



# AI SIMPLIFIED

## INSIGHTS YOU CAN USE



A WOMAN-OWNED  
SMALL BUSINESS

# MAKING BETTER DECISIONS WITH ARTIFICIAL INTELLIGENCE: THE DECISION-MAKING PROCESS AND WHERE AI FITS

This edition explores how artificial intelligence (AI) enhances decision-making in project management. We review organizational decision models, examine AI's role across different decision types, address ethical and practical limitations, and provide advanced AI prompts for both Waterfall and Agile methodologies.

## Overview of Organizational Decision-Making

Organizational decision-making is the process of selecting a course of action to achieve defined objectives. Simon (1977). Artificial intelligence engages in decision-making through a structured yet adaptive sequence that parallels human reasoning while leveraging computational speed and scale. Unlike human decision-making, which is often constrained by cognitive limitations and subjective biases, AI processes can systematically examine vast quantities of data in seconds, iteratively refining their understanding of the problem before generating recommendations.

The process begins with problem recognition, where the AI system identifies and defines the core issue requiring resolution. This stage is critical: the quality and precision of the initial problem definition directly influence the relevance and accuracy of the eventual solution. When a project manager inputs a well-structured description – including context, constraints, and expected outcomes – AI can move beyond surface-level symptoms to identify the underlying root causes. For example, if a project is experiencing delays, AI will not only note missed deadlines but will also assess whether those delays are linked to resource bottlenecks, vendor performance issues, or scope creep.

Once the problem is defined, AI transitions into information gathering. The scope and source of this data depend on the system's deployment environment. Within secure organizational firewalls, AI systems primarily rely on the employer's proprietary datasets, ensuring that recommendations align closely with internal processes, policies, and historical performance. In contrast, open-source or hybrid configurations may incorporate external data sources, industry benchmarks, and best practices from comparable projects. This dual sourcing can significantly broaden the range of potential solutions, but it also requires careful validation to ensure that externally derived recommendations align with the organization's specific context.

Following data acquisition, AI engages in alternative generation and evaluation. At this stage, the system analyzes potential courses of action, applying criteria such as expected effectiveness, implementation speed, and cost-efficiency. AI excels here because it can simultaneously process multiple scenarios, run predictive models, and conduct comparative analyses in a fraction of the time a human team might require. For instance, while a human analyst might take several hours to model three alternative scheduling strategies, an AI system can simulate dozens of possibilities in seconds, ranking them according to projected impact on project objectives.

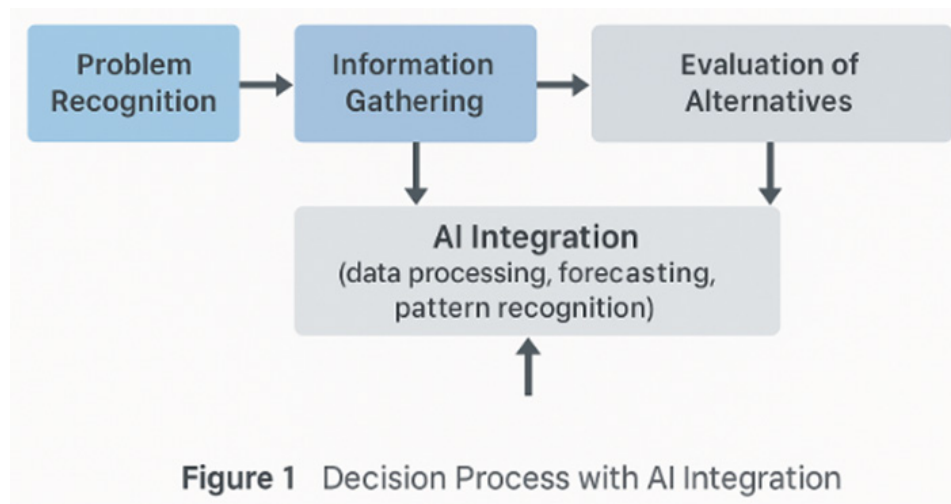
A distinguishing capability of AI at this stage is its integration of pattern recognition. Drawing from historical project data, industry trends, and organizational performance patterns, AI can detect correlations or recurring sequences that human decision-makers might overlook. For example, it might identify that projects with certain supplier dependencies historically encounter a 15% cost overrun unless a specific margin adjustment is made to procurement timelines.

The final stage involves recommending solutions and providing guidance on implementation. Here, AI not only proposes a preferred solution but also outlines a roadmap for execution. This can include step-by-step task sequences, resource allocation adjustments, risk mitigation strategies, and contingency plans. Importantly, AI-generated recommendations may challenge established practices – introducing unconventional approaches derived from data-driven insights rather than tradition or precedent. For example, an AI tool might suggest sequencing a testing phase earlier than usual based on historical evidence that earlier testing reduces late-stage defect costs by 30%. While such proposals can push teams outside their operational comfort zones, they often uncover opportunities for efficiency and innovation.



