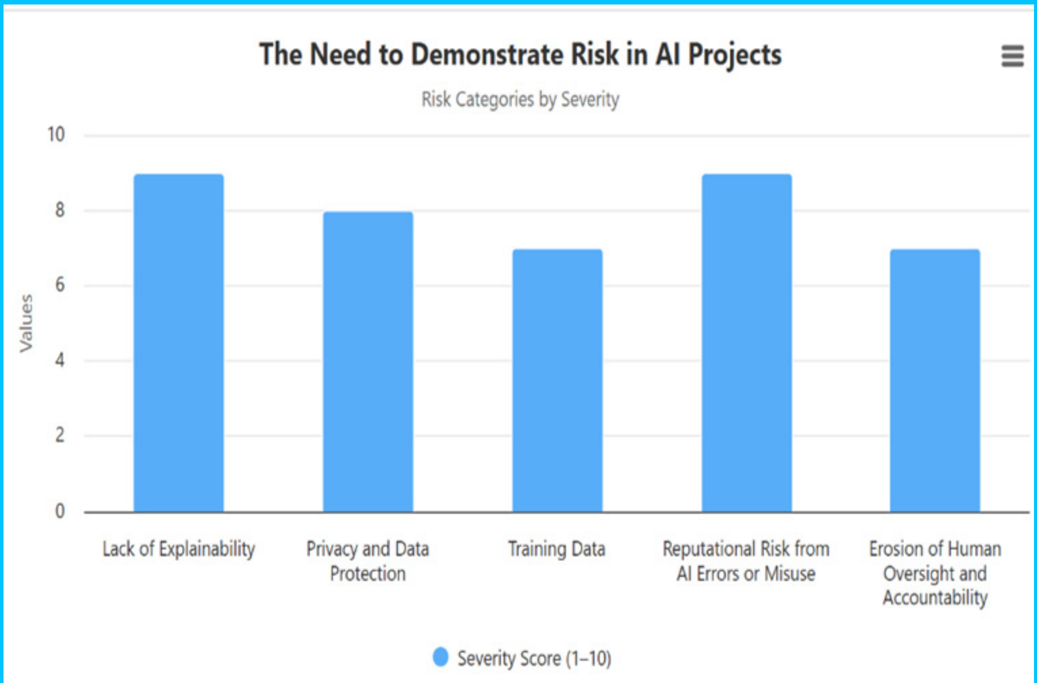

Data Quality and Algorithmic Bias

In all aspects of AI data and the examination of AI systems, the team and organization must anticipate that bias will, or at least can, surface. Bias is not theoretical; it can lead to discrimination and potential legal action by individuals whose rights the system violated by failing to consider them. This form of bias has often occurred in banking, with lenders refusing loans to qualified applicants based on their occupations or education levels. When using AI to review resumes and employment applications, there can be bias in candidate scores, with some candidates receiving higher or lower scores due to factors in their resumes or job histories. Bias can distort the data or analysis in these documents, creating a false impression. For this reason, organizations bring or retain a human-in-the-loop to monitor quality and an authoritative verifier to ensure the output is consistent and unbiased.

Any bias in the system erodes trust in AI and can hamper the evolution of these tools and the explanation of how and why this technology makes those decisions. If bias goes unchecked, the organization risks hiring the wrong people, failing to hire the most qualified people, or turning down qualified candidates for loans or other decisions, which could be devastating. Risk must be intentionally mitigated to prevent harm from this technology.

Lack of Explainability and Transparency

Both explainability and transparency help build stakeholder trust in AI. Explainability is the ability to understand how and why an AI system produces specific outputs. In contrast, transparency refers to understanding how an AI makes a decision or the steps a model takes to create a particular answer. Understanding algorithms and models helps build trust and confidence in using AI and in the answers it provides. A black box provides little or no information about why or how it makes decisions, offering little insight into the mechanisms it uses.

The following graphic compares the lack of explainability and transparency. These two topics are often discussed together, and some people use these terms interchangeably. This graphic shows that there are various types of openness to consider, such as the outcome process model and data, and to ensure that everyone understands why the AI tool responds in a specific manner. The better an organization can address explainability, the stronger the support it typically receives from stakeholders.

