



The Human Side of AI: Culture, Control, and Collaboration in the Age of Intelligent Work

Research & Insights from the
Culturati:AI Taskforce

**SPRING
2025**

Editor's Note

As AI reshapes the fabric of modern work, culture leaders are facing a cascade of urgent decisions—from how to integrate new tools into legacy systems to how to preserve trust, purpose, and humanity in increasingly automated environments. The Culturati:AI Taskforce commissioned this collection to equip senior executives with practical frameworks, cautionary insights, and forward-looking ideas on what it means to build, scale, and lead with AI—without losing the human thread.



Table of Contents

Leading the AI Transformation: A Human-First Approach to Enterprise Adoption

By Brandon Gadoci, Vice President of AI Operations, data.world

Page 1

Start here if you're looking for a practical blueprint for integrating AI across your organization. This article lays out a structured approach for operationalizing AI, focusing not just on tools but on cultivating trust, experimentation, and cross-functional buy-in. It introduces AI Operations as a discipline, offers a people-first adoption playbook, and outlines the three phases of enterprise rollout. This is the foundational piece that sets the tone for all others.

Navigating AI in the Workplace: A Practical Taxonomy for Business Leaders

By David Altounian, Ph.D., Vice Provost and Associate Professor of Entrepreneurship, Salve Regina University

Page 11

Start here if you want a clear, strategic framework for classifying and deploying AI responsibly across functions. This piece provides a five-tier taxonomy—Assistive, Augmentative, Semi-Autonomous, Autonomous, and Agentic—helping leaders evaluate risk, autonomy, and human oversight. It expands on the operational mindset introduced earlier, offering language and structure to assess use cases and avoid misapplication of high-autonomy systems.

AI Is Your New Coworker—But Is It Taxing Your Culture?

By Ethan Burris, Ph.D., Senior Associate Dean, McCombs School, The University of Texas at Austin

Page 27

Start here if you're concerned about AI's impact on trust, collaboration, and creativity inside teams. This essay explores AI's cultural and emotional consequences across the HR lifecycle—from recruiting to performance feedback to team formation. It challenges the assumption that efficiency always equals improvement and warns of the hidden costs of technological distancing. Builds on the earlier pieces by shifting focus from systems to lived employee experience.

AI – Built on Human Collective: Why is AI a ‘Human Resource’?

By Shakeel Rashed, Culturati:AI Taskforce Chair; Board of Advisors, Capital Factory

Page 38

Start here if you're in HR or responsible for safeguarding workplace culture. This piece argues for human resources to take a central role in governing AI, not just reacting to it. It links emerging developments like agentic AI, reasoning models, and humanoids to everyday people practices—and explains why AI reflects (and requires) our collective ethics, creativity, and bias awareness. A natural follow-on from the cultural focus of the previous article, but with a proactive lens.

The Paradox of AI: Empowering Agency or Eroding Control?

By Trei Brundrett, Co-Founder of Vox Media

Page 51

Start here if you're thinking about individual empowerment. This reflective piece looks at the paradox of agency in an AI-enabled world: are we gaining freedom through delegation, or losing critical thinking by outsourcing decisions? It's the most philosophical of the group and offers a timely reminder that AI isn't just changing systems—it's reshaping how we relate to ourselves and our own capabilities. Best read after the others to zoom back out and re-center on the human.



Leading the AI Transformation: A Human-First Approach to Enterprise Adoption

How To Guide Your Organization Through The Next Great Workforce Evolution

By Brandon Gadoci, Vice President of AI Operations at data.world

Executive Summary

The rise of AI presents leaders with a critical choice: proactively guide their workforce through this transformation or risk falling behind. This article presents a structured framework for enterprise-wide AI adoption that prioritizes people over technology. Drawing from real-world implementation experience at data.world, where this approach led to a 25% company-wide productivity increase, it outlines how to identify and support key personas, build internal champions, and measure success across three phases of adoption. Leaders will learn how to create an environment where AI augments human potential rather than replacing it, ultimately building a more capable and confident workforce.

When ChatGPT burst onto the scene in November 2022, it marked more than just the launch of the fastest-growing product of all time. It was the moment AI moved from the back office to the front office, creating a challenge that companies are still wrestling with today: how to actually implement and harness this incredibly powerful, once-in-a-lifetime, general-purpose technology that promises to change everything. According to a Harvard Business Review article from 2019, only about 8% of companies had AI adoption efforts inside their organizations. So when ChatGPT made the full power of AI available to everyone, companies weren't ready. And now, they're left asking, "What should we be doing with this AI thing?"

THE ENTERPRISE AI GAP

AI in the enterprise has been around much longer than people think. Generative AI democratized access, putting it in everyone's hands for the first time. With its potential came a gap in skills and experience, leading to fear, ambiguity, and hesitation. Organizations struggled to balance innovation with risk, facing challenges in education, governance, and ethical concerns. Bridging this gap became essential for responsible AI adoption.



Just as the HBR article concluded—and remains true six years later—the challenge companies face is part technology, but more importantly, culture. As this technology lands on employees' desks, they're left wondering if they'll be out of a job soon, or they're hesitant to admit they're using the technology for fear of being found out. While companies wrestle with whether their data is in good enough shape to leverage AI, they should be asking if their people are ready. As a leader in the enterprise, it's natural and important to pause in a moment of such significance to consider next steps—but that pause should be just that: a pause. Failure to act here will not only put you behind the technological curve but also fail on a much bigger level: embracing and preparing your people for the future of work.

I. From Fear to Innovation: Reframing AI in the Enterprise

Unlike previous technological revolutions that primarily affected specific tasks or departments, AI's impact reaches into the very heart of knowledge work—the thinking, creating, and decision-making that have long been exclusively human domains. This broader reach creates both unprecedented opportunities and deeper anxieties among employees at all levels of the organization.

The challenge facing enterprise leaders today isn't just technological adoption—it's psychological transformation. Every organization confronts a crucial choice: proactively guide their workforce through this transformation or let fear and uncertainty dictate the pace of change. History shows us that resistance to transformative technologies often follows a predictable pattern. Organizations initially react with restriction and control—much like high schools banning new

technologies out of fear—only to later scramble to embrace and teach these same tools. The lesson is clear: don't let fear drive your response to innovation.

"You can either lead your workforce through this metamorphosis or leave them in the cocoon to figure it out themselves. The latter creates a workforce that's fearful for their jobs, scared to experiment, and unwilling to share what's really working—no environment for innovation."

The path forward requires more than just implementing new tools—it demands creating an environment where experimentation is encouraged, where failure is viewed as learning, and where employees feel empowered to reimagine their roles with AI as a partner rather than a threat. This transformation begins with addressing the fundamental fears and uncertainties that naturally arise during periods of significant change, turning them into catalysts for innovation rather than barriers to progress.

The organizations that thrive won't be those that resist this transformation but those that actively guide their workforce through it. Success requires a structured approach that acknowledges both the technical and human dimensions of AI adoption—an approach we call AI Operations.

II. AI Operations: A Framework for Human-Centric Transformation

Thanks to my good friend (and former intern) Rachel Wood's work, there is an emerging discipline we're calling AI Operations. Unlike traditional IT-focused approaches that primarily concern the maintenance and monitoring of AI systems, AI Operations (AI Ops) is a structured framework that embeds AI into the heart of organizational operations, transforming how businesses operate at every level. It prioritizes starting with people and processes and ending with solutions and tools.

AI Ops requires a unique blend of skills: detective work, consulting, operational expertise, systems thinking, and technical development. The goal isn't to create super AI products but to empower "superhuman" employees by augmenting human capabilities. By eliminating mundane tasks, AI enables employees to focus on higher-value, creative, and strategic activities.

Twenty months ago at data.world, Brett Hurt (CEO), Matt Laessig (COO), and Bryon Jacob (CTO) tapped me on the shoulder to build upon Rachel's ideas and see what they would look like in practice. The result was remarkable: a 25% lift in productivity company-wide, 200% increase in AI tool usage on a weekly basis, increased revenue, and doubling of profit margins—all while maintaining the same number of employees.

For most organizations, AI Ops should report to the Chief Operating Officer (COO) to maintain an operational focus and drive business impact. This reporting structure ensures AI remains operationally relevant and has the cross-functional influence needed to drive enterprise-wide adoption. While a close partnership with the CTO ensures technical soundness, having AI Ops report to operations rather than IT emphasizes its role in transforming how work gets done rather than just implementing technology.

The promise of AI Ops extends far beyond efficiency gains. It offers a practical pathway for:

- Providing clear strategic direction for AI adoption
- Empowering teams by equipping them with AI tools that enhance their capabilities
- Unlocking innovation by lowering the cost and complexity of experimentation
- Reducing fear around AI by reframing it as a tool for augmentation, not replacement
- Creating "superhuman" employees who can achieve extraordinary results through AI collaboration

This is why I've come to believe that creating this discipline is the key to bridging the enterprise AI adoption gap. But before we go into more detail, let's first understand who we're talking about.

