



Building organizational and project readiness for AI: A dialectical approach

Daniel Erian Armanios^{a,*}, Christopher L. Tucci^b

^a BT Professor and Chair of Major Programme Management, Saïd Business School, University of Oxford

^b Professor of Digital Strategy & Innovation, Imperial Business School, Imperial College London

ABSTRACT

To date, most organizational and project management studies have focused on how to deploy artificial intelligence (AI) tools to improve outcomes. However, this wrongly assumes the preexisting project and organizational context is already prepared to incorporate such AI tooling productively. To address this issue, this essay proposes a provocation that we call CIPHER, a first-of-a-kind framework to increase organizational readiness for deploying AI in project work. CIPHER identifies six key tensions that occur with the use of AI: Cognitive, Informational, Projection, Haptic, Exchange, and Resource. These tensions occur either when humans directly interface with algorithms or indirectly where humans interface with an algorithmically driven robotic, immersive reality, and/or multi-agent system. For each tension, we identify a key project lever to navigate these tensions: intergenerational teaming, problem scoping, delivery-validation scaffolding, task separability, multi-agent governance, and model parsimony & transfer learning. Our aim is that through CIPHER, project leaders can harness the promise of AI while also more carefully managing its perils.

1. Introduction

Artificial intelligence (AI) in project management is an emerging area of interest. Studies thus far focus on where project managers perceive AI is most useful (e.g., [Holzmann et al., 2022](#)), how to ethically conduct project management research with AI (e.g., [Müller et al., 2024](#)), and how to use AI in specific processes such as project planning (e.g., [Barcaui & Monat, 2023](#)). Beyond project management, organizational scholars are beginning to analyze how different AI-human interactions influence task performance ([Agarwal et al., 2023, 2018](#); [Dell'Acqua et al., 2025, 2023](#)), and how to redesign work with such interactions in mind ([Choudhary et al., 2025](#); [Iantisi et al., 2020](#); [Puranam, 2021](#)). Overall, extensive experimentation is underway around how to use AI ([Mollick, 2024](#)), and the prognosis of such efforts ranges from optimism (e.g., [Project Management Institute, 2024a, 2024b](#)), to caution (e.g., [Geraldi et al., 2024](#)).

These task-specific studies assume the preexisting organization and project is already prepared to effectively deploy AI. So much so, the prevailing preoccupation is how specific targeted AI deployments advance project outcomes. For instance, prior studies have explored how AI can improve project performance ([Costantino et al., 2015](#); [Wang et al., 2012](#); [Yang & Zhao, 2025](#)), risk management ([Diekmann, 1992](#); [Jin & Zhang, 2011](#); [Mancini et al., 2023](#)), and stakeholder identification ([Mariani et al., 2023](#); [Miller, 2022](#)). These studies recognize that organizational capabilities and agility around deploying AI are important

([Mariani & Mancini, 2025](#); [Mikalef & Gupta, 2021](#); [Whyte et al., 2016](#)). However, where these skills come from and how organizations and projects should build them to enhance their AI readiness, or what some call absorptive capacity ([Cohen & Levinthal, 1990](#)), is largely absent. Put another way, while frameworks for AI deployment do exist such as SALIENT in healthcare ([Van Der Vegt et al., 2023](#)), these frameworks are often “inside-out” (AI-leading) where they start with the algorithm and extrapolate out into the implementation setting. Our approach is “outside-in” (project and organization-leading) where we start with the organizational and project site setting and preparedness to guide the most appropriate AI deployments. We see our approach as especially important as current investigations around the outcomes from AI deployments are mixed at best ([Challapally et al., 2025](#); [Dell'Acqua et al., 2023](#); [Korst et al., 2025](#)). This suggests building project and organizational readiness for AI is not a trivial undertaking, and yet so far, we lack a guiding framework for how best to enhance such preparedness for deploying AI tools effectively.

In this essay, we propose a dialectical framework that we call CIPHER¹ to conceptually identify how projects and organizations enhance their readiness to deploy AI tools. In particular, this framework proposes six key tensions, and six key accompanying levers for effectively deploying AI amidst such tensions. This complements work that develops a set of organizing principles, which are derived from the distinct features of AI and their potential to reshape organizational fundamentals based on the behavioral theory of the firm ([Armanios et al., 2025b](#)).

* Corresponding author.

E-mail addresses: daniel.armanios@sbs.ox.ac.uk (D.E. Armanios), c.tucci@imperial.ac.uk (C.L. Tucci).

¹ No relation to and not to be confused with the widely used CFIR implementation framework adopted in healthcare and other settings ([Van Der Vegt et al., 2023](#))

Together, they advance a more novel and coherent view of how to more effectively deploy AI within projects and organizations.

2. A brief primer on AI

Before we can advance a framework for making AI projects relevant, we need to first have a common understanding around what we mean by the term. Simply put, AI is how machines learn, reason, and self-adapt akin to human cognitive functioning (Nature Research Intelligence, 2025a). This is done through machine learning, which is a set of techniques that help computational systems learn from data and improve their performance that do not necessitate explicit, bespoke programming (Nature Research Intelligence, 2025b).

These techniques are usually classified into supervised, unsupervised, and reinforcement learning. Supervised learning is prediction from labeled or tagged data. Regression models are one such classic example that predicts outcomes based on modelled relationships between explicitly called or labelled variables. Unsupervised learning is detection of groupings and sequences in unlabeled data in which the relationship between variables is unknown. Clustering techniques, such as k-means clustering, are one such classic example that seeks to increase the internal coherence of a cluster while maximizing the distance between clusters to identify distinguishable topics within a corpus of natural text (Sinaga & Yang, 2020). Finally, there is reinforcement learning, which is prediction done through agents in an environment. These agents make decisions that are updated based upon feedback in the form of rewards or penalties earned based on their observed performance on a given task. Q-learning is one such classic example in which agents try a particular action under a particular condition. The agents evaluate two factors: the consequences of their action based on received payoffs (rewards or penalties), as well as the benefits or risks of the condition in which the action was taken (Watkins & Dayan, 1992). These three forms of learning and their combinations are what inform much of AI to date.

Large language models (LLMs, e.g., GPT-5 by OpenAI) is a form of AI that is receiving much attention and combines all three forms of machine learning (e.g., Minaee et al., 2024). These algorithms are initialized through unsupervised natural language processing (NLP) techniques that translate natural language into numerical representations to help these models understand structures and patterns in natural language. The downstream models are then fine-tuned using supervised and reinforcement learning via prompts. For text output, past sequences become “labels” for which future words are predicted—such automatic labelling of preceding text is a form of self-supervised learning. If the previous output is “The girl likes to play in the...”, “The girl likes to play in the” is the label from which future word predictions are conditioned. Using reinforcement learning, the model is trained to respond to different kinds of prompt instructions and adaptively improve its content based on the user. Humans then rate the quality of the output, and the model uses this performance feedback to update its weights and probabilities.