Throughout the process, human oversight remains essential. Even when AI demonstrates high accuracy in pattern detection and scenario modeling, final approval should rest with the project manager or designated decision authority. This ensures that proposed actions align with organizational culture, stakeholder expectations, and ethical considerations. AI is a powerful decision support partner, but it is most effective when paired with human judgment that can interpret nuances, manage trade-offs, and account for non-quantifiable factors.

As the AI tool suggests a solution, it will also discuss the process for implementing that solution and provide ideas to the project manager that might differ from what is typically done.



## Types of Decisions: Structured, Semi-Structured, and Unstructured

Decisions within organizations can be understood along a continuum ranging from highly structured to entirely unstructured. Each category has distinct characteristics, information requirements, and decision-making approaches, all of which influence how AI can support or enhance the process.

### *Structured Decisions*

Structured decisions are repetitive, rules-based, and governed by well-established protocols. They typically involve predictable inputs, a clear set of criteria, and a straightforward path to a correct or optimal outcome. Examples in project management might include approving recurring expense reports, processing standard procurement requests, or generating monthly performance dashboards. Gorry and Morton (1971) identified these as ideal candidates for algorithmic support because the rules for arriving at a decision can be explicitly programmed. In such cases, AI (particularly narrow AI or rule-based systems) can fully automate the process, thereby reducing the manual workload and eliminating the risk of human error associated with repetitive tasks.

### *Unstructured Decisions*

Unstructured decisions, on the other end of the spectrum, are complex, novel, and often high-stakes. They lack a clear decision rule, require judgment under uncertainty, and draw on experience, creativity, and sometimes intuition. Strategic planning, managing a reputational crisis, or deciding whether to pivot a project's scope in response to disruptive market changes all fall into this category. Mintzberg, Raisinghani, and Théorêt (1976) observed that unstructured decisions rely heavily on human interpretation of ambiguous or incomplete data. In such contexts, AI is best positioned as a

decision-support tool rather than a decision-maker – providing data analysis, pattern detection, and scenario modeling that inform, but do not replace, the nuanced reasoning of experienced leaders.

## ***Semi-structured Decisions***

Semi-structured decisions occupy the middle ground. They have identifiable components that can be structured, alongside elements that still require human judgment. For example, determining whether to accept a scope change request in a project involves structured steps, such as calculating the projected impact on budget and schedule, combined with qualitative considerations, including stakeholder relationships, contractual flexibility, and long-term strategic alignment. Kankanhalli, Zuiderwijk, and Tayi (2017) noted that these types of decisions are becoming increasingly prevalent in knowledge-driven sectors, where the volume of available data is high, but the context still demands human interpretation.

AI's role varies significantly across these categories. In structured contexts, automation can accelerate decision-making, minimize human error, and free project managers to focus on higher-value activities. In semi-structured contexts, AI can process large datasets, forecast likely outcomes, and present ranked options for review, enabling managers to focus on weighing trade-offs and contextual factors. In unstructured contexts, AI's analytical capabilities can still be valuable (providing risk models, trend analyses, and predictive insights). However, the final decision must account for human, cultural, and political factors that no algorithm can fully anticipate.

The key consideration is alignment between the decision type and the AI tool capability. Deploying an AI solution without first mapping its characteristics can result in a mismatch. For instance, using a fully automated AI system for a decision that involves high political sensitivity can lead to outcomes that are technically sound but strategically damaging. Conversely, relying solely on human judgment for a structured, data-heavy decision can waste time and resources that AI could have optimized. Understanding where a decision falls on the structured-unstructured spectrum is, therefore, a prerequisite for effectively integrating AI into organizational processes.

## **Challenges in Modern Decision Environments**

Decision-makers now operate under conditions characterized by volatility, uncertainty, complexity, and ambiguity. Eppler and Mengis (2004) noted that the abundance of information can overwhelm managers who struggle to isolate the most relevant insights from a continuous stream of data. Kahneman et al. (2021) further emphasized that compressed decision timelines increase the likelihood of errors rooted in cognitive biases. Bias continues to present a substantial barrier to effective decision-making. Bazerman and Moore (2013) and Milkman et al. (2009) document that both individual and group decision processes are susceptible to cognitive distortions, including confirmation bias, anchoring, and groupthink.

Even experienced managers may adopt suboptimal strategies under pressure, particularly in fast-moving environments. These findings collectively suggest a need for tools that can filter information, counteract bias, and facilitate timely judgment. AI, when implemented with a clear understanding of its limitations, offers a potential means of alleviating these decision-making challenges.

## **Mapping the Decision Process: Where AI Can Contribute**

Decision processes generally involve problem recognition, information collection, evaluation of alternatives, and the implementation of chosen actions. Shollo et al. (2020) found that AI tools, particularly those utilizing machine learning and natural language processing, enhance performance in data-intensive phases, such as information gathering and early analysis. For instance, AI can consolidate and analyze large data sets to support rapid risk assessment or detect emerging patterns that humans may overlook.