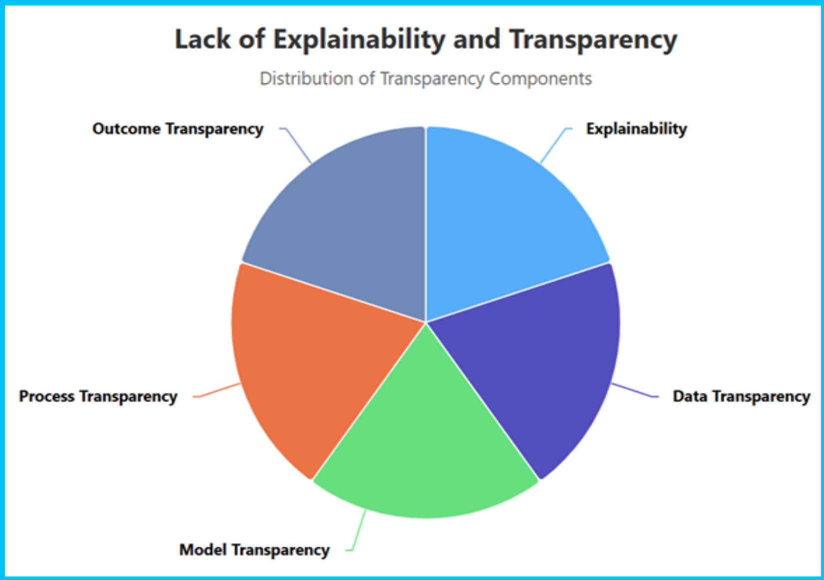

Explainability

Explainability highlights the AI decision-making process, enabling humans to understand how the model works. Strong explainability helps focus on the process by which the system generates a specific output or decision. However, explainability is challenging because models and algorithms can drift with new data and its effects. These adjustments affect the trade-offs between accuracy and interpretation. When organizations require higher explainability, they must trade off faster decision-making for stronger decisions that take longer.

Transparency

Transparency focuses on showing users the AI lifecycle, documenting the training data and providing behind-the-scenes insight into the process a tool uses to make decisions. Demonstrating the system’s transparency can be challenging because much of the technology is hidden. Transparency can focus on many different areas. Increasing transparency can be achieved by focusing on data, model, process, and outcome transparency.

Data transparency

Data transparency examines the training data used to train the model. The training data may not be as accurate as real-world data, leading to different outputs. Training data can come from internal documents selected to provide the system with breadth of information, or from a third-party vendor to supplement the internal data. Differences between training and real-world data can arise from how the system interprets bias or from differences in the computing power available at the time of training.

Model transparency

Model transparency explains how AI systems and algorithms operate and analyze data to produce outputs. Some outputs are clear and directly tied to the training process and content, such as data transparency. However, architecture and design affect the model outputs. Organizations want the model to be a "Whitebox," meaning it is clearly understandable across all areas; however, most function more like a "Blackbox," with hidden components and no explanation of how decisions are made. Making the model as clear as possible can benefit everyone in the decision-making process.

Process transparency

What are the steps and flow of the process? Are there decision points in the process? Where does the process decide, and where are humans involved? Does the process incorporate live information, or does it make the decision based solely on historical data? Understanding when the tool is making the decision and when a human is involved in influencing it is critical. Understanding the process can include the approval process and monitoring to shape the AI aspects.

Outcome transparency

Open tracking and reporting of the AI system in the real world constitute the best definition of outcome transparency. This includes using metrics to assess the population and auditing errors and bias to ensure that the system operates as intended. Additionally, if the system provides distorted decision feedback to the team, it may indicate a need to adjust the prompts used in the early stages of the process.

Privacy and Data Protection

Stakeholders understand their data must occasionally be transmitted electronically, but most want it to happen only when necessary and, at other times, for all personal information to be anonymous. AI systems typically rely on large datasets or third-party commercial data that may contain sensitive and personal information. This creates a serious privacy risk, as it makes it difficult to protect personal data from mishandling and breaches consistently. Sometimes, the assumption is that when buying from a third party, the vendor bears the product risk, but this is not always the case with AI. When implementing an AI system, the organization must examine the potential and protect the privacy of all parties.

Not understanding how much data is needed.