3. CIPHER: A dialectical framework for informing project-based use of AI

Given the wide application and potential of these models, we argue for the need to spend less time deciding whether AI is good or bad and spend more time accepting that both are possible and enhancing project and organizational readiness to steer these tools toward promise instead of peril. To do this, we employ a classic philosophical approach known as the dialectical method. Dating back to Ancient Greece, the aim of this approach is to advance knowledge by synthesizing across seemingly opposing assertions (Busco & Quattrone, 2018). Given AI’s simultaneous benefits and pitfalls, this is a useful approach for helping us ascertain how to progress through these interlocked dualities.

Our approach then is to identify project levers that help uncover and

decode these tensions akin to a metaphorical cipher, hence the naming of our framework. More specifically, we take each dialectical tension and couple it with a key project principle. From this, a project lever emerges that can inform how project professionals can more tangibly and simultaneously navigate the promise and perils of AI in their work. The resulting CIPHER framework focuses on six tensions that prominently appear in the use of AI – Cognitive, Informational, Projection, Haptic, Exchange, and Resource. A summary of this framework is presented in Fig. 1.

Cognitive: Team Readiness

Dialectic #1: AI enhances creativity and cognitive offloading (reduced critical thinking)
Project Principle #1: Account for when to accept and deny AI outputs in projects

While project-based studies to date recognize the value of AI on teams (Barbosa & Carvalho, 2025; Dell’Acqua et al., 2025), how to improve team readiness to effectively onboard and use AI to begin with is understudied. This is crucial because while AI enhances creativity, users are shown to reduce their critical thinking around AI outputs. We can see this through studies that compare individual vs. systemwide creativity from AI. Individual creativity is enhanced because AI reveals novel output not previously considered or known to that user. However, because AI only recombines what already exists in its training corpus, overall creativity may actually decrease (Doshi & Hauser, 2024). This occurs because humans are often not contributing novel discrepant inputs to AI models. In fact, several studies argue that when humans use AI, they engage in “cognitive offloading,” whereby instead of critically interrogating the output, they accept its conclusions without scrutiny (Gerlich, 2025; Lee et al., 2025). Given the humanlike ability that AI exhibits to convince users of its veracity, these problems are only likely to increase.

The key tension here is a *cognitive* one. To navigate this tension, project managers need to situate AI in contexts where some are more adept at using these tools to identify creative solutions, while others are more equipped to critically interrogate its outputs. Project teams need to have the means to both accept *and* deny AI outputs. How to do so is hidden in plain sight. Of those studies that find greater cognitive offloading and therefore less critical thinking around AI, they find that older individuals accept AI output less than younger individuals (Gerlich, 2025).

The lever here then is to situate AI tools in *intergenerational teams* that can allow deploying AI both creatively and critically. This is sometimes referred to as balancing digital and legacy skills (Lanzolla et al., 2021). Older individuals on the team are often seasoned project professionals whose experience can help better interrogate the validity of the output. Younger individuals are often experienced in the latest tools and technologies and so they can contribute more creative deployment of AI tools toward project problems. This leads to the following testable proposition:

Proposition 1 (Intergenerational teaming): AI deployment in projects will have greater performance when deployed in intergenerational project teams that vary in age, industry, and AI experience.

Informational: Individual Readiness

Dialectic #2: AI informs and deceives
Project Principle #2: Account for how AI searches and shapes knowledge

A growing set of studies on project management education show that AI is enhancing learning (Ingason et al., 2025). However, other studies highlight the lack of clear recommended education and training practice, and this lack of policy hampers individual readiness to effectively deploy AI in projects (Smit et al., 2025; Vettori & Warm, 2025).

This is especially problematic because AI, such as LLMs, are “stochastic parrots:” they generate language without knowing the meaning of what they process and output; they simply predict novel content that

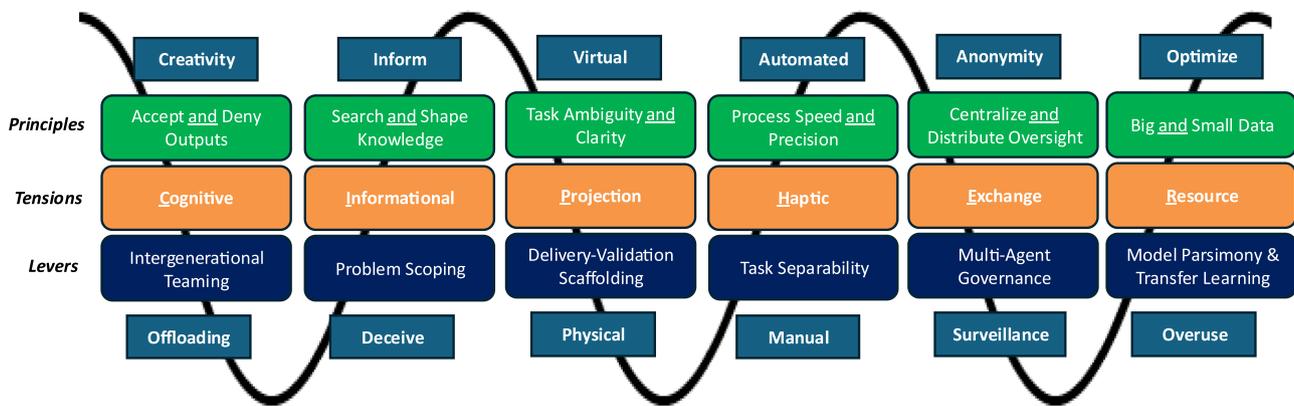


Fig. 1. CIPHER dialectical framework for deploying AI in projects.

is most likely to follow preexisting content based on its training data (Bender et al., 2021). This means while they can generate novel and potentially even creative output, they also “hallucinate” (i.e., provide realistic-sounding yet inaccurate information) and reflect the same biases of their training data and of their human labelers and developers (Buolamwini & Gebru, 2018). At best when humans mistakenly use such inaccuracies to inform organizational and project decision-making, AI can lead to *misinformation*, or unintentional processing errors / misunderstandings (Hannigan et al., 2024). However, some of these fabrications can rise to the level of *disinformation*, or purposive falsehoods designed to deceive their targets (Petratos, 2021). Given AI, such as GenAI, are designed to be humanlike in communications, distinguishing between useful vs. fabricated misinformation/disinformation is becoming increasingly difficult and consequential (Van Alstyne, 2023; Vanneste & Puranam, 2024).

The key project implication is that while most have been focused on how AI helps search for novel information, few have focused on how AI can shape the very information landscape through which we search (Gavetti et al., 2017). This is explored in arguably the most canonical organizational theory on this topic: the Behavioral Theory of the Firm (Cyert & March 1963). From this theoretical perspective, the fundamental starting point is that a problem emerges for which you are seeking a solution. At the onset, you use your existing knowledge, often stored in habit, standard operating procedures, and/or routines (Cyert & March 1963; Simon, 1997). If you do not have such knowledge, you first search proximately in more familiar “local” knowledge domains before more widely searching into more unfamiliar and costly “distant” domains until you find a “good enough” solution to the problem (Cyert & March 1963; March & Simon, 1958; Simon, 1955).