III. The Six Faces of AI Transformation

In every organization, we consistently observe six distinct personas that emerge during AI transformation, each bringing their own perspective and needs to the journey:

The Enthusiastic Pioneer eagerly embraces AI and actively seeks ways to integrate it into their workflow. Emma quickly learns new tools and shares discoveries with colleagues, often becoming a natural champion for AI adoption. While her enthusiasm is invaluable, she needs structured guidance to prevent burnout and ensure her energy translates into sustainable practices.

The Analytical Evaluator takes a measured approach, carefully questioning AI's value and feasibility. While this skepticism might initially seem like resistance, Sam's critical thinking often leads to more robust and well-considered AI implementations. By providing data-driven evidence and clear ROI metrics,

organizations can transform these evaluators into powerful advocates for thoughtful AI adoption.

The Concerned Professional harbors deeper concerns about AI's impact on job security and career evolution. These employees need clear communication about how AI will augment rather than replace their roles, along with concrete examples of how AI can enhance their career development. When properly supported, they often become compelling voices for balanced AI adoption.

The Curious Observer watches AI developments with interest but remains uncertain about how to begin engaging with the technology. These employees benefit most from structured learning opportunities and hands-on experimentation in low-risk environments. When paired with Enthusiastic Pioneers, they often discover creative applications for AI that others might miss.

The Process Defender prefers existing workflows and sees AI as unnecessary complexity. Rather than forcing change, successful organizations demonstrate how AI can enhance rather than disrupt proven processes. Small, incremental improvements that respect existing workflows often help these employees recognize AI's value on their own terms.

The Ethical Guardian raises essential questions about AI's implications and alignment with organizational values as well as ethical questions about AI in general (e.g. how it was trained, the amount of resources it consumes). Their concerns about data privacy, algorithmic bias, and responsible AI use are crucial for developing trustworthy AI practices. By including them in discussions about AI governance and ethical frameworks, organizations can build more responsible and sustainable AI operations.

IV. Building Your AI Champions Network

The success of AI adoption hinges not on technology alone, but on the passionate individuals who champion its potential throughout the organization. These champions emerge naturally from the ranks of your employees—they're the ones who already show curiosity about AI and bring contagious energy to new initiatives. But identifying these potential leaders is just the beginning; their effectiveness depends on how well you nurture and support their development.

Creating a formal AI Champions Program provides these emerging leaders with the structure they need to thrive. At data.world, we found that regular training sessions on AI tools and best practices form the foundation, but the real magic happens when champions are given opportunities to lead workshops and demonstrate their

expertise to colleagues. Recognition plays a crucial role too—whether through formal rewards for successful implementations or dedicated time for experimentation, acknowledging their contributions reinforces their importance to the organization's AI journey.

The most effective champion networks span across departments, creating a web of advocates who understand different business contexts and can speak to diverse use cases. A marketing champion might discover AI applications that inspire their colleagues in finance, while an operations champion might help customer service teams reimagine their workflows. This cross-pollination of ideas accelerates adoption and ensures AI solutions address real business needs rather than remaining isolated experiments.

Education forms the final pillar of champion development, but it must go beyond technical training. The most successful champions combine deep AI knowledge with strong change management skills. They learn to demonstrate tools effectively, address concerns with empathy, and guide colleagues through their AI learning journey. Most importantly, they become masterful storytellers, sharing success stories that inspire others to embrace AI's potential. Through their examples, these champions transform AI from an abstract technology into a tangible force for positive change within the organization.

V. The Three Phases of Enterprise AI Adoption

AI Operations guides organizations through a structured journey of AI adoption, one that unfolds like a well-orchestrated story rather than a rigid implementation plan. This journey begins with moments of inspiration, develops through systematic learning, and culminates in enterprise-wide transformation. Each phase builds upon the success of its predecessor, creating a natural progression that respects both human and organizational readiness.

The journey begins with what we call the "Wow" Phase—a period of discovery and excitement that captures the imagination of your workforce. During these early days, the focus is on quick wins and simple automations that demonstrate AI's potential without requiring major organizational changes. We've found that when employees first experience AI tools in their daily work, their skepticism often transforms into enthusiasm. A customer service representative discovers they can draft responses in seconds rather than minutes, or a marketing team member realizes they can generate creative concepts at unprecedented speed. These early wins require no complex infrastructure changes; they rely simply on readily available information and a willingness to experiment in low-risk scenarios.

As initial excitement gives way to practical application, organizations enter the Easy Scaling & Education phase. This is where scattered experiments evolve into structured deployment and systematic learning. The casual AI users of the "Wow" phase become regular practitioners, incorporating AI tools into their daily operations with increasing sophistication. However, this broader adoption brings new challenges. Employees who once dabbled in AI now need structured training to maximize its effectiveness. Even the skeptics, seeing their colleagues' success, begin to explore AI's potential. This phase also marks a shift in leadership focus, as questions of security, compliance, and governance move to the forefront. It's a delicate balance of maintaining momentum while establishing necessary guardrails.

The final phase—Enterprise Data Cleanup & AI Integration—represents AI's evolution from useful tool to operational necessity. Here, AI tools become seamlessly woven into the fabric of core systems, transforming from standalone solutions into integrated capabilities. Organizations begin to view their data through a new lens, recognizing that AI's effectiveness depends on data quality and governance. Cross-departmental collaboration reaches new heights as teams discover how their AI initiatives complement and enhance each other's work. This mature phase isn't just about technology integration—it's where long-term AI strategy takes shape, setting the stage for sustained innovation and competitive advantage.

VI. Beyond ROI: Measuring Success in AI Transformation

Throughout history, transformative technologies have reshaped industries long before their full return on investment could be precisely measured. When email first emerged, businesses hesitated, uncertain whether digital communication would justify the shift from traditional memos and phone calls. The internet faced similar skepticism—many questioned the value of having a website or digital presence.

"The organizations that will thrive in the age of AI won't be those with the most advanced technology or the largest data sets – they'll be the ones that successfully combine human intelligence with artificial intelligence, creating workplaces where both can reach their full potential."

Today, we're seeing the same pattern with AI. Organizations seeking immediate, quantifiable returns risk missing the broader, long-term impact AI can deliver. The value of AI—like its technological predecessors—often begins with qualitative improvements that lay the foundation for future financial gains. Success manifests first in the cultural sphere, where we see growing employee enthusiasm, increased cross-functional collaboration, and the organic emergence of AI champions. These

early indicators, while harder to quantify, signal the deep organizational changes that precede measurable business impact.

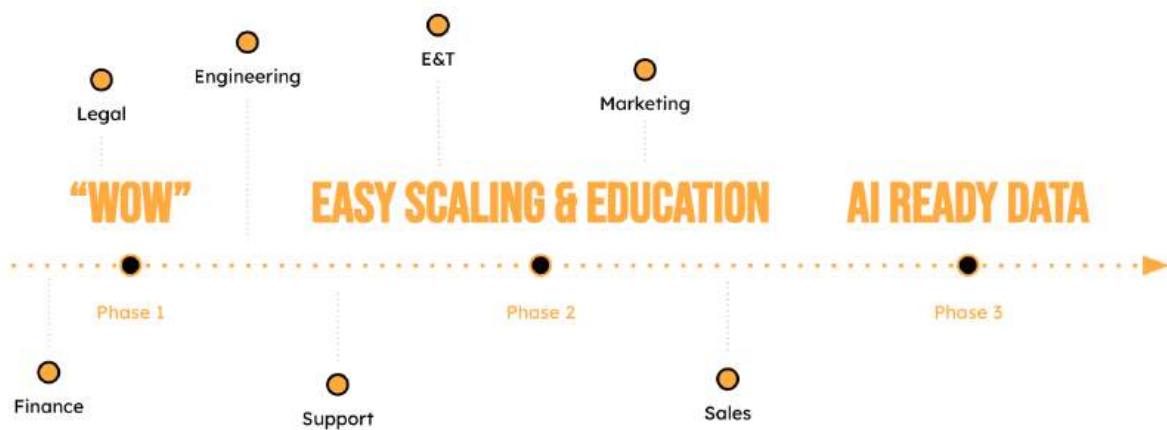
As adoption matures, capability metrics begin to tell a compelling story. We observe increasing AI literacy across roles, with employees not only using AI tools more frequently but also applying them in increasingly sophisticated ways. The speed at which teams develop and implement new use cases accelerates, creating a virtuous cycle of innovation and adoption. These capability improvements, while more tangible than cultural shifts, still precede the full financial impact of AI transformation.