Artificial intelligence requires substantial data to achieve the goals outlined in the project plan. Many organizations overestimate their data requirements, assuming they need millions of examples to train an AI system when they need precise information that may be far less than they anticipate. Techniques such as transfer learning and few-shot learning can produce effective models that perform well on much smaller datasets than one might expect. Many organizations purchase data from third-party data warehouses to resell to their customer base, freeing them to focus on creating and cleaning data for sale.

Determining data volume is not a one-size-fits-all approach across organizations or AI systems, because some organizations that rely more on visual, unstructured data may need to load thousands of images and graphics to train their systems to interpret them accurately. In contrast, another organization might use only one or two small databases to train its system to generate text output efficiently.

Lack of understanding of the specific types of data.

Because AI is advancing rapidly and numerous myths circulate, it is difficult for people to understand the specific types of data they need to begin using AI. The kind of AI associated with their system will depend on the organization's AI use case. If AI is going to function as a customer service assistant, it will be very different from creating marketing materials and slogans for media, social media, and marketing campaigns.

Different applications require fundamentally different data structures, which supports the idea that one size does not fit all. Each organization will need different resolutions and diversity, and natural language processing requires taxonomies that capture linguistic variation in context. Predictive analysis requires historical data to inform past organizational decisions and to generate projections that support the project and the project manager. Mismatched data can be a serious issue, potentially dooming the AI project or its predictions.

Not determining or identifying the internal or external data.

Organizations often fail in inventorying the valuable data they already possess. For years, organizations have produced large volumes of data, only to archive it in various formats and structures with little or no consideration for future use. The guidelines for most organizations focused on archiving data rather than on the structure and types of data required to maintain historical documents. Existing customer interactions, operational logs, transaction records, and intended documents may contain everything an AI project needs; however, many organizations are unsure whether they are archived. Many successful AI initiatives begin by mining information already available in company databases and systems and using that analysis to identify data gaps.

The decision between internal and external data should be connected to strategic choices, goals, and outputs. The data the organization stores should include all information necessary to conduct a robust analysis of the project's historical records, success factors, internal data, and decisions. Because internal data provides specificity about the organization's unique business context, operations, and decisions, it is a valuable source for analyzing and predicting its future. Unless an organization maps this information and ensures it retains key deliverables and knowledge, the data is likely to exhibit significant bias and inaccuracies.

Lack of understanding of the data sources and locations.

Another challenge is a lack of understanding of the data sources and locations. Because organizations often store information across departments and divisions, they may lack archiving guidelines or a central location. Without these things, data may need to be duplicated from one division to another for use and to comply with requirements that the data remain in that division's archives. For this reason, many organizations have adopted cloud systems to ensure that the entire organization has access to data and to minimize duplication. When organizations do not know the location of their data, it slows analysis and, in some cases, prevents them from realizing the full benefits of artificial intelligence because it cannot aggregate the specific information and outputs they need. The simple reason is that AI cannot find information when it is warehoused across multiple locations and hidden under layers of divisions.

Different data sources have varying levels of reliability. Some data sources are considered highly reliable due to their frequent updates and governance processes. In other situations, data may be unreliable because it has not been updated for months or years, leading many to assume it is outdated. Failing to understand the data sources and specific characteristics up front can lead to compliance violations due to insufficient protection of personal information and a lack of clarity about what this system will disclose to research parties.

Not understanding the data format.

The structure of the data can also affect artificial intelligence. Not all structures follow the same rules, depending on the type of data and the system with which they interact. There is both structured and unstructured data associated with AI. Structured data typically follows defined formats, whereas unstructured data often takes the form of free text, images, and videos and requires different processing approaches. Given the structured and unstructured approaches, organizations using AI must identify where they can realize the most significant opportunities and benefits from other types of information. AI models can learn from both types of structures, but will they produce the output the organization needs? A project that appears simple because you have the data becomes complex when you realize the data may be scattered across emails, PDF scans, and other graphical files in various locations.

The relationships and hierarchies within your data structure are often more important than the data itself. Sometimes, missing data can be obtained from third parties relatively easily; however, converting all internal data from various formats into a format usable by artificial intelligence can be challenging and costly. Understanding whether the data is a time-series hierarchical relationship or a graph-based structure can determine which AI approach will work and which will fail. For example, treating sequential customer journey data as independent snapshots loses critical temporal patterns, while ignoring the nested structure of organizational data means your model won't understand how departments, teams, and individuals relate to one another. All of these create challenges in determining how and when an organization will use data to its advantage.