Scaling this up, the landscape (the entire parameter space through which these agents search) is considered more “rugged” as the dimensions along which one searches and the interdependencies between these dimensions increase (Gavetti & Levinthal, 2000; Levinthal, 1997). Performance improves either through improving search within a landscape or shaping the landscape itself by reframing problems more tractably to generate a “smoother,” less complex information landscape (Gavetti et al., 2017).

Let us now bring this back to the project space with a canonical problem for which AI has long been applied: project planning. For simplicity, let us think about project planning as simply the ordering of tasks needed to complete a project, along with their interdependencies. Studies considering AI in project planning have long argued that this is most effective when possible conditions, initial and goal differences, and possible actions are limited in scope (Levitt et al., 1988). In other words, the information landscape is made “smoother” and more optimizable. Naturally then, using AI for project planning becomes problematic when these tools are used amidst highly complex and interdependent tasks (a more “rugged” landscape).

The key tension here is an *informational* one. In complex information environments where conditions, goals, and actions are more unbounded, AI tools have more room to make inferences that are informed (and skewed) by their own training data. A key lever then to address this informational tension is *scoping the problem* into a tractable set of factors, ensuring a more cognitively manageable and knowable landscape (Armanios et al., 2025a). If the project planner does not scope the problem, AI tools will tacitly do it for them, which is increasingly troublesome with greater complexity. Exploratory studies of scheduling with AI tools demonstrate that they tend to underestimate the time needed compared to experienced managers (Barcaui & Monat, 2023). Recent studies show that such inaccuracies and underperformance are only likely to increase in projects within cultural contexts whose languages are outside LLM training data (Koo, 2025), which are largely confined to English, Mandarin, and Western European languages (Longpre et al., 2024). Thus, when using these AI tools, project planners must scope the problem in ways that ensure alignment with the information sources used to train the AI tool of interest.

Problem scoping as a lever necessitates deeper reflection on the means-end chain, or the number of unmeasured assumptions made between desired AI-enabled project tasks and their intended outcomes. With more unmeasured assumptions, greater discretion is allocated to the algorithm’s causal inferences, and the accuracy and appropriateness of those inferences depend on the algorithm’s training data. This leads to the following testable proposition:

Proposition 2 (Problem Scoping): *AI deployment in projects will have greater performance when there are (a) fewer unmeasured causal links between project tasks and goals and (b) greater alignment between the algorithm’s training data and the project’s operating context.*

Projection: *Translational (Across Time and Knowledge) Readiness*

Dialectic #3: AI requires virtual and physical reality.

Project Principle #3: Account for when to increase task ambiguity and clarity

While AI can increase the clarity of how to complete future tasks by simulating them virtually, this also can increase the ambiguity in how and which physical interventions are needed. This is particularly challenging with AI-driven immersive reality (xR). For instance while designers can leverage virtual reality to help enhance construction safety, this assumes designers are familiar with construction site challenges and/or have experienced construction professionals to help fill knowledge and implementation gaps when that is not the case (Sacks et al., 2015).

Within immersive reality (xR), let us explore the difference between augmented reality (AR) and virtual reality (VR). Though the line between the two is increasingly blurred, VR aims to immerse the user in an entirely separate and simulated environment such as those of Metaverse applications. On the other hand, AR often overlays a digital world onto

existing physical surroundings. In short, VR is anchored more in an artificial setting, while AR overlays on a physical setting.

Why is this distinction relevant? On the one hand, algorithmically driven VR can uncover new possibilities to one's work and can be updated based on user interaction and input. Interacting with algorithmically informed avatars presents an emerging example of this form of learning and interaction (Hu et al., 2025; Park et al., 2023). VR increases the possibility of thinking differently around a task, but how to translate that into physical action becomes unclear. In short, VR increases task ambiguity and reduces task clarity, except arguably in instances where the virtual environment closely mirrors a common realistic scenario that does not need physical anchoring (Wu et al., 2019).

On the other hand, algorithmically driven AR helps you see more clearly how to do a specific task as the setting is mapped directly onto a physical environment. This is especially useful in addressing project work where 2D drawings must translate into 3D actions (Chalhoub & Ayer, 2018). For example, designer Adam Pickard converted paper IKEA manuals into an AR app (AssembleAR) precisely to address this issue (Morby, 2018). The savings from such a conversion on productivity are significant. The company Upskill, using the Skylight AR platform, worked with GE Renewable Energy (now GE Vernova) to analogously overlay instructions onto wind turbines and found a 34 % reduction in assembly time (Kloberdanz, 2017; Upskill, 2017). However while this improves task execution, AR has a constrained field of view, so the task must be carefully narrowed to fit within this field of view (Ren et al., 2016). So while AR increases task clarity, it reduces task ambiguity as it reduces the ability to see creative alternatives within such a narrow aperture.

The tension here is a *projection* one. Project managers must juggle two forms of projection (i.e., virtual and physical) simultaneously and do so across different time and knowledge gaps. The key lever here is *delivery-validation scaffolding*, or the tools that help link delivery (i.e., doing the work) to assurance (i.e., checking the work).

Traversing time gaps, immersive reality can record the exact actions of workers, so delivery and assurance can occur simultaneously. This is a significant breakthrough given the usual way in which assurance is done via a "scrutiny loop." When checking work, project personnel need to spend additional time answering follow-up requests to recall what occurred while also continuing the next work task, often leading to increased delays (Maylor et al., 2024). With actions recorded, the precise actions project personnel take when doing the work can be shared directly to those checking the work, helping break this cycle. Lockheed used this when deploying AR as part of the NASA Orion project (Winick, 2018). Besides reducing errors to zero and reducing labor by 90% (Microsoft, 2021), the use of immersive reality allowed for worker movements to be recorded as the work was delivered, allowing for simultaneous process assurance. In this way, immersive reality serves as useful scaffolding between delivery and assurance that reduces time gaps between the two.

Traversing knowledge gaps, assurance can nonetheless suffer information asymmetries between those doing the work and those checking the work. Simply "seeing" the work may not overcome such knowledge divides. The closest scaffolding to date for bridging this divide around commonly used algorithmic tools are what is known as "model cards" that seek to communicate key factors informing the deployment of algorithms (Crisan et al., 2022; Mitchell et al., 2019). The challenge with such scaffolding is they report variables that are useful for developers but not necessarily comprehensible to other workers seeking to properly implement these tools or to auditing personnel who need to confirm the work. New guidance therefore argues for clearer scaffolding akin to nutrition, prescription, or energy rating labels (Armanios et al., 2025b; Luccioni et al., 2024a; Nsoesie & Ghassemi, 2024).

The overarching point is that with algorithms interfacing with immersive reality, the labelling of the back-end algorithm or the immersive reality interface itself can serve as crucial scaffolding. Such scaffolding makes tasks easier to execute and attempts to reduce the

time and knowledge gaps needed to confirm the work's completion. How much scaffolding is needed, and where it can be tuned, is based on the nature and criticality of the project task at hand. Delivery-validation scaffolding thus requires deeper reflection at the interface between delivery and assurance. This is amplified in settings where AI interfaces with mixed reality, which juggles two forms of projection: one in the artificial and one in the physical. This leads to the following testable proposition:

Proposition 3 (Delivery-Validation Scaffolding): AI deployment in projects will have greater performance when deployed with appropriate scaffolding that bridges (a) time or (b) knowledge gaps.