The final dimension of success emerges in concrete business outcomes. Productivity improvements become measurable, innovation cycles compress, and customer satisfaction scores reflect the enhanced capabilities AI enables. Perhaps most importantly, organizations begin to see competitive advantages emerge—whether through faster market response, more personalized customer experiences, or more efficient operations. These business impacts, while the most quantifiable, represent the culmination of successful cultural and capability transformations.

The key to measuring AI success lies in understanding this progression and resisting the urge to focus solely on immediate ROI. Organizations that nurture cultural change, invest in capability building, and maintain patience while tracking business impacts position themselves to realize AI's full transformative potential. Just as email and the internet eventually became indispensable business tools, AI will prove its value—but only for organizations that take a holistic view of success.

VII. Embracing the Journey: AI Adoption as Continuous Evolution

The path to enterprise-wide AI adoption isn't a linear journey with a clear endpoint—it's an organic process that evolves differently across various parts of the organization. Just as each department has its unique workflows, challenges, and priorities, every individual brings their own perspective and pace to AI adoption. Some teams may eagerly embrace AI tools, while others require more time and evidence before integrating new technologies. This diversity in adoption rates isn't a barrier but a natural part of organizational evolution that requires careful nurturing and sustained support.



Understanding this organic nature of AI adoption is crucial for success. Marketing teams might rapidly embrace AI for content generation and campaign optimization, while finance departments may take a more measured approach, carefully validating AI outputs against established processes. Neither approach is wrong—they simply reflect the different contexts and requirements of each function. The key is recognizing and respecting these variations while providing appropriate support for each group's journey.

This understanding shapes how organizations should approach AI implementation. Rather than forcing a one-size-fits-all solution, successful organizations create flexible frameworks that accommodate different adoption speeds and styles. They recognize that some departments will serve as natural testing grounds for new AI applications, while others will benefit from watching and learning from these early experiences. This organic approach allows best practices to emerge naturally and spread through the organization in a way that feels authentic rather than forced.

The journey also demands a commitment to continuous learning and adaptation. As AI technology evolves at a rapid pace, organizations must foster a culture where ongoing education becomes second nature. This isn't just about technical training—it's about creating an environment where employees feel empowered to experiment, share their learnings, and build upon each other's successes. When teams see AI adoption as a journey of continuous improvement rather than a destination, they're more likely to embrace the technology in ways that truly transform their work.

VIII. The Leadership Imperative

The time for waiting is over. As AI continues to reshape the competitive landscape, organizations must move beyond experimentation to systematic, enterprise-wide AI adoption. Our research shows that companies that embrace AI now are 30% more likely to outperform their competitors in revenue growth.

The choice facing leaders today isn't whether to embrace AI – it's how to do so in a way that builds lasting competitive advantage while empowering their workforce. Through the structured approach of AI Operations, organizations can transform the challenge of AI adoption into an opportunity for unprecedented growth and innovation.

The path forward begins with quick wins that build confidence and momentum throughout the organization. By identifying and executing these early victories, leaders create tangible examples that demonstrate AI's potential while building trust and enthusiasm among their teams. These initial successes, however modest, serve as powerful catalysts for broader adoption.

But true transformation requires more than just technological implementation. Leaders must simultaneously invest in cultural readiness, recognizing that AI adoption is as much about people as it is about technology. This means creating an environment where experimentation is encouraged and learning from failure is celebrated. It means developing clear guidelines that provide security and direction while maintaining the flexibility needed for innovation to flourish.

Most importantly, successful AI transformation hinges on understanding that technological sophistication alone doesn't guarantee success. The organizations that will thrive in the age of AI won't be those with the most advanced technology or the largest data sets – they'll be the ones that successfully combine human intelligence with artificial intelligence, creating workplaces where both can reach their full potential.

The framework is proven, the path is clear, and the time to act is now. Your workforce is ready for AI – they just need you to lead them there.

References

- [Building the AI-Powered Organization • HBR, 2019](#)
- [Future of Jobs Report 2025: 78 Million New Job Opportunities by 2030 but Urgent Upskilling Needed to Prepare Workforces • World Economic Forum, 2025](#)



Navigating AI in the Workplace: A Practical Taxonomy for Business Leaders

By David Altounian Ph.D., Vice Provost, Associate Professor of Entrepreneurship, Salve Regina University

Executive Summary

A practical understanding of the different scales of AI use types (assistive, autonomous, etc) and a taxonomy for classifying AI use in the workplace is presented as a strategic framework for C-Level executives to manage and implement AI technologies across their organizations effectively. The taxonomy outlines five distinct AI use types: Assistive, Augmentative, Semi-Autonomous, Autonomous, and Agentic AI, each with unique characteristics and applications. By understanding these categories and their appropriate use cases, executives can make informed decisions about AI adoption, mitigate risks, and maximize the benefits of AI integration, ultimately driving innovation and competitive advantage in an increasingly AI-driven business landscape.

I. Introduction

Artificial Intelligence (AI) is becoming an integral part of modern workplaces, revolutionizing processes across various departments. As organizations increasingly adopt AI technologies, it's crucial to clearly understand the different ways that AI can be used with human engagement to complete tasks. This paper aims to present a comprehensive taxonomy of AI use types and provide guidelines for their implementation across various business functions.

The importance of classifying AI use types lies in the need to help organizations effectively leverage AI across different departments in ways that are additive and helpful while minimizing errors, misuse, or unintended costs. By understanding each AI type's distinct characteristics and applications, businesses can make informed decisions about which technologies to adopt and how to integrate them into their existing workflows.

There are numerous examples of overdependence or inappropriate applications of AI leading to negative outcomes for organizations. In a highly publicized case in 2023, two lawyers and their firm were sanctioned by a judge for submitting a legal brief

that included six fictitious cases generated by ChatGPT. The lawyers claimed that they made a 'good faith mistake' by trusting the generative AI output.

A technology news site was forced to issue corrections after an article published using AI included multiple significant inaccuracies, including incorrectly calculating the earnings over time, a result that would have been very simple to verify and correct prior to publication.

Relying on assistive AI for tasks requiring creativity or nuanced judgment often leads to suboptimal outcomes. For example, using an AI tool to generate customer service scripts without human review can result in generic, impersonal responses that alienate clients. Similarly, HR teams might misuse assistive AI to draft interview questions based solely on job descriptions, leading to irrelevant or superficial queries that fail to assess candidate fit. These misuses stem from treating assistive AI as a replacement for human insight rather than a tool to support it.

Augmentative AI, when overtrusted, can amplify errors or biases. A marketing team might adopt an AI tool to auto-generate social media posts, only to discover tone-deaf or culturally insensitive content that damages brand reputation. Legal departments risk deploying AI for contract drafting without verifying compliance with evolving regulations, resulting in outdated or legally risky clauses. These failures highlight the danger of assuming augmentative AI can operate independently of human expertise, particularly in dynamic or context-sensitive domains.

Semi-autonomous systems, when deployed without adequate human oversight, can create systemic issues. For instance, a customer service chatbot designed to resolve simple inquiries might escalate complex cases to human agents but fail to flag urgent issues, leading to unresolved complaints. In healthcare, a semi-autonomous diagnostic tool might flag false positives for serious conditions, overwhelming clinicians with unnecessary follow-ups. These scenarios underscore the importance of balancing automation with human judgment—semi-autonomous AI works best when paired with clear intervention protocols and continuous feedback loops.

By misaligning AI use types with task requirements, organizations risk operational inefficiencies, reputational harm, and legal exposure. The key lies in applying AI strategically.

While terms like automation, agency, and augmentation are commonly used to describe AI applications in the workplace, they fall short of providing a comprehensive framework for managing the adoption and use of AI technologies. These concepts, while valuable, often oversimplify the complex interplay between human workers and AI systems, potentially leading to misguided implementation strategies.

A more nuanced approach is necessary to effectively integrate AI into workplace processes. This approach should consider the technical capabilities of AI systems and the specific context of each task, the potential risks and benefits, and the broader organizational and ethical implications. An AI Use Type Taxonomy is proposed which offers a more comprehensive framework that addresses these complexities. This taxonomy encourages a thoughtful evaluation of factors such as task criticality, required human judgment, regulatory requirements, and user acceptance. It also promotes a balanced view of AI integration that recognizes both the potential for enhancing human capabilities and the continued importance of human oversight and decision-making in many workplace scenarios.

Furthermore, having an understanding of the different types of AI to human engagement models facilitates a more strategic approach to AI adoption by highlighting the spectrum of AI applications available. This prevents organizations from falling into the trap of viewing AI implementation as a binary choice between full automation and no automation. Instead, it encourages a tailored approach that can evolve as AI technologies advance and as the organization's needs and capabilities change. By providing a more comprehensive framework for understanding and implementing AI in the workplace, this taxonomy enables organizations to make more informed decisions about AI adoption, leading to more effective and responsible use of these powerful technologies.