Jarrahi (2018) and Dellermann et al. (2019) converged on the idea that AI functions best as a complement rather than a replacement for human decision-making. Hybrid models, in which algorithms handle pattern recognition and humans apply judgment and ethical reasoning, often produce superior outcomes. This comparison highlights the importance of mapping decision stages deliberately, ensuring that AI is applied where its speed, scalability, and data-processing capabilities offer meaningful advantages, while preserving essential human judgment.

## When to Use a *Project Manager's Decision Guide*

The *Project Manager's Decision Guide* is a lightweight, repeatable guide that shows how to use AI for specific help in making project decisions. The guide assists project managers in determining whether to apply pre-developed prompts and suggestions or to develop a customized prompt for decisions that extend beyond the scope of the *Project Manager's Decision Guide*.

This guide examines when, how, and whether to use AI, based on the complexity of the decision, the repeatability of the decision (i.e., how often the same decision recurs), and data relevance and availability, including the volume, quality, and timeliness of the data.

The guide classifies each decision and highlights AI involvement, including informing, recommending, or taking action. All AI-made decisions must include human oversight to verify that it is focused on the correct topic, decision, and suggestions.

### ***Inform***

Using AI to inform others can help reduce knowledge gaps within the project team when implementing new technology. As artificial intelligence becomes more mainstream, organizations will likely have gaps in their knowledge and may need a bridge to fill those areas.

The ability for AI to inform can also include an automatic focus that allows it to track various project metrics, and the system creates a push notification, enabling everyone to stay informed about the project's progress.

Enabling AI to track specific data and trigger alerts enables leadership to remain informed about project developments in real-time. This timely information provides additional opportunities to make informed decisions or take corrective action before potential harm affects the plan or team. An example would be that the budget exceeds a preset threshold of 10%, triggering a response to the project manager so they can review the expenditure more closely.

### ***Recommend***

The ability for AI to conduct analysis in seconds and recommend solutions is a wealth of knowledge at the project manager's fingertips. AI scans various knowledge bases and creates a list of solutions, ranking recommendations based on which solution is potentially the best response, and which is not. After reviewing multiple options, the project manager or team can then ask AI to recommend approaches to implement the solution, identify the risks associated with the different solutions, and suggest migration or contingency strategies.

An example of this is when AI recommends a change to the plan or a shift in resources to expedite the project and mitigate potential risk. Rapid analysis provides leadership with actionable options at the earliest possible stage. Typically, leaders want to know about a problem as early as possible so they can build an action plan to mitigate or prevent any potential harm.

## Act

The last connection to the *Project Manager's Decision Guide* focuses on taking action automatically or behind the scenes. It is the ability to set permissions that allow AI to carry out specific-level decisions under strict guardrails set by the organization, or the project manager. Automatic decisions are monitored through audits, dashboards, and exception reporting, which ensures that human oversight remains fully in control of the outcomes. This does not imply that AI assumes full decision-making authority; instead, it engages in low-risk, predefined actions within established parameters.

Some examples of automation using AI include adjusting a resource calendar when team members input information, such as time off for a vacation, which triggers an immediate response from the resource to the project manager, rather than waiting until the next meeting and hearing about it from the team member. Additionally, re-prioritizing backlog items based on updated customer metrics is becoming increasingly common. Product owners are coming into work, and the product backlog indicates a need for refinement based on customer feedback. The feedback triggers the instant refinement of the backlog.

Acting is only appropriate if the organization feels comfortable with AI capabilities and trusts the decisions it makes in specific areas. Each decision should build trust. Although monitoring can be intensive at first, it becomes less so as trust increases.

## Types of AI Relevant to Project Work

One of the primary benefits of utilizing AI in project management is its ability to support informed decision-making. In this section, the various types of AI will be examined, along with the kinds of work AI is expected to perform.

### Narrow AI

**Definition:** Narrow AI, also known as Weak AI, focuses on systems engineered to perform a single, specific task with high efficiency and accuracy. This is considered the workhorse side of AI.

**Positives:** In project management, Narrow AI is utilized for routine, well-defined tasks, such as chatbots answering stakeholder queries, recommendation engines optimizing resource assignments, or document classifiers organizing project files. This tool focuses on these main jobs.

**Negatives:** Narrow AI does have some constraints and limitations. This tool system can generalize a topic that requires specifics and may produce incorrect solutions when working with unfamiliar scenarios. Because Narrow AI is trained with a narrow focus, it excels in its designated area but is limited when it needs to be more versatile.

### Predictive AI

**Definition:** Predictive AI focuses on systems and outcomes that analyze historical and current data, as well as using statistical machine learning models. This AI focuses on making predictions and providing probabilities to help predict trends and future events.

**Positives:** The value of Predictive AI in analyzing scheduling risks and focusing on stakeholder responses is powerful. This tool analyzes past project performance, resource usage patterns, and timelines to anticipate potential bottlenecks or overruns. Putting this information in the hands of the project manager or development team is effective in enabling earlier decisions and solutions. Predictive AI helps not only with risk and project performance but also with cost estimating and stakeholder behavior modeling.