Haptic: Workflow Readiness

Dialectic #4: AI requires automated and manual labor.

Project Principle #4: Account for when to increase process speed and precision

While many studies have focused on how automation enhances productivity (or not) (Kromann et al., 2020), recent studies have only started to recognize how context, such as industry, drives such effectiveness (Anglani et al., 2023). This is particularly challenging in AI-driven robotics, where decisions must be made around when manual interventions are needed just as AI is increasing the ability to automate processes.

To help crystallize this in a project setting, let us take two different projects that use AI-controlled robotics to automate project work. The first is automation in warehouses such as those in Amazon. While we choose to focus on examples where automation is increasingly ubiquitous, namely project logistics, such insights are transferable to other project functions. For instance, studies in construction demonstrate automation is also increasingly used in offsite manufacturing of pre-fabricated components or in operations and cleanup in hazardous settings too dangerous for humans (Trevelyan et al., 2016; Vähä et al., 2013).

While automation has benefits, such as reduced walking required, it has also resulted in greater human costs such as increased work intensity, speed, and repetitive stress injuries (Gutelius & Theodore, 2019; Vallas et al., 2022). For instance, one report of 1,484 frontline Amazon warehouse workers across 451 facilities in 42 states shows that 41 % of workers report being injured, 69 % take unpaid time off due to pain and exhaustion, 41 % feel pressure to work faster always or most of the time, and 60 % experience more workplace monitoring (Gutelius & Pinto, 2023). While robots can help move larger sets of packages at great speeds, human workers are still needed to pick and fulfill individual orders. Robots dictate speed, but humans are still required to ensure precision.

The second is automation of the genomic process at the Francis Crick Institute in London (Wynter, 2023). Recent advancements in genomics, namely in high-throughput screening methods, have expedited the drug discovery process through faster iteration across vast libraries of possible gene sequences to identify promising candidates (Adams et al., 1991; Barry III et al., 2000; Debouck & Metcalf, 2000). Sample preparation is a part of the process especially prone to human error and contamination. With automation, a scientist can now prep samples in half an hour rather than half a day and do so in a more reproducible manner (Wynter, 2023).

The key tension here is a *haptic* (touch) one. While algorithmically driven robotics allow for unprecedented automation and programmability, especially around standardized tasks, their dexterity remains limited (Schirmer et al., 2025). In the Amazon example, human dexterity is still needed, which increases the pressure on human workers. In the genomics example, dexterity exigencies have been better considered in the workflow to ensure a more standardizable and repeatable process that exerts less pressure on human workers.

The lever then is *task separability*, or the degree to which tasks can be partitioned so that some are algorithmically driven with only minimal or

periodic human intervention or supervision. Task separability as a lever requires deeper reflection on how one sequences project tasks between humans and algorithmically driven machines such as robots, especially around the haptic needs of the task. If the task needs more granular touch and motion, the haptic requirements are high and automating parts next to such work will put more strain and pressure on the human part of that process. This will result in performance gains quickly plateauing, if gains are seen at all. However, automation may be needed even if the haptic requirements are high due to extreme operating conditions. For such instances—as in the Crick Institute—project task flow could perhaps be redesigned to ensure the tasks with high haptic requirements “bookend” the workflow (i.e., human preparation of the inputs followed by automation that is adequately spaced out, with output handed off for more fine-tuned human post-processing). This leads to following testable propositions:

Proposition 4 (Task Separability): AI deployment in projects will have greater performance when deployed in (a) workflows with low haptic requirements or (b) workflows that increasingly “bookend” the high haptic task requirements at either the beginning or end of the workflow.

Exchange: Stakeholder Readiness

Dialectic #5: AI enables control and freedom.

Project Principle #5: Account for when to centralize and distribute oversight

AI can increase both control and freedom, and this dialectic is particularly challenging when AI is the backend for multi-agent systems (e.g., blockchain, marketplaces, or drone swarms). Such systems allow agents to freely operate and yet can also restrict the agent’s range of actions (Wooldridge, 2009). While multi-agent systems are increasingly deployed in projects, especially to improve outcomes around globally distributed project teams relying on scrum management (Lin et al., 2015; Yan et al., 2000), how to build stakeholder readiness for such deployment remains unclear.

Let’s take blockchain as an increasingly prevalent example of multi-agent systems used in projects, especially in industrial manufacturing contexts (Nguyen et al., 2022). Blockchain can link various sensors to collect increased breadth and depth of data with little oversight, ranging from factory conditions, screw sizes, all the way to managerial keyboard strokes, to name only a few measurable project metrics. At the same time, blockchain can be governed via smart contracts that are programmed with activation and termination conditions that can verify the completion of agreed-upon tasks and restrict transactions if conditions are not met (Lumineau et al., 2021a, 2021b). This is more generally akin to restricting the range of actions that an agent can perform, similar to establishing norms amongst agents (Wooldridge, 2009). With breakthroughs in the GenAI space that incorporate multi-agent systems, such as Manus AI (<https://manus.im/>), this interface between AI and multi-agent systems as seen in blockchain will only become more ubiquitous in projects.

The key project question clearly becomes how to govern such relationships: err on the side of greater control or freedom? To gain insight, let us consider two blockchain projects. While these projects are increasingly ubiquitous yet more operational in nature, such blockchain activities have similarly been used to deploy more end-to-end project work such as blockchain systems used for ensuring vaccine refrigeration across a supply chain (i.e., cold chain) (Mendonça et al., 2021). Beyond blockchain, such insights are transferrable into other settings such as multi-drone (agent) systems that are increasingly deployed in construction projects for surveying, safety, inspections, and even for demolition (Li & Liu, 2019).

The first blockchain project is Ant Group’s collaboration with HelloBike (Yang, 2023). HelloBike’s challenge was convincing investors of the promise of the e-bike sharing industry where several previous entrants had misled investors with fraudulent claims. In looking to prove the credibility of their claims and data, HelloBike partnered with Ant

Group to provide a reliable third-party certification of the data via AntChain. HelloBike uploaded e-bike usage data to a blockchain system via a sensor device on every bike that reported location and usage, along with cross-checked payment data. Investors were then added as nodes onto the system, so they could see the data in real time.

The second is a partnership between IBM and Maersk to create the platform TradeLens. The challenge for Maersk was how to reliably track and automate container shipping to improve efficiency (Lal & Johnson, 2018). The impetus for this was the sheer number of handoffs between different entities that must occur for a shipping order to be completed, and the fact that this was still a predominantly paper-based process that was open to fraud. The solution developed was a more permission-based blockchain where everyone knows who is on the system and their activity but “firewalled” to prevent the spread of commercially sensitive materials across competitors (Jovanovic et al., 2022).