This paper will explore five main categories of AI use types: assistive AI, augmentative AI, semi-autonomous AI, autonomous AI, and agentic AI. We will examine their definitions, characteristics, and applications across marketing, finance, operations, legal, human resources, and medical fields. Additionally, we will provide guidelines for AI implementation in the workplace and discuss future trends and challenges.

II. Current Scales of AI Autonomy

Various scales have been developed to classify levels of AI autonomy across different domains. These scales can be viewed as subsets or specific applications of the broader AI Use Type Taxonomy proposed in this paper. Here, we examine four prominent scales currently in use:

The Four-Tier Autonomy Model for Human-AI Collaboration:

The Four-Tier Autonomy Model emerged from recent research into human-AI teamwork dynamics, particularly in shared workspaces where collaboration between humans and AI agents is critical. Developed in 2024 by researchers exploring situational autonomy, this framework categorizes AI systems based on their level of

initiative within collaborative teams. The model's tiers—No Autonomy, Low Autonomy, Moderate Autonomy, and High Autonomy—provide a nuanced lens for evaluating how AI agents contribute to task execution while balancing human oversight.

Currently, the model is applied in collaborative environments like logistics, customer service, and industrial automation, where human-AI teams must adapt to dynamic conditions. For instance, a Moderate Autonomy AI might propose inventory adjustments based on real-time data in warehouse management, requiring human confirmation before implementation. The model's alignment with broader AI use types—such as Assistive, Augmentative, and Semi-Autonomous AI—highlights its practical utility in designing workflows that balance automation with human judgment.

The model's source, a 2024 ACM paper titled "*An Evaluation of Situational Autonomy for Human-AI Collaboration*," underscores its academic roots and relevance to modern workplace challenges.

The SAE Autonomous Vehicle Scale:

SAE International, the association of engineering and technology specialists in the field of mobility, introduced the SAE Autonomous Vehicle Scale, formally defined in the SAE J3016 standard (first published in 2014), which has become a framework for standardizing autonomous driving and operation capabilities terminology. This six-level scale, developed by SAE International, a global engineering organization, addressed the growing need for clarity as automakers and tech companies began experimenting with advanced driver-assistance systems (ADAS) and fully autonomous vehicles. At the time, terms like "self-driving" and "autonomous" were used inconsistently, leading to confusion among consumers, regulators, and engineers. The SAE scale provided a structured way to classify vehicles based on their ability to perform dynamic driving tasks, ranging from Level 0 (No Driving Automation) to Level 5 (Full Driving Automation).

Historically, the scale was pivotal in aligning industry practices. Before its introduction, automakers like Tesla and Audi marketed systems with ambiguous names (e.g., "Autopilot" or "Traffic Jam Assist"), which often misled consumers about the level of autonomy. The SAE framework clarified these distinctions, ensuring that Level 2 systems (e.g., Tesla's Autopilot) were recognized as partial automation requiring constant driver oversight, while Level 3 systems (e.g., Audi's Traffic Jam Assist) allowed conditional autonomy under specific conditions. By 2021, SAE refined the terminology further, distinguishing driver controlled driving systems (Levels 0–2) from automated driving systems (Levels 3–5), emphasizing the shift from human-centric to machine-centric control.

Today, the SAE scale is globally adopted by automakers, regulators, and safety organizations. For example, Level 4 autonomous taxis like Waymo's services in Phoenix operate without human drivers in predefined zones, while Level 3 systems (e.g., Honda's Traffic Jam Assist) enable hands-free driving in traffic jams but require driver readiness to intervene.

The scale's alignment with broader AI use types—such as Assistive, Augmentative, Semi-Autonomous, and Autonomous AI—highlights its versatility:

- **Level 1 (Driver Assistance):** Requires constant driver oversight
- **Level 2 (Partial Driving Automation):** Requires constant driver oversight
- **Level 3 (Conditional Driving Automation):** Autonomous under certain conditions
- **Level 4 (High Driving Automation):** Autonomous AI with human control capability
- **Level 5 (Full Driving Automation):** Fully autonomous with no human control

This scale provides a more granular breakdown within the automotive context but still aligns with the broader categories of our taxonomy.

The Graded Medical AI Autonomy Model from *The Lancet*:

The Graded Medical AI Autonomy Model, first articulated in a 2020 *Lancet Digital Health* article titled "Approaching autonomy in medical artificial intelligence," emerged as a response to the healthcare sector's need for a nuanced framework to classify AI systems. At the time, medical AI was rapidly advancing, yet the industry lacked standardized terminology to describe varying levels of autonomy, leading to confusion about risks, liabilities, and clinical applications. The model's developers—clinicians and researchers—recognized that a binary distinction between "assistive" (non-autonomous) and "fully autonomous" AI systems failed to capture the spectrum of real-world scenarios, particularly in high-stakes environments like diagnostics or treatment planning.

The model drew inspiration from analogous frameworks in transportation and aviation, where autonomy levels (e.g., the U.S. Department of Transportation's five-tier scale) had already proven critical for risk stratification. By adapting this concept to healthcare, the *Lancet* model introduced three key determinants of autonomy-related risk: (1) the responsible agent for medical event monitoring, (2) the availability of human backup decision-makers, and (3) the specificity of the medical domain and population. This approach acknowledged that AI autonomy is not absolute but

context-dependent—systems designed for narrow, well-defined tasks (e.g., diabetic retinopathy screening) could operate with higher autonomy than those tackling complex, variable conditions.

Today, the model is widely referenced in medical AI research and regulatory discussions. For example, assistive AI tools like radiology analysis software provide clinicians with diagnostic suggestions but require human interpretation, aligning with the model's emphasis on clinician responsibility. Semi-autonomous systems—such as AI-driven patient monitoring tools—flag anomalies for human review, balancing automation with oversight. Fully autonomous systems remain rare but are being explored in controlled settings, such as automated insulin dosing for type 1 diabetes, where predefined protocols minimize risks.

Emphasizing the interplay between AI autonomy and human responsibility provides a flexible framework for healthcare organizations to adopt AI technologies while prioritizing patient safety and ethical practice.

The Legal Liability Autonomy Scale:

The Legal Liability Autonomy Scale emerged in 2024 as a response to the growing complexity of assigning responsibility for AI-driven harm. Proposed in a peer-reviewed paper titled *Levels of Autonomy: Liability in the age of AI Agents* (OpenReview, 2024), this five-tier framework draws inspiration from governance models for autonomous vehicles, such as the UK's Automated Vehicles Act 2024. The scale addresses a critical gap in liability law by categorizing AI systems based on their decision-making autonomy, thereby guiding legal discussions about accountability when AI causes harm.

The 2024 paper recognized that traditional tort law—rooted in human agency—struggled to address AI systems that learn and adapt beyond their initial programming. By introducing a tiered autonomy scale, the model provides a structured way to allocate liability, mirroring how vehicle autonomy levels (e.g., SAE J3016) inform regulatory and legal frameworks.

Currently, the scale is applied in legal and regulatory contexts to assess liability risks. For example, Level 5 (Full Autonomy) aligns with systems like autonomous taxis, where liability shifts to operators under strict regimes, as seen in the UK's Act. Level 3 (Multi-Step Planning) applies to semi-autonomous tools in healthcare or finance, where AI makes decisions but requires human oversight. The scale's tiers—ranging from No Autonomy (human-controlled tools) to Full Autonomy (self-modifying systems)—help stakeholders determine whether liability falls on developers, operators, or users.

As AI adoption expands, the Legal Liability Autonomy Scale is a foundational tool for legal and ethical discussions, ensuring that liability frameworks evolve alongside

technological advancements. Its integration into regulatory debates underscores the need for clarity in an era where AI agents increasingly operate with minimal human oversight.

The NIST AI Use Taxonomy: A Human-Centered Approach

In 2024, the National Institute of Standards and Technology (NIST) introduced the AI Use Taxonomy: A Human-Centered Approach to provide a structured method for classifying AI implementations. Recognizing the rapid growth of AI across sectors like finance, healthcare, and manufacturing, NIST sought to establish a framework that fosters trustworthy and responsible AI development. Historically, traditional Human-Computer Interaction (HCI) models focused on direct human control, but the advent of AI necessitates a paradigm shift.

The NIST AI Use Taxonomy centers around sixteen key "human-AI activities," each detailing how an AI system augments or replaces human effort. Activities range from "Content Creation," where AI generates new artifacts like video or software code, to "Vehicular Automation," automating physical transportation. Examples of this can be seen with tools for subtitle creation, text-to-image software, self-driving cars, trucks, and trains. While these activities are defined distinctly, NIST notes that real-world AI applications often involve a combination of activities. The AI Use Taxonomy also emphasizes that humans, AI systems, or both can perform each of the activities. This model's primary goal is to provide a common terminology and understanding for human-AI interactions, independent of AI techniques or domains. The taxonomy helps organizations navigate the complexities of human-AI teaming by focusing on the interaction's intended outcome and how the AI system contributes to achieving it.