**Negatives:** All the positives associated with Predictive AI focus heavily on the consistency and quality of the historical data, which must be of the highest quality. It is essential to acknowledge that AI is not a crystal ball; its recommendations and predictions are not definitive. A project manager must still verify whether the projections are accurate. One can treat production as guidance rather than a set of deterministic answers. Overconfidence in such models without human judgment and context can lead to misguided project decisions.

## Generative AI

**Definition:** Generative AI, which excites everyone in project management as a potential source of data patterns, can produce new content drafts such as texts, images, code, or video based on prompts. It leverages generative models, such as large language models. Using ChatGPT sparks ideas about how they will impact the project environment.

**Positive:** The ability to utilize Generative AI to create reports, summarize meeting notes, generate decision options, develop various solutions, and facilitate communication with stakeholders is enough to excite anyone. The ease of use of this tool, even after minimal training, demonstrates that individuals across project types, educational backgrounds, and operational constraints can adopt it as a functional project assistant.

**Negatives:** Please make no mistake, Generative AI is not a creative being; it is a tool that creates narratives from pattern tendencies and refines what words typically follow other words. These abilities make it seem like this is new, but the tool follows a pattern or algorithm when writing. It is for this reason that AI Detectors can pick up what is written by AI and not by a human. Because AI is not creating anything new, it requires the user to edit and refine all documents, solutions, and patterns for factual and contextual accuracy before using or recommending them.

Generative AI has raised significant concerns across all areas of business, including project management. Many believe that AI will make decisions, eliminating the need for a project manager to oversee the project. Since AI does not make decisions on its own, it requires a human to review the information and make decisions based on the culture or specific situation.

## Cognitive AI (Adaptive AI for Learning Systems)

**Definition:** Cognitive AI (or Adaptive AI) represents one of the most advanced forms of artificial intelligence in project management. Unlike Narrow or Predictive AI, which operate within fixed parameters or depend heavily on historical datasets, Cognitive AI is designed to continuously learn from its interactions, adapt to changing conditions, and improve decision-making over time. This learning occurs through iterative exposure to both successful and unsuccessful outcomes, allowing the system to refine its internal models and adjust future recommendations.

At its core, Cognitive AI attempts to mirror the adaptability of human decision-making. Just as a project manager develops better instincts through experience, Cognitive AI builds its “experience” by analyzing patterns in historical project data, monitoring the outcomes of its own recommendations, and incorporating feedback from human users. Over time, the system becomes more adept at anticipating bottlenecks, identifying risks before they escalate, and recommending proactive countermeasures.

**Positives:** The most significant advantage of Cognitive AI in project management is its self-improving nature. Once deployed, it requires far less manual rule-setting or retraining than static AI models. For example, suppose a project’s resource allocation strategy consistently results in schedule delays. In that case, Cognitive AI can detect that pattern, test alternative allocation models, and recommend improved approaches in similar future scenarios, without the need for extensive reprogramming.

This adaptability is particularly valuable in complex, fast-moving environments where conditions shift rapidly. Cognitive AI can integrate new data streams (live supplier updates, evolving customer requirements, or emerging market trends) into its decision-making framework almost in real time. In doing so, it not only responds to current challenges but also anticipates potential disruptions, giving project managers a head start in mitigation planning.



As its knowledge base expands, Cognitive AI can also offer context-aware prioritization. For example, it might recognize that during regulatory compliance projects, meeting fixed deadlines is more critical than minimizing costs, and, therefore, recommend schedule-preserving actions even if they increase short-term expenses. This ability to weigh multiple objectives dynamically is what makes Cognitive AI more strategic compared to other AI types.

**Negatives:** Cognitive AI is not without challenges. One of the most persistent concerns is bias reinforcement. Because it learns from historical patterns, it can inadvertently amplify flawed decision-making tendencies present in its training data. For example, if cost-cutting measures have historically been favored over schedule adherence – even in situations where meeting deadlines was critical – the AI may continue to prioritize cost savings, potentially causing regulatory breaches or reputational harm. Without human oversight, this bias can remain undetected until significant damage occurs.

There is also the broader job displacement concern. As Cognitive AI grows more capable, some responsibilities that once required human expertise may shift to automated processes. While this can improve efficiency, it also raises organizational and ethical questions about the role of project managers and the need to maintain a human touch in leadership, stakeholder engagement, and culture-sensitive decision-making.

From a technical standpoint, Cognitive AI demands substantial computing resources. Deep analysis and real-time learning can require significantly more processing power than simpler AI systems. Users who have engaged with advanced “deep research” modes in tools like ChatGPT have experienced this firsthand, sometimes waiting 10–15 minutes for a complex output instead of the near-instant responses typical of lighter queries. In large organizations with multiple concurrent AI workloads, this can create performance bottlenecks if infrastructure is not scaled accordingly.

## The Non-Negotiable Role of Human Oversight

Given these risks, human oversight is essential – not just for final approvals but for guiding the AI’s learning trajectory. Project managers must monitor which recommendations are adopted, which are rejected, and why. Inputting this feedback into the system ensures that its evolving decision framework aligns with both the organization’s strategic priorities and cultural norms.