Many individuals mistakenly believe that training data is always superior to live data for training AI. The training data typically underpins the model and algorithm, and the live data enables them to operate realistically, incorporating real-world challenges organizations face.

Training Data Compared to Live Data

A common phrase in computer systems is "garbage in, garbage out." This means that when programming a computer system, or in this case, training an artificial intelligence system, if you're using garbage and putting garbage into that system, then only garbage will come out of that system. Comparing training data with live data can be challenging and frustrating.

Training Data

Training data is a historical data set used to teach an AI model patterns, relationships, and decision-making rules during the development phase. Training data is typically compiled by the system's trainers and hand-selected to highlight specific qualities and characteristics and to align with the outputs the tool will produce later. This data has been collected and is typically labeled and annotated with correct answers to help the system understand specific knowledge across various topics. For example, if the tool is trained to recognize trees, it has been trained on thousands of tree images, so when it encounters a tree, it can identify it quickly and accurately. The tree does not appear as a cat or a dog; it seems like a tree only when a specific output is requested.

Training data can be expensive to obtain, sourced from historical transactions and reviewed for gaps or biases, including any missing data associated with that information. Most of the time, when looking at training data, it is looking at historical information or things that have already taken place or happened. The training data includes fraudulent checks because it is trained to detect them as they pass through the system. Using historical data and examples helps the system learn from patterns, and pairing them with responses and human oversight can make this training highly effective.

The following graph compares training data with real-world, live data. Each is a complementary dataset that enables algorithms and models to function correctly.

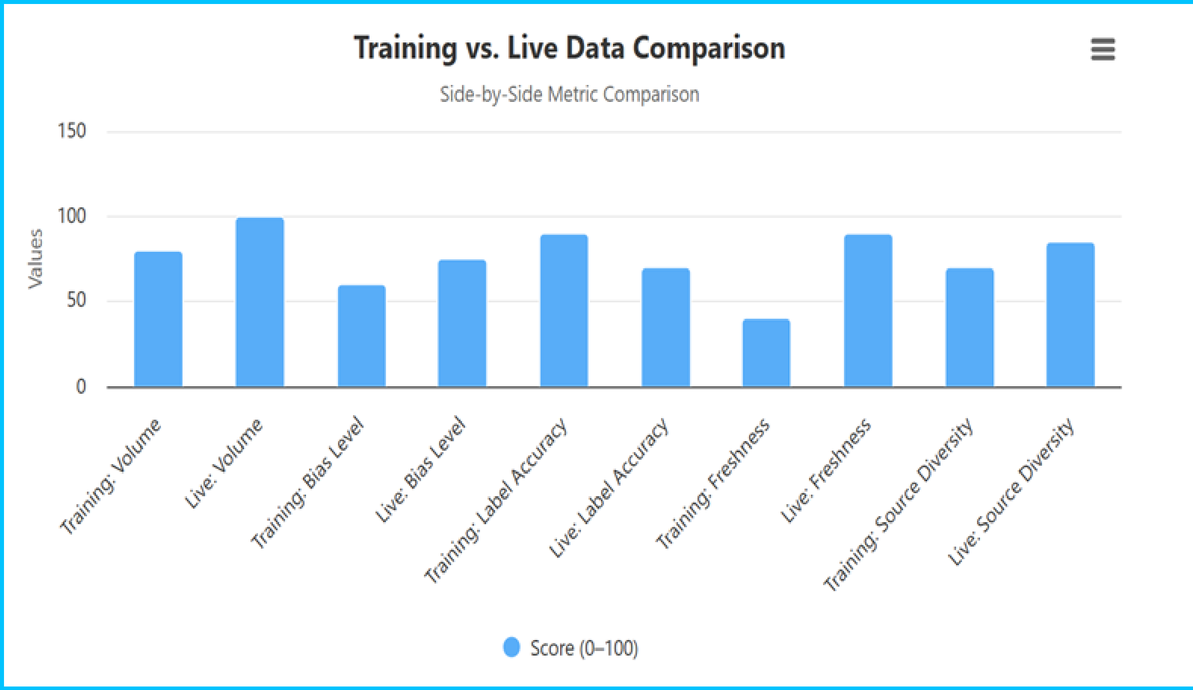

Live data

Live data is the opposite of training data. It is real-time data that is immediately available. This is new data that the model has not previously seen or used to make real-time predictions or decisions. It lacks the luxury of years of historical support and backing, and, in some cases, does not have human oversight or humans-in-the-loop to detect deviations. Live data could be an immediate email from a customer, indicating frustration with a product that is not performing as expected.