What varies in these two examples is the use of more distributed (providing more freedom) vs more centralized (providing more control) governance (Hanisch et al., 2025). In the case of AntChain and HelloBike, the approach was a more distributed form of governance whereby every investor was added as nodes in the system and could access all e-bike sensor data. AntChain simply served as a reputable third-party certifier that authenticated the data provided from HelloBike. AntChain has no direct commercial interest in the sector other than collecting transaction revenues. For TradeLens, the platform represented more centralized governance. Maersk was directly tracking the transactions and had a direct commercial interest in the data, perhaps leading to the platform’s eventual discontinuation because there was not adequate uptake of the platform by other market players (Cecere, 2022; Maersk, 2022). Presumably, market players were worried about allowing Maersk to entrench and expand their market position by becoming the de facto information standard for the industry.

Regardless of the result, the tension here is an *exchange* one. From the two cases, we see the lever is how choices are made around *multi-agent governance*. One can decide whether a third-party certifies the platform, the project creates its own platform, or whether to use an entirely public platform such as Ethereum. Depending on the strength of stakeholder norms, one can decide whether the agents are given more or fewer constraints. In the aggregate, these toggles help decide whether the governance mechanisms across the agents are more distributed (free) or more centralized (monitored) in their exchanges.

Multi-agent governance as a lever requires deeper reflection as to the kind of system one wants to create to manage exchanges between project stakeholders, which may be amplified in settings where AI backends multi-agent systems. The ultimate decision will depend on overall stakeholder trust, norms, and performance stakes. If stakeholders have high trust or the stakes are high (i.e., little room for error and catastrophic consequences if things go wrong), one may want to operate their own multi-agent platform. This may even depend on whether stakeholders trust the agents themselves. Recent studies are alarmingly showing that AI agents can act as autonomous stakeholders in and of themselves. In so doing, these agents can act in their own interest, even if against the interest of the deploying project or organization (Lynch et al., 2025). If stakeholders have stronger norms, such as governments who are tightly bound by stringent regulations and requirements, then programming greater constraints on agents in the multi-agent system may be more warranted. This leads to the following testable proposition:

Proposition 5a (Multi-Agent Governance – Trust & Performance Stakes): If stakeholders have greater trust or performance stakes are high, self-operated and managed multi-agent systems will perform better than third-party certified or public multi-agent systems.

Proposition 5b (Multi-Agent Governance - Norms): If stakeholders have stronger norms, more constraints on agents will improve multi-agent system performance.

Resource: Ecological Readiness

Dialectic #6: AI optimizes and overuses computational resources
Project Principle #6: Account for when to use "Big Data" and "Small Data"

We are increasingly seeking to ensure our projects are sustainable, and the hope is that AI can help us achieve those outcomes more effectively (Khan et al., 2025). However, whether data practices and models are prepared for the increase in compute needed to operate the algorithms promising greater project performance and sustainability remains unclear.

This is becoming especially critical as AI paradoxically can both optimize and overuse ecological resources. In 2021, the amount of energy to tune a single BERT model (one of Google's key language models) required the same energy as a trans-American flight (Bender et al., 2021). As of 2024, generating a single image across the 8 most downloaded AI models accessible on Hugging Face averaged to about a 50% smartphone charge (Luccioni et al., 2024b).² Scaling this up, the International Energy Agency predicted energy from data centers in 2026 would be equivalent to the consumption of Japan (International Energy Agency, 2024). Water demand is equally drastic. Training GPT-3, an earlier ChatGPT model, was projected to withdraw over half the water consumption of the entire UK (4.2 – 6.6 billion m³) and, from that, consume (via evaporation) 0.38–0.60 billion m³ (Li et al., 2023). Yet at the same time, AI can also improve our grid and energy efficiency to better optimize our energy usage (Kaack et al., 2022; Wilson et al., 2024, 2025).

The key tension here is a *resource* one. To navigate this tension, project managers need to focus not simply on how AI can help them better plan their work, but also whether they have the ecological resources for such tooling. How then can we harness the potential resource benefits of AI in ways that work within increasingly constrained resource budgets? To do this, we need to interrogate a tacit assumption often made with AI—the need for “Big Data” (i.e., a large corpora of data) to realize benefits from AI.

Assume we are implementing a healthcare project focused on a rare form of cancer. Perhaps we do not need to train on large corpora of more generalist Internet-based data but fine-tune an existing model on narrower specialist oncological data. This is the premise of transfer learning models used to improve algorithmic accuracy, especially amongst marginalized populations with limited “small” data (Gao & Cui, 2020). This “piggybacking” can reduce computational resources.

In addition, do we really need billions of parameters to solve our problem? Perhaps we are only interested in a few parameters. In this case, we only need to focus on a limited set of criteria that can be done with more parsimonious AI models that do not need the complexity and largesse of typical GenAI models. Gandy et al. (2025) show how simple bicriteria estimations can help more inclusively ascertain which bridges to prioritize for rehabilitation or reconstruction. By focusing only on bridge condition (state of bridge disrepair) and equity (minimizing bridge condition between disadvantaged and advantaged communities), the authors could solve, using a single laptop, much of the equity issues for the entire U.S. bridge program, and do so with minimal differences in overall bridge condition.

The more general takeaway is that project managers must now navigate *resource* tensions when using AI through more thoughtful exploration of how much data and how many parameters they really need. Perhaps specialist data is adequate and trainable with existing models, or perhaps the project needs only a select few parameters and not billions. In scoping for small data or fewer criteria, the key lever is *model parsimony* (fewer variables) and/or *transfer learning* (enhanced

² Whether distilled models such as DeepSeek use substantially less energy to train remains to be seen (Heaven, 2025).

training of existing models on specialist “small” data) to more sustainably harness the value of AI in projects. This leads to the following testable proposition:

Proposition 6 (Model Parsimony & Transfer Learning): AI deployment in projects will have greater performance when (a) fewer parameters are prioritized (parsimony) or (b) specialist data is used (transfer learning).

4. Conclusion

This essay aims to be a provocation on where we place attention when deploying AI within projects and organizations. The prevailing preoccupation has been on evaluating the algorithms themselves and how they enhance project outcomes. However, given the mixed performance from AI deployments to date (Challapally et al., 2025; Dell'Acqua et al., 2023; Korst et al., 2025), we advocate focusing more carefully and holistically around how to increase organizational and project readiness to ensure more effective AI deployment. The result of our efforts is CIPHER, a novel dialectical framework that focuses on six key tensions: Cognitive, Informational, Projection, Haptic, Exchange, and Resource. Each tension characterizes a different set of opposing forces and is linked to a different key project principle and lever. We distill this framework into testable propositions that we hope will motivate future empirical research around how to enhance readiness for effective AI project deployment.

Given the speed at which AI developments are occurring, literally on a daily basis, updates to this CIPHER framework are all but certain. This is precisely why this is not the last word, but only a beginning “postcard from the edge” towards a more systematic and coherent conversation around the deployment of AI in projects. Perhaps such incorporation of AI could be the basis for a more general action-based perspective on projects and organizing (Winch, 2025). Regardless of where this collectively takes us, our sincere intent is that future empirical research interrogates this agenda-setting exercise to identify boundary conditions and to even add new tensions and levers to this framework.

CRediT authorship contribution statement

Daniel Erian Armanios: Writing – original draft. **Christopher L. Tucci:** Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We acknowledge the constructive feedback of Agni Orfanoudaki, Benn Lawson, and Graham Winch, our editor, on this manuscript. We also thank the faculty of the Major Programmes Area at the University of Oxford for their helpful feedback. Finally, we thank the numerous executives and civil servants who helped us stress-test this framework. All errors and omissions are our own.