In practice, this means setting guardrails around what Cognitive AI can decide autonomously, defining clear escalation points for human review, and periodically auditing the AI’s decision history for unintended patterns. For example, an organization might permit the AI to adjust task sequencing in a project schedule automatically but require human sign-off before reallocating the budget between major workstreams.

When deployed thoughtfully, with a balance of adaptive capability, transparency, and human oversight, Cognitive AI can serve as a long-term strategic partner in project management. It offers the promise of not just faster decisions, but smarter ones. Decisions that grow more precise and contextually relevant with every project they touch.

# 10 Advanced Prompts for Running Traditional/Waterfall Projects

## 1. Requirements Prioritization

“Analyze the following project requirements list and assign a priority score to each based on alignment with objectives, dependencies, estimated effort, and predicted value delivery. Provide a ranked list and justify the top five priorities with evidence-based rationale.”

## 2. Timeline Estimation

“Given the project work breakdown structure (WBS), historical task durations, and resource availability data, generate a risk-adjusted project timeline. Identify potential bottlenecks and propose a schedule with milestones, incorporating best-case, worst-case, and most-likely completion scenarios.”



### 3. Resource Assignment

“Review the provided resource skill matrix, past performance data, and workload availability. Assign each task in the WBS to the most suitable resource while avoiding overutilization. Recommend alternative assignments where gaps in skills or capacity are identified.”

### 4. Risk Forecasting

“Analyze the project plan, status reports, and dependency maps to identify potential risks that could lead to delays, cost overruns, or quality issues. Prioritize risks based on likelihood and impact, and recommend mitigation strategies for the top five.”

### 5. Scope Change Impact Modeling

“Evaluate the impact of the proposed change request on the project’s cost, schedule, and resource plan. Model at least two implementation paths and quantify trade-offs for each. Recommend whether to approve, defer, or reject the change.”

### 6. Budget Performance Prediction

“Analyze project cost reports, committed expenditures, and procurement timelines to forecast the end-of-project spending. Identify any budget variances and recommend corrective actions to stay within financial targets.”

### 7. Critical Path Monitoring

“Review the project schedule and recalculate the critical path based on updated progress data. Highlight any tasks that have become time-sensitive due to slippage and recommend actions to protect the project’s completion date.”

### 8. Quality Risk Prediction

“Based on the planned sequence of tasks and historical defect patterns, predict which deliverables are at the highest risk of failing quality checks. Recommend adjustments to the quality assurance plan to address these risks.”

### 9. Documentation Completeness Checks

“Review all project documentation for compliance with required standards. Identify missing sections, inconsistencies, or misaligned references, and provide a checklist of corrections needed before the next stage gate.”

### 10. Stakeholder Communication Optimization

“Analyze stakeholder engagement records, decision-making timelines, and feedback frequency. Recommend the optimal communication format, frequency, and content strategy for each stakeholder group to maintain alignment and accelerate approvals.”

## 10 Advanced Prompts for Running Agile/Scrum Projects

### 1. Backlog Item Prioritization

“Analyze the backlog items and rank them based on predicted value delivery, implementation effort, and historical sprint outcomes. Provide a prioritized list and justify the selection of the top five items based on projected impact for the upcoming sprint.”

## 2. Sprint Capacity Forecasting

“Using past velocity data, team availability, and story complexity ratings, forecast the maximum number of backlog items that can be completed in the upcoming sprint. Recommend which items to commit to and which to defer.”

## 3. Task Assignment Optimization

“Review the sprint backlog and assign tasks to team members based on skills, past completion rates, and current workload. Identify any potential skill gaps or workload imbalances and recommend targeted resolutions.”

## 4. Risk Trend Detection

“Analyze sprint progress data, daily stand-up notes, and impediment logs to detect emerging risks. Summarize the three most significant emerging risk trends, analyze their potential impact on sprint objectives, and propose evidence-based mitigation actions.”

## 5. Scope Change Effect in Mid-Sprint

“Evaluate the impact of adding the proposed backlog item to the current sprint. Model how it will affect the completion probability for committed work and recommend whether to include, swap, or defer the item.”

## 6. Predictive Burndown Tracking

“Review real-time progress data and historical completion rates to predict if the sprint burndown will reach zero on time. Identify pacing issues and recommend adjustments to ensure sprint completion.”

## 7. Defect Probability for User Stories

“Analyze user story complexity, past defect density, and test coverage metrics to predict which stories have the highest defect probability post-release. Recommend additional QA steps for these stories.”

## 8. Continuous Improvement Insights

“Aggregate performance metrics from the last five sprints, including cycle time, lead time, and spillover rates. Identify recurring inefficiencies and recommend process improvements for the next retrospective.”

## 9. Stakeholder Feedback Integration

“Review stakeholder feedback from demos and surveys, extract recurring themes, and rank them by potential product impact. Recommend backlog adjustments to address the most valuable feedback first.”