In clinical practice, live data may include a new diagnosis made by the physician; however, asking the AI system to review the current symptoms and test results and to generate predictions and recommendations for diagnosis and therapy protocols could be beneficial for this patient. All these live interactions must occur in real time and cannot be delayed by minutes, hours, or days for an output.

When comparing training data with live data, the fundamental difference lies in the shift from controlled to real-world conditions. Training data for fraud detection AI was carefully selected to include balanced examples of both fraudulent and legitimate transactions, ensuring the model can detect both. However, when choosing these examples as a human, it is possible to oversimplify or bias the selection process, without intending to and even being aware that it has happened.

Training data is also typically static and clean. This indicates that the data is correctly formatted, the emails are complete, the information is clear, and all data has been categorized by issue or type. These steps ensure the content is free of typos, formatted correctly, and easy for a system to understand and analyze.

Live data, on the other hand, can be very messy, with missing or incorrect information, misspellings, incorrect data entered, or customers misunderstanding what was being asked of them. Live data can demand that every prediction matters because it affects a real person or transaction, and there's no option but to wait and study something for a period. On the other hand, because training data is historical, others have had time to review and analyze it before it is used.

Reputational Risk from AI Errors or Misuse

Reputational risk concerns damage to the organization's name or brand resulting from an AI tool malfunction. Organizations have received a black eye for implementing an AI tool that advised customers to do something that was not ethically or policy-wise correct. In some cases, the AI tool made recommendations to the customer that did not make logical sense in the business context.

AI tools are not 100% perfect and may contain errors in content research and analysis. Consequently, keeping a human in the loop for high-risk, high-exposure areas is essential. A human-in-the-loop approach gives people greater confidence that the AI tool is being monitored for technical glitches and reduces the risk of the organization being held accountable or becoming a story on the 5 o'clock news. For example, a chatbot that provides incorrect medical advice would be unacceptable for a hospital to rely on when treating patients. Even if you quickly resolve the technical issue, reputational harm persists because people remember that the hospital allowed an AI to misdiagnose an illness. Other organizations have experienced this, such as insurance companies using the AI tool to deny claims. Although the company responded, corrected the claims, and adjusted the AI tool, its reputation was still damaged because the AI tool rejected qualified applicants. People tend to remember the negatives of AI and its impact on them more than the organization's apology for its shortcomings.

Erosion of Human Oversight and Accountability

As organizations increasingly rely on AI, there's a real risk that human judgment erodes, especially when leaders delegate too much decision-making authority to automated systems. When employees consistently see AI systems making accurate predictions and recommendations, they gradually stop questioning the output and begin rubberstamping the decision without genuine review. These are the challenges associated with using an AI system that is not, and will not be, 100% correct. A loan officer who initially reviews AI credit risk assessments carefully eventually begins approving or denying based on the AI tool's recommendation. No longer does this individual view factors through human eyes; instead, they assume the tool has correctly analyzed the application.

A bias toward automation and technology is nothing new. People have accepted what shows up on Google and the Internet as truth, regardless of whether it has any credible evidence. Automation bias refers to humans becoming passive validators rather than active decision-makers. When AI makes an error, no one catches it because everyone assumes the system must be correct. After all, it has been right in the past.

When comparing human decision-making with an AI tool, most people focus on efficiency by measuring how quickly humans can process a task versus how quickly the AI's recommendations are made. Inadvertently, AI makes recommendations in seconds, compared with its human counterpart; however, no one typically examines whether the decision is better than the AI tool's. AI suggests responses built into the system that push for careful human review, which serves as a safety net, and human oversight becomes performance rather than a subjective existence on paper, but not in practice.