References

- Adams, M. D., Kelley, J. M., Gocayne, J. D., Dubnick, M., Polymeropoulos, M. H., Xiao, H., Merril, C. R., Wu, A., Olde, B., & Moreno, R. F. (1991). Complementary DNA sequencing: Expressed sequence tags and human genome project. *Science*, 252, 1651–1656.
- Agarwal, N., Moehring, A., Rajpurkar, P., & Salz, T. (2023). Combining human expertise with artificial intelligence: Experimental evidence from radiology. *National Bureau of Economic Research*.
- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Boston, Massachusetts: Harvard Business Review Press.

- Anglani, F., Pennetta, S., Reaiche, C., & Boyle, S. (2023). Crossing digital frontiers with cultural intelligence - a new paradigm for project managers. *International Journal of Project Management*, 41.
- Armanios, D. E., McCulloch, M., Ventresca, M. J., & Itani, M. (2025a). Major program Value creation and Capture: The S3 Framework for assessing and minimizing risk propagation to maximize opportunities for value creation. *California Management Review*, 67, 117–143.
- Armanios, D. E., Tucci, C. L., & Roberts, S. (2025b). *Organizing principles for AI: A Provocation*. SSRN eLibrary.
- Barbosa, M. T. J., & Carvalho, M. M. (2025). Enhancing sensing capabilities in project-based organizations: The role of knowledge transfer between commercial and operational teams. *International Journal of Project Management*, 43.
- Barcaui, A., & Monat, A. (2023). *Who is better in project planning? Generative artificial intelligence or project managers?*, 4. Project Leadership and Society.
- Barry, C. E., III, Slayden, R. A., Sampson, A. E., & Lee, R. E. (2000). Use of genomics and combinatorial chemistry in the development of new antimycobacterial drugs. *Biochemical pharmacology*, 59, 221–231.
- 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 2021* (pp. 610–623).
- Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial Gender classification. In S. A. Friedler, & C. Wilson (Eds.), *Proceedings of the 1st Conference on Fairness, Accountability and Transparency* (pp. 77–91). PMLR.
- Busco, C., & Quattrone, P. (2018). In search of the "Perfect One": How accounting as a maieutic machine sustains inventions through generative "in-tensions. *Management Accounting Research*, 39, 1–16.
- Cecere, L. (2022). *Tradelens discontinues operations. Why you should care*. <https://www.forbes.com/sites/loracecere/2022/12/05/tradelens-discontinues-operations-why-you-should-care/>, Forbes.
- Chalhoub, J., & Ayer, S. K. (2018). Using Mixed Reality for electrical construction design communication. *Automation in Construction*, 86, 1–10.
- Challapally, A., Pease, C., Raskar, R., & Chari, P. (2025). *The genai divide: State of ai in business 2025*. MIT. https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf.
- Choudhary, V., Marchetti, A., Shrestha, Y. R., & Puranam, P. (2025). Human-AI ensembles: When can they work? *Journal of Management*, 51, 536–569.
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35, 128–152.
- Costantino, F., Di Gravio, G., & Nonino, F. (2015). Project selection in project portfolio management: An artificial neural network model based on critical success factors. *International Journal of Project Management*, 33, 1744–1754.
- Crisan, A., Drouhard, M., Vig, J., & Rajani, N. (2022). Interactive model cards: A Human-centered approach to model documentation. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 427–439). Seoul, Republic of Korea: Association for Computing Machinery.
- Cyert, R. M., & March, J. G. (1963). *A behavioral theory of the firm*. Englewood Cliffs, N.J: Prentice-Hall.
- Deboux, C., & Metcalf, B. (2000). The impact of genomics on drug discovery. *Annual Review of Pharmacology and Toxicology*, 40, 193–208.
- Dell'Acqua, F., Ayoubi, C., Lifshitz-Assaf, H., Sadun, R., Mollick, E.R., Mollick, L., Han, Y., Goldman, J., Nair, H., Taub, S., 2025. The cybernetic teammate: A field experiment on generative AI reshaping teamwork and expertise.
- Dell'Acqua, F., McFowland, E., III, Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayner, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality*. Harvard Business School Technology & Operations Mgt. Unit Working Paper.
- Diekmann, J. E. (1992). Risk analysis: Lessons from artificial intelligence. *International Journal of Project Management*, 10, 75–80.
- Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10.
- Gandy, C., Orfanoudaki, A., & Armanios, D. E. (2025). *Bridging gaps: Equity in infrastructure decisions*. SSRN eLibrary. <https://doi.org/10.2139/ssrn.5163140>
- Gao, Y., & Cui, Y. (2020). Deep transfer learning for reducing health care disparities arising from biomedical data inequality. *Nature Communications*, 11.
- Gavetti, G., Helfat, C. E., & Marengo, L. (2017). Searching, shaping, and the quest for superior performance. *Strategy Science*, 2, 194–209.
- Gavetti, G., & Levinthal, D. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45, 113–137.
- Geraldi, J., Locatelli, G., Dei, G., Söderlund, J., & Clegg, S. (2024). AI for management and Organization research: Examples and reflections from project studies. *Project Management Journal*, 55, 339–351.
- Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15, 6.
- Gutelius, B., & Pinto, S. (2023). *Pain points: Data on work intensity, monitoring, and health at amazon warehouses*. Chicago: University of Illinois Chicago. <https://doi.org/10.25417/uic.24435124.v2>
- Gutelius, B., & Theodore, N. (2019). *The future of warehouse work: Technological change in the U.S. logistics industry*. Berkeley: UC Berkeley, Center for Labor Research and Education. <https://escholarship.org/uc/item/90p6f8rg>.
- Hanisch, M., Goldsby, C., Fortin, M., & Rogerson, M. (2025). Governing within semirigid limits: Navigating the centralization–Decentralization paradox in blockchain-based platforms. *Journal of Operations Management*.