## 10. Cross-Team Coordination

“Analyze dependencies between multiple Agile teams and forecast where delays in one team could block another. Recommend coordination strategies and timing adjustments to avoid bottlenecks.”

# Ethical and Practical Limits of Current AI Capabilities

## Definition

The ethical and practical limits of current artificial intelligence capabilities refer to the boundaries that define how AI can be responsibly developed, deployed, and managed in real-world contexts. Ethically, these limits are shaped by principles such as human well-being, fairness, privacy, transparency, accountability, and safety (Cebulla et al., 2023; Fischer & Frennert, 2025; Knight, 2025). These principles establish a moral framework for guiding the use of AI; however, they are neither static nor universally interpreted.

Their meaning often shifts according to context, stakeholder perspectives, and cultural norms, which can create challenges in translating high-level ideals into operational practices. As a result, ethical limits are not only determined by the formal codes and guidelines that govern AI but also by the ability of those guidelines to adapt to lived experiences, emerging risks, and unforeseen consequences.

Practically, these limits encompass the technical, organizational, and social constraints that affect the safe and effective operation of AI. They include issues such as the persistence of workplace hazards even after automation (Cebulla et al.), the unpredictability of human-AI interactions (Fischer & Frennert), and the gaps in evidence-based guidance for ethical governance (Knight). Practical limits acknowledge that AI systems function within complex environments shaped by the intersection of technical performance, human behavior, and institutional processes.

These boundaries highlight the need for ongoing monitoring, iterative risk assessment, and adaptive governance to ensure that AI not only functions as intended but also aligns with societal values and organizational responsibilities.

## Introduction

Artificial intelligence is becoming increasingly integrated into workplaces, public services, and everyday life. It is used to automate decision-making, manage workforces, provide predictive analysis, and operate robotics systems. These developments have created significant opportunities for increased efficiency, enhanced safety in specific environments, and innovation in products and services. However, the adoption of AI also presents ethical and practical challenges that affect individuals, organizations, and society.

This discussion draws on three complementary perspectives. Cebulla, Szpak, and Knight examine workplace health and safety (WHS) risks associated with AI adoption, emphasizing the need for systematic risk assessment at each stage of implementation. Fischer and Frennert review empirical studies on how people experience AI and robots, proposing an “experiential ethics” approach that grounds ethical reflection in lived encounters. Knight examines the development of AI ethics guidelines, highlighting the necessity for evidence-based, context-specific, and adaptable processes. Together, these works highlight that ethical and practical limits are not solely theoretical concerns but emerge from the intersection of governance structures, workplace realities, and human experiences.

## Ethical Challenges in Principle and Practice

### Common Ethical Foundations

Across the three studies, there is strong agreement on core ethical principles for AI governance. These include human well-being, fairness, privacy protection, transparency, accountability, and safety. Cebulla et al. reference the Australian Government’s AI Ethics Principles, condensing them into three overarching categories: well-being, safety, and accountability. Fischer and Frennert identify similar values as recurring themes in both empirical research and policy

discussions. At the same time, Knight notes that such values often dominate AI ethics guidelines, which are frequently adapted from established frameworks, such as biomedical ethics and the Belmont Principles.

Although these principles provide an essential ethical foundation, they frequently remain abstract and challenging to apply in specific contexts. Fischer and Frennert note that the meaning of principles such as transparency or fairness changes depending on the context in which AI is used. Interpretations may vary among developers, managers, and end users. Cebulla et al. show that even when principles are well-articulated, they can fail to address nuanced psychosocial risks, such as changes in workload distribution or the erosion of peer relationships. Knight further highlights that without clear strategies for implementation, principles risk serving as symbolic statements rather than operational guidance.

## Limits of Abstract Principles

The inherent flexibility of interpretation in many ethical terms represents both a strength and a limitation. For example, “transparency” might mean algorithmic explainability for a software developer, but for an employee, it may refer to explicit notification that AI is being used for monitoring performance. Fischer and Frennert emphasize that the full implications of such principles often emerge only after deployment, when AI systems are fully integrated into daily practice.

Cebulla et al. demonstrate that ideals such as fairness can fail to capture the localized impacts of AI, including the unequal distribution of tasks or disparities in who bears the brunt of increased cognitive demands. Knight notes that many guidelines fail to define these principles with sufficient precision to make them operational across diverse settings. Such ambiguity can lead to inconsistent application and erode trust in AI governance frameworks.

## Practical Limits in Workplace and Societal Contexts

### AI in the Workplace: Safety and Well-Being Risks

The research by Cebulla et al. identifies several hazards associated with the adoption of AI in the workplace. These include increased cognitive load, diminished job control, privacy concerns, and overreliance on automated decision-making. Although AI can improve safety by removing humans from physically dangerous tasks, it can also create new forms of harm, particularly to mental health and social cohesion in the workplace. For example,

AI-based monitoring systems can create an environment of constant surveillance, which can lead to stress and a decline in trust.