When an AI system fails, it becomes nearly impossible to pinpoint accountability because responsibility is distributed among multiple data scientists and stakeholders who approved and endorsed the AI system's output. The data scientist claims they built what was requested using the provided data, the project manager argues that they trusted the technical experts, and the employee says they followed the system's guidelines, which they are trained in. An example is a school system that uses an AI tool to assist its teachers. A teacher allowed AI to create an inappropriate test and then blamed the tool rather than taking responsibility, stating they had not checked the test for credibility and reliability.

Because individuals no longer make decisions and conduct analyses individually, they can experience a gradual loss of knowledge (D-skilling) that undermines their ability. Because the AI tool handles human decision-making and performance so skillfully, it is easy for humans to become lethargic due to a lack of practice and creativity, and to become dangerously dependent on the tool rather than their own expertise. In addition, AI-generated responses may lack communication skills to handle complex situations. When AI encounters a scenario outside its training or makes an error, the human-in-the-loop should catch and correct it. However, if the person no longer has the skills required, they are no longer the right person.

After the AI tool is used for an extended period, the organization may become lax in maintaining documentation and processes for manual operations, should they ever be needed. Because AI is implemented extremely efficiently, it becomes difficult, if not impossible, to continue operating if the AI system fails or must be taken offline. For example, organizations no longer know how to perform non-technical activities because technical tools have been used for so long. An example of this is a cashier who cannot make change unless the cash register tells them how much to return. Making change is a common task for cashiers; however, modern technology tells the cashier exactly how much change to give the customer. If that technology malfunctions, most cashiers will not be able to make change. This is because cashiers become D-skilled in those abilities over time.

Bias

Bias in artificial intelligence is simultaneously a technical, ethical, and governance problem. While data bias can arise from problematic or insufficient data, organizations need to ensure that their models and data are as bias-free as possible when developing an AI tool. The tool mirrors the organization's data and inputs to train and operate the system. Any bias that could harm the hiring process or discriminate against people in promotions or otherwise must be removed from the data to optimize fairness and sound decision-making.

The following graph shows biases in an AI system. However, depending on an organization's objectives, these biases may not align with them. Bias can be expanded to include other areas, such as sampling bias, deployment bias, evaluation bias, feedback-loop bias, and representation bias. For illustrative purposes, the issues addressed here are model, prejudice, and measurement bias, which are common in most AI systems. However, the type of bias your project is experiencing must be identified, and mitigation strategies must be developed to minimize its negative impact.

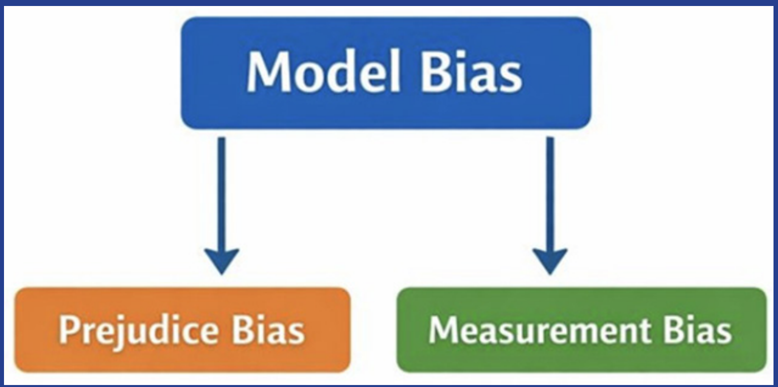

Model Bias

Model bias can come in two primary forms: prejudice bias and measurement bias. Both forms of bias can discriminate against internal and external candidates for promotions or for being offered a job in the first place.

Prejudice bias

Prejudice bias occurs when societal stereotypes are embedded in data features. Because the world is full of biases and stereotypes, it is easy for data to be tainted by those same attitudes and information. Because artificial intelligence draws on global data, it is easy for bias to become embedded in its foundational information and to produce unacceptable outputs in the future.

Measurement bias

Measurement bias occurs when variables misrepresent the output being measured, as in scoring individuals for promotion or hiring decisions. Because individuals are typically evaluated with a numerical metric during hiring, it is easy for AI to review a resume and conduct a preliminary score. If a resume is exceptionally low, a recruiter or HR professional may skip it and avoid wasting time trying to understand why it is low; they may accept the AI tool's score and move on to those with higher scores. If the scoring was misrepresented due to measurement bias, this may be a quality candidate who was mis-scored in the HR system. The HR professional will never know unless they conduct their own resume audits.