It's intended for use by those involved in AI design, development, deployment, evaluation, human factors, and governance. By classifying AI applications into specific activities, the taxonomy enables cross-domain insights, highlights common measurement needs, and facilitates the development of use cases.

While valuable for framing human-AI interaction outcomes, the NIST AI Use Taxonomy has notable limitations for managers seeking to assess AI autonomy in deployment. Its primary focus on human-AI activities (e.g., content creation, vehicular automation) rather than autonomy levels leaves critical gaps for decision-makers. For instance, the taxonomy does not explicitly classify AI systems by their capacity for independent decision-making, a key factor in evaluating risks, compliance, and operational feasibility. Managers relying solely on this framework may struggle to distinguish between assistive AI (e.g., data analysis tools) and semi-autonomous AI (e.g., automated customer service chatbots), which require varying levels of human oversight. This ambiguity could lead to misaligned expectations about AI capabilities, such as deploying fully autonomous systems in contexts where human intervention is legally or ethically necessary.

These scales demonstrate that there is a need for a comprehensive framework encompassing various domain-specific autonomy classifications. While each scale

focuses on particular aspects relevant to its field (e.g., driver involvement in automotive, liability in legal contexts), they all reflect a progression from human-controlled to increasingly autonomous systems.

Our taxonomy provides a unifying structure that can be applied across industries, allowing for a more standardized approach to discussing and implementing AI systems. It accommodates the nuances of different domains while offering a common language for cross-sector comparisons and policy development.

As AI technology evolves, these scales may need to be updated. However, the fundamental categories of our taxonomy—Assistive, Augmentative, Semi-Autonomous, Autonomous, and Agentic AI—provide a flexible framework that can adapt to future developments in AI capabilities and applications.

III. A Taxonomy Scale for AI Use in the Workplace

In order to effectively manage and deploy AI in the workplace, it's helpful to categorize AI systems based on the degree of human interaction required, focusing on the balance between automated capabilities and human oversight. This taxonomy proposes a tiered approach, encompassing five distinct AI use types: Assistive AI, which provides information and support; Augmentative AI, which enhances human expertise with insights and suggestions; Semi-Autonomous AI, which performs tasks independently but requires human intervention for critical decisions; Autonomous AI, which operates entirely without human input; and Agentic AI, which exhibits goal-directed autonomy and integrates diverse tools.

Assistive AI: Empowering Human Decision-Making Through Information and Task Support

Assistive AI refers to systems designed to enhance human capabilities by providing information, streamlining processes, or automating simple tasks. These tools are characterized by their supportive role, never making independent decisions but rather empowering users with the insights or functionality needed to make informed choices and execute tasks more efficiently. From a historical perspective, assistive AI builds upon earlier iterations of expert systems and decision support systems but leverages advancements in machine learning and natural language processing to offer more intuitive and adaptive assistance.

Assistive AI finds broad application in various sectors. In healthcare, for instance, diagnostic tools analyze patient data to highlight potential concerns, but the physician ultimately synthesizes the information and makes the final diagnosis. Similarly, in customer service, AI-powered chatbots handle routine inquiries and provide support documentation, freeing human agents to address complex issues. Consider a marketing manager researching potential campaign strategies: they

might employ an assistive AI tool to identify relevant market data and summarize customer feedback, which informs their strategic decisions. Engineers might use an AI-powered code completion tool to accelerate the writing of code, but they still need to determine what logic the code will incorporate. This is different from an AI code generator that attempts to write larger sections of code from a natural language command from the engineer. In both instances, AI serves as a catalyst for human expertise, providing insights and capabilities that augment human capabilities, leading to more successful project outcomes.

Augmentative AI: Enhancing Human Expertise Through Collaborative Insights

Augmentative AI represents a transformative approach to human-AI collaboration, designed to amplify human capabilities by providing actionable insights, suggestions, or complex analyses that complement professional expertise. Unlike autonomous systems, these tools operate as partners, offering context-aware recommendations while leaving final decision-making authority to humans. Historically rooted in early decision-support systems, modern augmentative AI leverages advancements in machine learning and natural language processing to deliver real-time, adaptive assistance. For example, marketing professionals use tools like Perplexity AI to generate tailored social media post ideas aligned with campaign goals, such as boosting engagement or driving conversions. The AI analyzes trends, audience preferences, and past performance data to suggest content strategies, but the marketer retains control over creative direction and execution.

In finance, augmentative AI tools like BlackRock's Aladdin analyze vast market datasets to identify investment opportunities or risks, offering portfolio managers data-driven recommendations. These systems process complex financial models and historical trends to highlight potential strategies, enabling professionals to make informed decisions without replacing their judgment. Similarly, in healthcare, AI-powered diagnostic assistants like IBM Watson Health provide clinicians with evidence-based treatment options by synthesizing medical literature and patient data. While the AI identifies patterns and potential diagnoses, doctors evaluate the suggestions against their clinical experience to finalize care plans. These examples illustrate how augmentative AI bridges the gap between data analysis and human expertise, enhancing productivity without compromising accountability. By maintaining human oversight, such systems ensure that nuanced, context-dependent decisions remain in the hands of professionals.

Semi-Autonomous AI: Balancing Automation with Human Oversight for Enhanced Decision-Making

Semi-autonomous AI represents a middle ground between full automation and purely human-driven processes, offering a balance of AI efficiency and human judgment for critical decisions. These systems can perform tasks with some degree of independence, but they are designed to require human oversight or intervention in situations that demand nuanced understanding, ethical considerations, or unpredictable environmental adaptation. Semi-autonomous AI is often deployed in complex environments where the cost of errors is high, or where ethical considerations require a human in the loop.

Examples of semi-autonomous AI systems can be found in both finance and healthcare. In finance, fraud detection systems employ AI to analyze transactions and flag suspicious activities, but human investigators review these flagged transactions before taking any actions, such as freezing an account. This hybrid approach leverages AI's ability to quickly analyze vast amounts of data while relying on human expertise to discern legitimate anomalies from fraudulent behavior. In healthcare, AI-powered diagnostic tools can analyze medical images to detect potential issues, such as tumors or fractures. However, radiologists must review the AI's findings and incorporate them into a broader assessment of the patient's condition before making a final diagnosis and treatment plan. This semi-autonomous model allows for greater efficiency in the diagnostic process while ensuring that a human expert is ultimately responsible for critical decisions.

Autonomous AI: Closed-Loop Decision-Making Without Human Intervention

Autonomous AI represents the pinnacle of machine independence, operating entirely without human input across all phases of a task—from perception and analysis to decision-making and execution. These systems employ closed-loop decision-making, where environmental feedback directly informs subsequent actions, enabling continuous adaptation without external guidance. Unlike semi-autonomous or augmentative AI, autonomous systems are self-contained, relying on preprogrammed logic and real-time data processing to navigate complex scenarios. For example, Waymo's fully autonomous taxis exemplify this paradigm, navigating urban environments, avoiding obstacles, and transporting passengers without human drivers. The AI processes sensor data, adjusts routes dynamically, and executes maneuvers autonomously, demonstrating how closed-loop systems can achieve human-like autonomy in structured domains.

Autonomous AI is also deployed in industrial and logistics settings, where precision and consistency are paramount. Amazon's Kiva warehouse robots autonomously sort

and transport inventory, optimizing storage layouts and order fulfillment without human oversight. These systems use machine learning to adapt to changing inventory levels and operational demands, ensuring efficient workflow management. Similarly, high-frequency trading algorithms execute trades autonomously based on predefined market conditions, leveraging real-time data to maximize profitability. While autonomous AI excels in repetitive, rule-bound tasks, its limitations emerge in unpredictable or ethically complex scenarios. For instance, fully autonomous weapons systems face intense scrutiny due to the irreversible nature of their decisions, highlighting the need for ethical safeguards even in highly automated domains. By design, autonomous AI prioritizes efficiency and reliability but requires rigorous testing and fail-safe protocols to mitigate risks in high-stakes applications.

Agentic AI: Autonomous, Goal-Driven Action with Tool Integration

Agentic AI represents a paradigm shift in artificial intelligence, moving beyond mere task execution to exhibiting goal-directed autonomy, the capacity to independently define, plan, and execute complex projects to achieve specific objectives. A key feature of agentic AI is its ability to not only operate independently but also to leverage tools and data sources outside of its initial training data to achieve those defined goals. In practice, this means that agentic AI systems can not only perform assigned tasks but also determine the steps, data, and tooling necessary to see a project through to completion without human intervention. To achieve these outcomes, agentic AI incorporates mechanisms for self-monitoring, feedback loops, and adaptation, so that it can learn from its successes and failures to continuously improve its ability to achieve the desired outcomes.