These findings suggest that existing workplace health and safety (WHS) frameworks, which traditionally emphasize physical risks, may insufficiently address cognitive, emotional, and relational harms. If these frameworks are not updated to account for AI-related risks, organizations may fail to safeguard employees from subtle yet equally significant consequences of technological change.

### Experiential Limits: Human-AI Interaction

Fischer and Frennert identify six recurring dimensions of human experiences with AI and robots: appreciation of imperfection, formation of emotional connections, discomfort with lack of transparency, invisible human work, shifting responsibilities, and willingness to trade privacy for benefits. These dimensions demonstrate that the real-world impact of AI extends beyond technical functionality, influencing perception, trust, and workplace culture.

For instance, invisible labor (the human effort required to maintain or train AI systems) challenges the common assumption that automation entirely removes human involvement. Similarly, shifting responsibilities from humans to machines can blur accountability, particularly in high-stakes contexts. These insights align with Cebulla et al.’s emphasis on the persistence of risks over the AI system’s life cycle, and they reinforce Knight’s call for adaptive governance that responds to changing circumstances.



## Limits in Governance and Guideline Development

Knight reveals that many AI ethics guidelines lack clear definitions of their intended audience, fail to adapt principles to specific contexts, and do not provide mechanisms for ongoing evaluation. While most guidelines state high-level values, they often fail to connect these values to practical actions supported by evidence. This governance gap reflects a prevailing tendency to treat ethics as a static set of rules rather than a dynamic process responsive to technological and social change.

Knight argues for integrating procedural guidance, learning objectives, and adaptive governance structures into ethics guidelines. This aligns with Fischer and Frennert's emphasis on incorporating feedback from experiences and with Cebulla et al.'s recommendation for continuous risk monitoring.

## Towards an Integrated Understanding of AI's Limits

### Stage-Based Risk Identification

Cebulla et al. recommend integrating ethical principles with the AI Canvas framework to assess potential risks at each stage of AI implementation, from the initial concept through design, testing, and operational use. This approach ensures that both physical and psychosocial hazards are identified early and monitored over time. Embedding ethical considerations into each stage enables organizations to anticipate and prevent harm more effectively.

However, this process must be flexible enough to respond to unforeseen challenges. Fischer and Frennert demonstrate that many ethical issues arise only during the real-world use of AI systems. In contrast, Knight emphasizes the importance of iterative review and stakeholder participation throughout the AI system's lifecycle. Integrating these perspectives fosters a more robust and adaptive risk management framework.

### Experiential Feedback Loops

Fischer and Frennert's experiential ethics framework emphasizes the importance of systematically gathering and analyzing user experiences to detect emerging risks. This process can reveal concerns such as discomfort with system opacity or changes in responsibility allocation before they escalate.

Cebulla et al. note that risks can persist or reappear throughout the life of an AI system, making continuous monitoring essential. Knight suggests that guidelines should incorporate formal mechanisms for integrating such feedback, ensuring that ethical oversight adapts to evolving circumstances and user experiences.

### Evidence-Based Guideline Design

Knight's PRISMA-ETHICS framework calls for grounding AI ethics guidelines in empirical evidence. This involves drawing on case studies, research findings, and stakeholder input to create substantive, procedural, and educational recommendations. Evidence-based design enhances the credibility and usability of guidelines across sectors.

Fischer and Frennert's focus on contextual, lived experiences, and Cebulla et al.'s structured risk assessment model complements this approach. Together, they support the development of guidelines that are adaptable, inclusive, and capable of guiding ethical decision-making in complex, changing environments.

## Conclusion

The interplay between abstract principles, workplace realities, and human experiences shapes the ethical and practical limits of AI. Cebulla et al. reveal that AI can introduce new hazards even as it mitigates traditional risks. Fischer and Frennert demonstrate that the ethical significance of AI emerges in lived encounters, often in ways not anticipated by policy frameworks. Knight shows that without evidence-based, participatory, and adaptive guideline development, ethical governance risks remain aspirational rather than actionable.

An integrated approach that incorporates stage-based risk identification, experiential feedback loops, and evidence-based guideline design represents the most effective path forward. Such an approach ensures that AI governance is both principled and responsive, safeguarding human well-being while enabling responsible innovation.