The only way to ensure that bias is removed, or at least mitigated, in these situations is to consistently audit the AI tool's decisions and keep humans in the loop. Keeping humans in the loop reduces the likelihood of bias. It increases the possibility that, if bias occurs, individuals can identify and correct it before anyone is harmed or discriminated against.

REFERENCES

Alrasheed, K. A., & Alshammari, I. (2025). Enhancing construction management with machine learning and BIM. *Discover Civil Engineering*.

Baerenbold, R., et al. (2023). Reducing risks in megaprojects: The potential of reference class forecasting. *Journal of Benefit-Cost Analysis*.

Chen, L., et al. (2025). Transparent and reliable construction cost prediction using machine learning with uncertainty quantification. *Sustainable Computing: Informatics and Systems*.

Coffie, G. H., et al. (2023). Toward predictive modelling of construction cost overruns using SVM. *Cogent Engineering, 10*(1), 2269656.

Flyvbjerg, B., Hon, C., & Fok, W. H. (2017). Reference class forecasting: Eliminating bias in project forecasting (Working paper).

Gartner. (2019, March 20). *80% of today's project management tasks will be eliminated by 2030 as AI takes over* (Press release).

Li, C., et al. (2025). A systematic review of learning-based software defect prediction. *Journal of Information Technology*, 36(3), 1–24.

PMI. (n.d.). Integrating risk and earned value management. *PMI Knowledge Library*.

PMI. (2021). *Pulse of the Profession® 2021: Beyond agility*. Project Management Institute.

PMI. (2025). *Pulse of the Profession® 2025: Boosting business acumen*. Project Management Institute.

Reuters. (2025, June 25). More than 40% of agentic AI projects will be scrapped by 2027, according to Gartner. *Reuters Technology*.

Soliman, E., & El-Gohary, N. (2024). Integrating risk management with earned value for better cost/schedule forecasting. *Journal of Project Leadership*.

Stradowski, S., et al. (2023). Machine learning in software defect prediction: A business-oriented review. *Information and Software Technology, 161*, 107257.

Vanhoucke, M. (2012). Integrating schedule risk analysis with earned value management. *The Measurable News*, 2012(1), 13–20.

THE MATHIS GROUP, INC.

A WOMAN-OWNED SMALL BUSINESS (WOSB)

Providing quality, customized training and consulting services that inspire, educate, and equip organizations to be better tomorrow than they are today.

DR. KEITH MATHIS, PMP, PMI-ACP, CSP-SM, CSP-PO
WANDA MATHIS, M.ED. PMI-ACP

PROJECT MANAGEMENT TRAINING

OVER 60 PROJECT MANAGEMENT COURSES REGISTERED WITH PMI

PRESENTATIONS THAT EDUCATE, MOTIVATE, AND INSPIRE

Since 1993, The Mathis Group has been helping organizations change worker productivity and behavior.

PROJECT MANAGEMENT
MARKETING
MOTIVATION
ORGANIZATIONAL BEHAVIOR
LEADERSHIP
CUSTOMER SERVICE

COMPANY MANDATE

The Mathis Group provides training and consulting that will impact the organization and individual while maintaining an outstanding reputation for success and integrity.

VALUES STATEMENT

Every person has worth and should be treated with respect.

AREAS OF EXPERTISE

- Curriculum Design
- Project Management
- Organizational Behavior and Development
- Management
- Agile Project Management
- Strategic Planning
- Executive Coaching
- Performance
- Team Building
- Emotional Intelligence
- Leadership
- Customer Service
- Supervisory Leadership
- Hybrid Project Management

9515 N Spring Valley Dr
Pleasant Hope, MO 65725
800-224-3731
417-759-9110
(voice/fax)

www.themathisgroup.com

keith@themathisgroup.com
wanda@themathisgroup.com

DUNS Number: 007722098
CAGE: 3C1N9
GSA Contractor Number: GS-02F-0010V

GSA Contract Holder

ACCREDITED BUSINESS