Agentic AI systems are beginning to be deployed in a number of different domains. For example, in cybersecurity, agentic AI systems might be used to autonomously identify and remediate security vulnerabilities, leveraging a variety of security tools, threat intelligence feeds, and attack simulation platforms to fortify an organization's defenses. Similarly, in sales and marketing, agentic AI systems can manage and optimize marketing campaigns, leveraging customer data, marketing automation platforms, and A/B testing to increase customer acquisition and revenue. The true power of these AI systems is their ability to define tasks, execute them, and adapt to evolving scenarios so that they can achieve complex goals without human intervention.

IV. Usage Type Decision Factors

Determining the appropriate AI Use Type for workplace tasks is crucial for maximizing benefits while minimizing risks. The choice between Assistive, Augmentative, Semi-Autonomous, Autonomous, or Agentic AI should be based on several key factors:

1. **Task Complexity and Variability:** Simple, repetitive tasks are better suited for Autonomous AI, while complex tasks requiring nuanced understanding benefit from Assistive or Augmentative approaches. For instance, data entry or inventory management can be automated, but customer communications often need human oversight.
2. **Potential Impact and Risk:** Tasks with high-stakes outcomes or significant potential for error should use Assistive or Augmented AI. Using fully Autonomous AI for critical decision-making or customer-facing communications can lead to reputational damage or legal issues.
3. **Need for Human Judgment:** Tasks requiring empathy, creativity, or ethical considerations are best suited for Assistive or Augmentative AI. For example, HR support can benefit from AI-driven personalization, but final decisions should involve human judgment.
4. **Data Availability and Quality:** Autonomous AI requires extensive, high-quality data for accurate performance. Tasks with limited or unreliable data are better suited for Assistive or Augmentative approaches that incorporate human expertise.
5. **Regulatory and Compliance Requirements:** Industries with strict regulations may require human oversight, making Assistive or Semi-Autonomous AI more appropriate. This is particularly relevant in fields like healthcare, finance, and legal services.
6. **User Acceptance and Trust:** Consider the comfort level of both employees and customers with AI-driven processes. Assistive or Augmentative AI can help build trust by maintaining human involvement.

Additional criteria for selecting the appropriate AI Use Type include:

1. **Task Criticality:** Higher criticality suggests more human involvement (Assistive or Augmentative).
2. **Required Response Time:** Tasks needing immediate action may benefit from Autonomous AI if properly trained.
3. **Consistency Requirements:** Highly standardized processes are suitable for Autonomous AI, while variable tasks benefit from human input.
4. **Learning and Adaptation Needs:** Dynamic environments may require Augmentative or Semi-Autonomous AI that can learn from human feedback.
5. **Transparency and Explainability:** When decision rationale is crucial, Assistive or Augmentative AI allows for human interpretation and explanation.

By carefully considering these factors and criteria, organizations can avoid the pitfalls of inappropriately applying fully Autonomous AI to tasks better suited for human-AI collaboration. This approach ensures that AI enhances rather than replaces human capabilities in the workplace, leading to more effective and responsible AI integration.

| Criteria | Assistive AI | Augmentative AI | Semi-Autonomous AI | Autonomous AI | Agentic AI |
|--|--------------------------------------|-------------------------------------|---------------------------------------|---|--|
| Task Complexity & Variability | Complex, nuanced tasks | Moderate complexity, variable tasks | Moderate to complex, some variability | Simple, repetitive tasks | Complex, dynamic, multi-step tasks |
| Potential Impact & Risk | Low to moderate risk | Moderate risk | Moderate to high risk | High risk (requires safeguards) | Very high risk (requires strong oversight) |
| Need for Human Judgment | High need for human judgment | Moderate need for human judgment | Limited need for human judgment | Minimal need for human judgment | Adaptive human oversight |
| Data Availability & Quality | Limited/unreliable data acceptable | Moderate data quality required | High data quality required | Extensive, high-quality data needed | Real-time data and tool integration |
| Regulatory Compliance | Highly suitable in regulated sectors | Suitable in regulated sectors | Conditional suitability | Limited use in regulated sectors | Significant regulatory concerns |
| User Acceptance & Trust | Highest user acceptance | High user acceptance | Moderate user acceptance | Lower user acceptance | Requires substantial trust building |
| Task Criticality | High criticality, human involvement | High criticality, human involvement | Variable criticality | Low criticality (typically) | High to variable criticality |
| Required Response Time | Flexible | Flexible | Fast, but can pause | Immediate action needed | Real-time or scheduled actions |
| Consistency Requirements | Variable | Variable to moderate | Moderate to high | Highly standardized processes | Adaptive consistency requirements |
| Learning & Adaptation | Limited | Moderate, human feedback needed | High, feedback-driven | Limited to initial training data | Continual, environment-driven adaptation |
| Transparency & Explainability | High importance | Moderate importance | Limited importance | Minimal importance (limited human engagement) | Explainability increasingly important |

V. Examples of Use Cases

The following table presents examples of how Assistive, Augmented, Semi-Autonomous, Autonomous, and Agentic AI can be leveraged across various functional areas and industry segments to illustrate the practical application of the AI Use Type Taxonomy.

| | | |
|--|--|--|
| Marketing <ul style="list-style-type: none">● Assistive AI: Content research tools that help marketers find relevant information for campaigns.● Augmentative AI: Social media post generation systems that suggest content based on campaign goals.● Semi-Autonomous AI: Customer segmentation and targeting algorithms that analyze data and propose strategies for human approval. | Finance <ul style="list-style-type: none">● Augmentative AI: Financial forecasting and risk assessment tools that provide insights to human analysts.● Semi-Autonomous AI: Fraud detection systems that flag suspicious transactions for human review.● Autonomous AI: High-frequency trading algorithms that execute trades based on market conditions without human intervention. | Operations <ul style="list-style-type: none">● Assistive AI: Inventory management assistants that provide real-time stock information to human managers.● Augmentative AI: Supply chain optimization tools that suggest improvements based on complex data analysis.● Autonomous AI: Warehouse robotics systems that independently manage inventory movement and organization. |
| Legal <ul style="list-style-type: none">● Assistive AI: Legal research and case law analysis tools help lawyers find relevant precedents.● Augmentative AI: Contract review and drafting assistance systems that recommend enhancements and identify potential issues.● Semi-Autonomous AI: Compliance monitoring systems that track regulatory changes and alert human experts to necessary actions. | Human Resources <ul style="list-style-type: none">● Assistive AI: Candidate sourcing tools that help recruiters find potential hires based on specified criteria.● Augmentative AI: Resume screening and initial interview scheduling systems that support human decision-making in the hiring process.● Semi-Autonomous AI: Employee engagement analysis tools that provide insights on workforce sentiment. | Medical <ul style="list-style-type: none">● Assistive AI: Medical literature search and summarization tools .● Augmentative AI: Diagnostic support systems that analyze patient data and suggest potential diagnoses for physician consideration.● Semi-Autonomous AI: Patient monitoring and alert systems that track vital signs and notify human staff of significant changes. |

VI. Future Trends and Challenges

As AI technology advances, it's natural to anticipate changes in how we categorize and understand its applications, demanding an evolution of the AI use type taxonomy. One key trend will be the seamless integration of AI across multiple use types within single systems, blurring the lines between traditionally distinct categories. For instance, a customer service platform may combine assistive AI to gather information, augmentative AI to suggest solutions, and semi-autonomous AI to resolve routine inquiries, creating a unified, multifaceted system. Additionally, there will be a growing focus on explainable AI (XAI) to enhance transparency and trust, enabling users to understand how AI systems arrive at decisions. We'll also see the development of more sophisticated agentic AI systems capable of handling complex, multi-step tasks, exhibiting greater autonomy and adaptability.

This evolution brings significant challenges that organizations must address proactively. Ethical considerations surrounding AI decision-making and potential biases must be paramount, ensuring fairness and accountability. Finding the right balance between automation and human oversight will also be crucial, leveraging AI's efficiency while preserving human judgment in critical areas. Concerns about job displacement and the changing nature of work must be tackled through upskilling initiatives and thoughtful workforce planning. Ensuring data privacy and security in AI-driven systems is paramount, requiring robust safeguards and adherence to evolving regulations. By navigating these challenges thoughtfully, organizations can harness AI's transformative potential while mitigating its risks, fostering a future where humans and AI collaborate effectively and ethically.

VII. Conclusion

The taxonomy of AI use types presented in this paper provides a framework for understanding and implementing AI across various business functions. By categorizing AI systems as assistive, augmentative, semi-autonomous, autonomous, or agentic, organizations can better align their AI strategies with specific departmental needs and overall business objectives.