- Hannigan, T. R., McCarthy, I. P., & Spicer, A. (2024). Beware of botshit: How to manage the epistemic risks of generative chatbots. *Business Horizons*, 67, 471–486.
- Heaven, W. D. (2025). How DeepSeek's AI compares to established models. *Science Friday*.
- Holzmann, V., Zitter, D., & Peshkess, S. (2022). The expectations of project managers from artificial intelligence: A Delphi study. *Project Management Journal*, 53, 438–455.
- Hu, L., Yuan, S., Qin, X., Zhang, J., Lin, Q., Zhang, D., Rajmohan, S., Zhang, Q., 2025. MEETING DELEGATE: Benchmarking LLMs on attending meetings on our behalf. arXiv preprint arXiv:2502.04376.
- Iantisi, M., Lakhani, K. R., & Cohen, S. J. (2020). *Competing in the age of ai*. Findaway World.
- Ingason, H. T., Aaltonen, K., Asmundarson, A. S., Fridgeirsson, T. V., Huemann, D., Huemann, M., Kujala, J., Lampela, H., Mancini, M., Mariani, C., & Ringhofer, C. (2025). Personalised learning in project management education: Insights from an artificial intelligence-driven chatbot. *Project Leadership and Society*, 6.
- International Energy Agency. (2024). *Electricity 2024: Analysis and forecast to 2026*. International Energy Agency. <https://iea.blob.core.windows.net/assets/6b2fd954-2017-408e-bf08-952fdd62118a/Electricity2024-Analysisandforecastto2026.pdf>.
- Jin, X. H., & Zhang, G. M. (2011). Modelling optimal risk allocation in PPP projects using artificial neural networks. *International Journal of Project Management*, 29, 591–603.
- Jovanovic, M., Kostić, N., Sebastian, I. M., & Sedej, T. (2022). Managing a blockchain-based platform ecosystem for industry-wide adoption: The case of TradeLens. *Technological Forecasting and Social Change*, 184, 121981. <https://doi.org/10.1016/j.techfore.2022.121981>.
- Kaack, L. H., Donti, P. L., Strubell, E., Kamiya, G., Creutzig, F., & Rolnick, D. (2022). Aligning artificial intelligence with climate change mitigation. *Nature Climate Change*, 12, 518–527.
- Khan, A. N., Bodla, A. A., Moin, M. F., & Soomro, M. A. (2025). Guest editorial: Building a sustainable future with AI: The next frontier in project management. *International Journal of Managing Projects in Business*, 1–9.
- Kloberdanz, K. (2017). *Looking smart: Augmented reality is seeing real results in industry*. GE. <https://www.ge.com/news/taxonomy/term/6060>.
- Koo, W. W. (2025). Cross-lingual effects of AI-generated content on human work. *Scientific Reports*, 15, Article 30949.
- Korst, J., Puntoni, S., Tambe, P., Smith, B., Odishoo, C., & Urbina-McCarthy, D. (2025). *Accountable acceleration: Gen ai fast-tracks into the enterprise*. Wharton Human-AI Research, University of Pennsylvania. <https://ai.wharton.upenn.edu/wp-content/uploads/2025/10/2025-Wharton-GBK-AI-Adoption-Report-Full-Report.pdf>.
- Kromann, L., Malchow-Moller, N., Skaksen, J. R., & Sorensen, A. (2020). Automation and productivity—a cross-country, cross-industry comparison. *Industrial and Corporate Change*, 29, 265–287.
- Lal, R., & Johnson, S. (2018). *Maersk: Betting on blockchain*. Harvard Business School Case, 9-518-089.
- Lanzolla, G., Pesce, D., & Tucci, C. L. (2021). The digital transformation of search and recombination in the innovation function: Tensions and an integrative framework. *Journal of Product Innovation Management*, 38, 90–113.
- Lee, H.-P.H., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., Wilson, N., 2025. The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of Knowledge Workers, CHI, Yokohama, Japan.
- Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science*, 43, 934–950.
- Levitt, R. E., Kartam, N. A., & Kunz, J. C. (1988). Artificial-intelligence techniques for generating construction project plans. *Journal of Construction Engineering and Management-Asce*, 114, 329–343.
- Li, P., Yang, J., Islam, M.A., Ren, S., 2023. Making ai less "thirsty": Uncovering and addressing the secret water footprint of ai models. arXiv preprint arXiv:2304.03271.
- Li, Y., & Liu, C. L. (2019). Applications of multirotor drone technologies in construction management. *International Journal of Construction Management*, 19, 401–412.
- Lin, Y., Descamps, P., Gaud, N., Hilaire, V., & Koukam, A. (2015). Multi-Agent System for intelligent scrum project management. *Integrated Computer-Aided Engineering*, 22, 281–296.
- Longpre, S., Singh, N., Cherep, M., Tiwary, K., Materzynska, J., Brannon, W., Mahari, R., Obeng-Marnu, N., Dey, M., Hamdy, M., 2024. Bridging the data provenance gap across text, speech and video. arXiv preprint arXiv:2412.17847.
- Luccioni, S., Gamazaychikov, B., Hooker, S., Pierrard, R., Strubell, E., Jernite, Y., & Wu, C.-J. (2024a). Light bulbs have energy ratings—So why can't AI chatbots? *Nature*, 632, 736–738.
- Luccioni, S., Jernite, Y., & Strubell, E. (2024b). Power hungry processing: Watts driving the cost of AI deployment?. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency. Association for Computing Machinery* (pp. 85–99). Brazil: Rio de Janeiro.
- Lumineau, F., Wang, W. Q., & Schilke, O. (2021a). Blockchain governance—A new way of organizing collaborations? *Organization Science*, 32, 500–521.
- Lumineau, F., Wang, W. Q., Schilke, O., & Huang, L. (2021b). *How blockchain can simplify partnerships*. Harvard Business Review. <https://hbr.org/2021/2004/how-blockchain-can-simplify-partnerships>.
- Lynch, A., Wright, B., Larson, C., Ritchie, S.J., Mindermann, S., Perez, E., Troy, K.K., Hubinger, E., 2025. Agentic misalignment: How LLMs could be insider threats. arXiv preprint arXiv:2510.05179.
- Maersk. (2022). A.P. Moller - Maersk and IBM to discontinue TradeLens, a blockchain-enabled global trade platform. <https://www.maersk.com/news/articles/2022/11/29/maersk-and-ibm-to-discontinue-tradelens>.
- Mancini, M., Mariani, C., & Manfredi, C. M. (2023). Nuclear decommissioning risk management adopting a comprehensive artificial intelligence framework: An applied case in an Italian site. *Progress in Nuclear Energy*, 158.
- March, J. G., & Simon, H. A. (1958). *Organizations*. New York: Wiley.