## References

- Alon-Barkat, S., & Busuioc, M. (2022). Human–AI interactions in public sector decision making: “Automation bias” and “selective adherence” to algorithmic advice. *Journal of Public Administration Research and Theory*, 33(1), 153–169. <https://doi.org/10.1093/jopart/muac007>
- Bazerman, M. H., & Moore, D. A. (2013). *Judgment in managerial decision making* (8th ed.). John Wiley & Sons.
- Brynjolfsson, E., & McElheran, K. (2016). The rapid adoption of data-driven decision-making. *American Economic Review*, 106(5), 133–139. <https://doi.org/10.1257/aer.p20161016>
- Cebulla, A., Szpak, Z., & Knight, G. (2023). Preparing to work with artificial intelligence: Assessing WHS when using AI in the workplace. *International Journal of Workplace Health Management*, 16(4), 294–312. <https://doi.org/10.1108/IJWHM-09-2022-0141>
- Dellermann, D., Ebel, P., Söllner, M., & Leimeister, J. M. (2019). Hybrid intelligence. *Business & Information Systems Engineering*, 61(5), 637–643. <https://doi.org/10.1007/s12599-019-00595-2>
- Elbanna, S. (2006). Strategic decision-making: Process perspectives. *International Journal of Management Reviews*, 8(1), 1–20.
- Elbanna, S., & Child, J. (2007). The influence of decision, environmental and firm characteristics on the rationality of strategic decision-making. *Journal of Management Studies*, 44(4), 561–591. <https://doi.org/10.1111/j.1467-6486.2006.00670.x>
- Eppler, M. J., & Mengis, J. (2004). The concept of information overload: A review of literature from Organization Science, Accounting, Marketing, MIS, and related disciplines. *The Information Society*, 20(5), 325–344. <https://doi.org/10.1080/01972240490507974>
- Fischer, B., & Frennert, S. (2025). Towards an experiential ethics of AI and robots: A review of empirical research on human encounters. *Technological Forecasting & Social Change*, 219, 124264. <https://doi.org/10.1016/j.techfore.2025.124264>
- Gerlings, J., Shollo, A., & Constantiou, I. (2021). Reviewing the need for explainable artificial intelligence (xAI). In *Proceedings of the 54th Hawaii International Conference on System Sciences (HICSS)*. arXiv. <https://doi.org/10.48550/arXiv.2012.01007>
- Gomez, C., Cho, S. M., Ke, S., Huang, C.-M., & Unberath, M. (2024). Human–AI collaboration is not very collaborative yet: A taxonomy of interaction patterns in AI-assisted decision making from a systematic review. *Frontiers in Computer Science*, 6, Article 1521066. <https://doi.org/10.3389/fcomp.2024.1521066>

- Gorry, G. A., & Morton, M. S. (1971). A framework for management information systems. *Sloan Management Review*, 13(1), 55–70.
- Haugaard, M. (2021). The four dimensions of power: Conflict and democracy. *Journal of Political Power*, 14(3), 425–447. <https://doi.org/10.1080/2158379X.2021.1878411>
- Inkpen, K., Chappidi, S., Mallari, K., Nushi, B., Ramesh, D., Michelucci, P., Mandava, V., Vepřek, L. H., & Quinn, G. (2023). Advancing human–AI complementarity: The impact of user expertise and algorithmic tuning on joint decision making. *ACM Transactions on Computer–Human Interaction*, 30(5). <https://doi.org/10.1145/3534561>
- Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human–AI symbiosis in organizational decision making. *Business Horizons*, 61(4), 577–586. <https://doi.org/10.1016/j.bushor.2018.03.007>
- Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A flaw in human judgment*. Little, Brown Spark.
- Kankanhalli, A., Zuiderwijk, A., & Tayi, G. K. (2017). Open innovation in the public sector: A research agenda. *Government Information Quarterly*, 34(1), 84–89. <https://doi.org/10.1016/j.giq.2016.12.002>
- Knight, S. (2025). Understanding the use of evidence in AI ethics guidelines development through a PRISMA-ETH-ICS informed scoping review of guidelines. *Computers and Education Open*, 6, 100281. <https://doi.org/10.1016/j.caeo.2025.100281>
- Li, J., Yang, Y., Liao, Q. V., Zhang, J., & Lee, Y.-C. (2025). As confidence aligns: Exploring the effect of AI confidence on human self-confidence in human–AI decision making. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM. <https://doi.org/10.1145/3706598.3713336>
- Milkman, K. L., Chugh, D., & Bazerman, M. H. (2009). How can decision making be improved? *Perspectives on Psychological Science*, 4(4), 379–383. <https://doi.org/10.1111/j.1745-6924.2009.01142.x>
- Mintzberg, H., Raisinghani, D., & Théorêt, A. (1976). The structure of “unstructured” decision processes. *Administrative Science Quarterly*, 21(2), 246–275. <https://doi.org/10.2307/2392045>
- Simon, H. A. (1977). *The new science of management decision* (3rd rev. ed.). Prentice-Hall.
- Vasconcelos, H., Jörke, M., Grunde-McLaughlin, M., Gerstenberg, T., Bernstein, M., & Krishna, R. (2023). Explanations can reduce overreliance on AI systems during decision-making. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW) 2023*. ACM. <https://doi.org/10.1145/3579605>
- Zhang, S., Yin, C., Tory, M., Zhang, P., Yu, J., Lu, Y., Padilla, L. M., Wang, D., Xu, X., Yao, B., & Caterino, J. (2024, February 26). Rethinking human–AI collaboration in complex medical decision making: A case study in sepsis diagnosis (Version 2). arXiv. <https://doi.org/10.48550/arXiv.2309.12368>



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](http://www.themathisgroup.com)

keith@themathisgroup.com  
wanda@themathisgroup.com

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