As AI continues to transform the workplace, it is crucial for organizations to develop comprehensive AI integration strategies that consider not only the technological capabilities but also the ethical, legal, and human factors involved. By following the guidelines outlined in this paper and staying informed about emerging trends and challenges, businesses can harness the power of AI to enhance productivity, innovation, and employee satisfaction while mitigating potential risks.

The future of work will undoubtedly be shaped by the continued evolution and integration of AI technologies. Organizations that successfully navigate this landscape by thoughtfully implementing appropriate AI use types across their

operations will be well-positioned to thrive in an increasingly AI-driven business environment.

A final note: this paper was written using augmented and assisted AI. The research and section development was done using Perplexity to augment my literature review work and helping with writing sections. Grammarly was used to assist with grammar, tone, and spelling. The use of Perplexity helped uncover additional elements and scales that were not known prior to the use of the tool.

References:

1. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.200-1.pdf>
2. https://www.ey.com/en_gl/ai/navigating-gen-ai-a-six-category-taxonomy-for-strategic-use-cases
3. <https://dl.acm.org/doi/10.1145/3613904.3642564>
4. <https://www.faistgroup.com/news/autonomous-vehicles-levels/>
5. <https://openreview.net/forum?id=EH6SmoChx9>
6. <https://www.reworked.co/digital-workplace/choosing-the-right-ai-for-your-business-goals/>
7. <https://www.brookings.edu/articles/generative-ai-the-american-worker-and-the-future-of-work/>
8. <https://www.alexandrasamuel.com/ai/how-to-pick-the-right-ai-for-each-task>
9. <https://www.linkedin.com/pulse/generative-ai-workplace-choosing-right-model-future-khalid-turk-icewc>
10. <https://www.ibm.com/think/topics/ai-in-the-workplace>



AI Is Your New Coworker— But Is It Taxing Your Culture?

By Ethan Burris, Ph.D., Senior Associate Dean, McCombs School, The University of Texas at Austin

Imagine this: You log into your company’s performance portal one morning, and instead of a traditional “Good job!” from your manager, you see a concise message—“Our algorithm indicates that your performance this quarter exceeds expectations by 12%. Congrats!” Your heart skips a beat. Is this praise genuine recognition or a cold statistical judgment rendered by a faceless machine? In today’s high-speed digital workplace, artificial intelligence is not only streamlining processes—it’s fundamentally reshaping our organizational culture. Sometimes for better. Sometimes for worse.

In organizations across the globe, AI is rapidly taking over roles once reserved for humans: from screening resumes and scheduling meetings to providing performance feedback and even shaping team dynamics. But as we race toward a future where algorithms determine promotions and assign team members, a critical question emerges: What happens to the human touch in our work lives? More specifically, how do employees feel when a machine is heavily involved in determining life-changing evaluations, or when an unseen algorithm decides whether they land a job or get passed over?

Although much of the conversation around AI centers on driving efficiency, increasing productivity, and utilizing a thought partner to augment strategic decision making, one externality is how the increased use of machines for what were human interactions might inadvertently foster isolation, favoritism, and a breakdown of trust and human connection. In short, I argue that we should not only consider the technological impact but also the very real emotional and psychological responses of employees who find themselves on the receiving end of AI-driven processes.

I. The Changing Face of Organizational Culture in the Age of AI

Redefining “Culture” in a Digital Era

Organizational culture has long been described as the collective values, beliefs, and behaviors that shape how work gets done within a company. Traditionally, this culture was nurtured in face-to-face interactions, water-cooler chats, and shared

rituals. Today, however, AI is rewriting that script. When a machine is responsible for everything from candidate screening to performance management, culture isn't just what you see in the break room—it's also encoded in lines of code.

Consider the classic definition: culture is “how we do things around here.” But what happens when “things” are determined by algorithms? When the decisions that affect our careers are made not by a trusted manager, but by a black-box system, the essence of “how we do things” can shift dramatically. As noted by Adrian Ward in his research on technology's impact on human cognition, our reliance on digital tools can both free us from menial tasks and gradually erode the informal, personal connections that underpin a vibrant organizational culture.

The Emotional Dimension: How It Feels to Be Evaluated by a Machine

For many employees, the idea of being judged by an algorithm can be unsettling. It raises questions about fairness, transparency, and the role of empathy in the workplace. Studies on virtual communication reveal that when feedback is delivered through impersonal, automated channels, employees are less likely to engage in constructive dialogue. And this communication is still delivered by people! The absence of a human touch in utilizing algorithms not only diminishes the warmth of the praise but can also degrade the sense of community that culture is supposed to create.

Imagine a scenario where you are working on the early stages of a project. Before, you might have brainstormed some ideas, tapped some trusted colleagues to offer early feedback, gathered a meeting to beat up the ideas, and finally settle on a path forward. Now, it can often feel safer to test preliminary ideas with an AI. You don't face the judgment of colleagues who may criticize less baked ideas. Yet, you also don't develop the trust of colleagues by presenting less polished work and offering a chance to co-create something special.

This short example offers one of both the promise and the perils of using AI for tasks that have historically been accomplished in teams. But there are many others.

II. The Upside of AI in the Workplace: Efficiency, Innovation, and Fairness

1. Efficiency & Automation: Liberating the Human Mind

One of the most touted benefits of AI is its capacity to automate repetitive tasks, liberating employees from administrative drudgery. Consider the process of resume

screening: traditionally, HR professionals had to sift through hundreds—if not thousands—of applications to identify suitable candidates. Then they would get together in hiring committees to advance a smaller number forward for interviews. With AI, this process is accelerated, allowing recruiters to focus on strategic, important elements of resumes that lead to better quality hires.

In many organizations, AI-driven tools are already taking on tasks like scheduling meetings, managing emails, and even drafting preliminary performance reviews. This automation not only boosts efficiency but also opens up time for employees to engage in more creative, value-added work. For instance, managers can now dedicate more time to mentoring and coaching, rather than drowning in data analytics.

However, as liberating as automation can be, it also sets the stage for a paradox. When machines take over routine functions, the expectation is that employees will have more “human” time to innovate and collaborate. But if those very interactions are increasingly mediated by technology, the opportunities for spontaneous, creative exchanges may diminish over time.

2. Data-Driven Fairness: The Promise of Objectivity

A significant promise of AI is its potential to reduce human biases in decision-making. In theory, algorithms—if designed properly—can evaluate candidates, performance, and team dynamics based solely on quantifiable data. For example, in the recruiting process, AI can anonymize resumes, stripping away demographic information such as gender, ethnicity, or even age, to help ensure that only the candidate’s skills and experiences are considered. Moreover, they can also strip out other sets of biases about successful leadership traits like extraversion, past performance as an individual contributor, or whether leadership can be learned.

Research suggests that when AI systems are fed diverse and representative data sets, they can promote fairness by eliminating unconscious human biases. In a well-designed system, every candidate gets a chance, and performance evaluations become a reflection of measurable contributions rather than subjective impressions. Yet the reality is that most organizations do not have access to these kinds of representative data sets, nor do they devote resources to evaluate how representative they are.

3. Personalization at Scale: Tailored Employee Experiences

AI’s capacity for data analysis also enables a more personalized approach to employee development. Imagine an AI tool that continuously monitors your work habits, identifies skill gaps, and then recommends targeted learning modules or

stretch assignments. Such systems can provide employees with real-time feedback and development opportunities tailored to their individual needs.

Personalization isn't limited to learning and development. AI can also be employed to enhance well-being by analyzing patterns of stress, burnout, or even disengagement. Early-warning systems can flag potential issues, allowing managers to better coach employees before problems escalate. In an ideal world, this leads to a more supportive work environment where employees feel seen, understood, and empowered to grow.

4. Innovation Through Enhanced Collaboration

When AI handles the mundane, employees are freed to focus on innovation and creative problem-solving. This can lead to the formation of high-performing teams where each member contributes their unique strengths. By using AI to analyze team compositions—identifying complementary skills, cognitive styles, and even interpersonal dynamics—organizations can form teams that are primed for success.

Yet, innovation is not just about individual brilliance. It's about collective synergy. AI can support this by enabling new forms of collaboration, such as real-time language translation in global teams or sentiment analysis that highlights emerging issues before they become conflicts. The result can be a more agile, responsive organization that's better equipped to navigate a rapidly changing business landscape.

III. The Downside: Isolation, Bias, and the Erosion of Trust

1. Isolation: When Efficiency Comes at a Cost

For all its efficiencies, AI can also lead to a disconcerting loss of the human element. When a machine becomes the primary conduit for hiring decisions, informal feedback or performance evaluations, employees can feel reduced to mere numbers. The warmth and nuance of human judgment—elements that foster trust and connection—are hard to replicate in an algorithm.