- Mariani, C., & Mancini, M. (2025). Harnessing AI for value: Examining the impact of AI capabilities and the mediating role of organizational agility on project value proposition. *International Journal of Managing Projects in Business*, 18, 112–143.
- Mariani, C., Navrotska, Y., & Mancini, M. (2023). Unsupervised machine learning for project stakeholder classification: Benefits and limitations. *Project Leadership and Society*, 4.
- Maylor, H., Smith, V., & Tully, P. (2024). *Further education and skills: Changing systems of change*. Oxford Saïd Business School and Education and Training Foundation. https://www.et-foundation.co.uk/wp-content/uploads/2024/05/ETF025_Changing-Systems-of-Change-Online_v5.pdf.
- Mendonça, R.D., Gomes, O.S., Vieira, L.F., Vieira, M.A., Vieira, A.B., Nacif, J.A., 2021. Blockcoldchain: Vaccine cold chain blockchain. arXiv preprint arXiv:2104.14357.
- Microsoft, 2021. Lockheed Martin builds for the next stage in human space exploration with Microsoft HoloLens 2, <https://www.microsoft.com/en/customers/story/857748-lockheed-martin-discrete-manufacturing-hololens-en-united-states>.
- Mikalef, P., & Gupta, M. (2021). Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. *Information & Management*, 58.
- Miller, G. J. (2022). Stakeholder roles in artificial intelligence projects. *Project Leadership and Society*, 3.
- Minaee, S., Mikolov, T., Nikzad, N., Chenaghlu, M., Socher, R., Amatriain, X., Gao, J., 2024. Large language models: A survey. arXiv preprint arXiv:2402.06196.
- Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for Model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 220–229). Atlanta, GA, USA: Association for Computing Machinery.
- Mollick, E. (2024). *Co-intelligence: Living and working with ai*. London: WH Allen.
- Morby, A., 2018. IKEA assembly made easier through augmented-reality app, Dezeen, <https://www.dezeen.com/2018/03/23/ikea-assembly-made-easier-through-augmented-reality-app/>.
- Müller, R., Locatelli, G., Holzmann, V., Nilsson, M., & Sagay, T. (2024). Artificial intelligence and project management: Empirical overview, State of the art, and guidelines for future research. *Project Management Journal*, 55, 9–15.
- Nature Research Intelligence. (2025a). Artificial intelligence. *Nature*. <https://www.nature.com/research-intelligence/nri-topic-summaries/artificial-intelligence-for-12-4602>.
- Nature Research Intelligence. (2025b). Machine learning. *Nature*. <https://www.nature.com/research-intelligence/nri-topic-summaries/machine-learning-for-12-4611>.
- Nguyen, L. D., Bröring, A., Pizzol, M., & Popovski, P. (2022). Analysis of distributed ledger technologies for industrial manufacturing. *Scientific Reports*, 12, Article 18055.
- Nsoesie, E. O., & Ghassemi, M. (2024). Using labels to limit AI misuse in health. *Nature Computational Science*, 4, 638–640.
- Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (pp. 1–22).
- Petratos, P. N. (2021). Misinformation, disinformation, and fake news: Cyber risks to business. *Business Horizons*, 64, 763–774.
- Project Management Institute. (2024a). *First movers' advantage: The immediate benefits of adopting generative AI for project management*. Project Management Institute (PMI), Inc.
- Project Management Institute. (2024b). *Pushing the limits: Transforming project management with genai innovation*. Project Management Institute (PMI), Inc.
- Puranam, P. (2021). Human–AI collaborative decision-making as an organization design problem. *Journal of Organization Design*, 10, 75–80.
- Ren, D., Goldschwendt, T., Chang, Y., & Höllerer, T. (2016). Evaluating wide-field-of-view augmented reality with mixed reality simulation. In *2016. IEEE Virtual Reality (VR)* (pp. 93–102). IEEE.
- Sacks, R., Whyte, J., Swissa, D., Raviv, G., Zhou, W., & Shapira, A. (2015). Safety by design: Dialogues between designers and builders using virtual reality. *Construction Management and Economics*, 33, 55–72.
- Schirmer, F., Kranz, P., Rose, C. G., Schmitt, J., & Kaupp, T. (2025). Towards dynamic Human–Robot collaboration: A holistic framework for assembly planning. *Electronics*, 14, 190.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99–118.
- Simon, H. A. (1997). *Administrative behavior: A study of decision-making processes in administrative organizations* (4th ed). New York: Free Press.
- Sinaga, K. P., & Yang, M.-S. (2020). Unsupervised K-means clustering algorithm. *IEEE access : Practical innovations, open solutions*, 8, 80716–80727.
- Smit, M., Wagner, R. F., & Bond-Barnard, T. J. (2025). Ambiguous regulations for dealing with AI in higher education can lead to moral hazards among students. *Project Leadership and Society*, 6.
- Trevelyan, J., Hamel, W. R., & Kang, S.-C. (2016). Robotics in hazardous applications. *Springer handbook of robotics* (pp. 1521–1548). Springer.
- Upskill, 2017. Upskill & GE renewable energy: Wind turbine wiring productivity study using Skylight, YouTube, <https://www.youtube.com/watch?v=n5LhQqggGTE>.
- Vähä, P., Heikkilä, T., Kilpeläinen, P., Järviuoma, M., & Gambaio, E. (2013). Extending automation of building construction - survey on potential sensor technologies and robotic applications. *Automation in Construction*, 36, 168–178.
- Vallas, S. P., Johnston, H., & Mommodova, Y. (2022). Prime suspect: Mechanisms of labor control at Amazon's warehouses. *Work and Occupations*, 49, 421–456.
- Van Alstyne, M.W., 2023. Free speech and the fake news problem. Available at SSRN 4414261.
- Van Der Vegt, A. H., Scott, I. A., Dermawan, K., Schnetler, R. J., Kalke, V. R., & Lane, P. J. (2023). Implementation frameworks for end-to-end clinical AI: Derivation of the SALIENT framework. *Journal of the American Medical Informatics Association*, 30, 1503–1515.
- Vanneste, B. S., & Puranam, P. (2024). Artificial intelligence, trust, and perceptions of agency. *Academy of Management Review*. amr. 2022.0041.
- Vettori, O., & Warm, J. (2025). The race for AI skills as an obstacle course: Institutional challenges and low threshold suggestions. *Project Leadership and Society*, 6.
- Wang, Y. R., Yu, C. Y., & Chan, H. H. (2012). Predicting construction cost and schedule success using artificial neural networks ensemble and support vector machines classification models. *International Journal of Project Management*, 30, 470–478.
- Watkins, C. J., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8, 279–292.
- Whyte, J., Stasis, A., & Lindkvist, C. (2016). Managing change in the delivery of complex projects: Configuration management, asset information and 'big data'. *International Journal of Project Management*, 34, 339–351.
- Wilson, C., Agnew, M., Amanta, F., Van Fan, Y., Kumar, P., & Seger, M. (2024). In *Evidence Synthesis of Indirect Impacts of Digitalisation on Energy and Emissions, 2024 10th International Conference on ICT for Sustainability (ICT4S)*. (pp. 116–127). IEEE.
- Wilson, C., Fan, Y. V., & Amanta, F. (2025). AI's indirect impacts on climate outweigh concerns over its direct energy footprint. *The Oxford Institute for Energy Studies - Forum*, 10–16.
- Winch, G. M. (2025). An action theory of the project. *Project Management Journal*, 56, 299–309.
- Winick, E. (2018). NASA is using HoloLens AR headsets to build its new spacecraft faster. *MIT Technology Rev*.
- Wooldridge, M. J. (2009). *An introduction to multiagent systems* (2nd ed.). Chichester: Wiley.
- Wu, W., Hartless, J., Tesei, A., Gunji, V., Ayer, S., & London, J. (2019). Design assessment in virtual and mixed reality environments: Comparison of novices and experts. *Journal of Construction Engineering and Management*, 145, Article 04019049.
- Wynter, F., 2023. Automata partners with The Francis Crick Institute on integrated automation for genomics. Automata, <https://automata.tech/company-news/automata-partners-with-the-francis-crick-institute/>.
- Yan, Y. H., Kuphal, T., & Bode, J. (2000). Application of multiagent systems in project management. *International Journal of Production Economics*, 68, 185–197.
- Yang, Y. T., & Zhao, B. (2025). How building information modelling mitigates complexity and enhances performance in large-scale projects: Evidence from China. *International Journal of Project Management*, 43.
- Yang, S.A. and N. Jain (2023). "AntChain's Blockchain as a Service: Digitising Industry Collaboration." London Business School Case, CS–23–011.