Research on technology's influence on human interaction warns us about the “technological distancing” effect: as more processes become automated, the subtle cues and interpersonal bonds that form the backbone of effective teams begin to dissipate. Employees may find themselves feeling like they have more feedback and yet are more isolated than ever, not only from their peers but also from the very processes that determine their career progression.

2. Overreliance on the Machine: The Risk of “Algorithm Dependence”

There is a fine line between leveraging AI as a tool and becoming overly dependent on it. Research highlights a concerning trend: as organizations lean more heavily on algorithmic decision-making, they risk diluting human judgment. When managers and employees begin to defer entirely to AI recommendations, the nuanced understanding that comes from human experience and intuition may be lost. In short, algorithms still need to be checked and have active engagement from its users.

This overreliance on technology can lead to a “black-box” mentality, where decisions are accepted with less scrutiny, less debate among team members with differing perspectives, and less opportunity to evolve understanding of the problem. In such cases, the use of technology can speed decision making, but in a process that should ideally be collaborative, the result is an erosion of trust.

3. The Loss of Spontaneity: Killing Creativity with Algorithms

One of the most underappreciated aspects of a vibrant organizational culture is the spontaneous exchange of ideas—the offhand comments, the impromptu brainstorming sessions, the “half-baked” ideas that, when shared, spark innovation. Yet, when employees are increasingly relying on AI to refine their ideas before presenting them to peers, that raw, unfiltered creativity can be lost.

In an environment where every thought is first filtered through an algorithm, the fear of judgment can stifle the very experimentation that drives innovation. Employees might hesitate to share nascent ideas for fear of being “graded” by an AI system, or worse by their peers. This can lead to a culture of perfectionism that ultimately undermines growth and collaboration.

4. Trust & Privacy: The “Big Brother” Effect

Another significant downside of integrating AI into everyday work processes is the potential for over-surveillance. When AI is used to monitor productivity, analyze sentiment, or even track employee interactions, it can create a pervasive sense of being watched—akin to a “Big Brother” scenario. This constant monitoring can lead to heightened anxiety, reduced openness, and a chilling effect on creativity.

Employees who feel that every keystroke, email, or chat message is being scrutinized are less likely to engage in honest, spontaneous interactions. Over time, this erosion of privacy not only damages trust but can also lead to a culture of self-censorship,

where employees are more concerned with “playing it safe” than taking the risks necessary for innovation.

IV. AI and the HR Life Cycle: A Detailed Analysis

The impact of AI – both positive and negative – on organizational culture is most pronounced when examined through its role throughout the HR life cycle. Let’s break down how AI shapes—and sometimes warps—each phase from recruiting and hiring to daily collaboration.

A. Recruiting and Hiring: Objectivity vs. Alienation

The Promise: Speed and Standardization

In the recruiting process, AI tools can scan thousands of applications in a fraction of the time it would take a human recruiter. By using natural language processing and predictive analytics, these systems can quickly identify candidates whose skills and experiences align with the job requirements. Moreover, anonymizing applications can, in theory, strip away markers of identity that might lead to bias.

For candidates, this offers the promise of a more standardized, fair process. Knowing that the same set of standards are applied to weigh qualifications objectively can be reassuring. A fair process, after all, is one where everyone is evaluated by the same yardstick.

The Downside: The “Black-Box” Effect

However, candidates often react negatively when they learn that a machine, rather than a human, made the initial decision about their future. By definition, the algorithm can only pick up on so many important nuances, and humans might differentially weight important criteria that surface during the recruitment process. The lack of human touch can create a sense of disenchantment with the entire recruiting process. Applicants may wonder: Was I rejected because of an arbitrary cutoff? Could a human have seen potential where the algorithm did not? Given that upwards of 30% of applicants reapply for different roles within the same company, repeated rejections could only enhance this disillusionment.

For those who do advance in the hiring process, the lingering knowledge that an AI was responsible for screening their application may color their entire perception of the company. Trust is built on transparency and empathy—qualities that are hard to convey through lines of code.

B. Performance Feedback: When Metrics Replace Mentorship

The Upside: Continuous Improvement

One of the advantages of integrating AI into performance management is the ability to provide real-time feedback. Instead of waiting for an annual review, employees might receive continuous updates on their performance metrics through dashboards, customized AI generated messages, and coaching feedback for improvement. This immediacy can help individuals adjust course quickly, learn from their mistakes, and seize opportunities for success.

For managers, these systems can serve as an early warning mechanism, flagging potential issues before they escalate into major problems. The idea is that data-driven insights will lead to more objective, constructive feedback sessions that help employees grow. Platforms now exist for managers to become better mentors through customized reports for each of their employees, including language to use to better coach performance.

The Downside: The Cold Calculus of a Machine

Yet, when the primary feedback comes from an algorithm, the human element of empathy and encouragement can be lost. Receiving a “performance score” without the context of a conversation or the warmth of human recognition can leave employees feeling isolated and demoralized. Even constructive criticism, when delivered by a machine, may lack the nuance needed to motivate and guide.

Employees have expressed that knowing they are being “judged” by an AI can inhibit follow-up discussions. The opportunity to ask for clarification, share context, or simply receive a supportive pat on the back is diminished when the feedback is impersonal. Over time, this can lead to a culture where employees feel their efforts are reduced to mere numbers on a dashboard.

C. Team Composition: Assembling the Dream Team—or Not

The Promise: Data-Driven Team Formation

AI can analyze vast amounts of data—from skills and experience to personality traits and work habits—to suggest optimal team compositions. By identifying complementary strengths, these systems can theoretically form teams that are more innovative and effective. For example, pairing a creative thinker with someone who excels at operational execution might produce a synergy that neither could achieve alone.

In a perfect scenario, employees would appreciate that the team they're placed in is chosen based on objective data rather than subjective managerial whims. It could also mitigate favoritism, ensuring that every employee has a fair shot at being part of a high-performing unit.

The Downside: The Loss of Autonomy

Despite these advantages, many employees feel that being “assigned” to a team by an algorithm diminishes their sense of agency. When team formation is reduced to data points and predictive models, individuals may worry that their unique, intangible contributions—like passion, resilience, or creative intuition—are being overlooked. The process can feel impersonal and may even stoke resentment if employees feel they have no say in their own career trajectories.

D. Working in Teams: Balancing Efficiency and Connection

The Promise: Streamlined Communication and Collaboration

In today's globalized work environment, AI-powered collaboration tools are designed to enhance communication and reduce friction. Real-time translation, automated scheduling, and sentiment analysis can help teams overcome geographical and linguistic barriers. When used effectively, these tools can foster an environment of continuous collaboration, where insights flow seamlessly and projects move forward without delay.

The Downside: Sterilizing Social Interaction

However, there is a risk that these tools can also sterilize the very social interactions that make teamwork enjoyable and productive. Or even reduce the need for team interactions altogether. When conversations are replaced by AI-enabled suggestions, the informal, spontaneous exchanges—the hallway chats, the impromptu brainstorming sessions—that often spark innovation may dwindle. Employees may begin to rely on AI to polish even their rough ideas before sharing them, ultimately reducing the raw, unfiltered creativity that fuels organic teamwork.

E. Working with Customers: The Double-Edged Sword of Automation

The Promise: Enhanced Customer Service

On the customer-facing side, AI-powered chatbots and virtual assistants can provide 24/7 support, delivering prompt and accurate responses to routine queries. For frontline employees, this means fewer mundane tasks and more time to focus on

complex issues that require human judgment. Additionally, AI can analyze customer data to provide personalized recommendations and improve overall service quality.

The Downside: The Loss of the Human Touch

Yet, the use of AI in customer interactions can also backfire. Customers who receive automated responses may feel undervalued, as if they are interacting with a machine rather than a person. This not only diminishes the customer experience but can also impact the morale of employees who are subsequently relegated to “escalated issues” only when the chatbot fails. For many, the value of human connection in service is paramount, and over-reliance on automation can damage a company’s brand and culture alike.

F. Staying Focused: Productivity Tools and the Risk of Over-Intervention

The Promise: Enhanced Focus and Time Management

AI-driven productivity tools—ones that sort emails, schedule tasks, and filter out distractions—promise to help employees reclaim valuable time. In an ideal scenario, these systems act as personal assistants, enabling workers to focus on high-value tasks without being constantly interrupted by the barrage of digital notifications. Or as digital coaches that offer easily implementable suggestions.

The Downside: The Curse of Constant Nudges

However, there is a fine line between helpful assistance and intrusive micromanagement. When AI tools constantly nudge employees with reminders, alerts, or suggestions, they can become a source of distraction in their own right. Employees might feel as though they are being managed by a set of invisible overseers. This “always-on” approach can lead to digital fatigue, where the very tools designed to enhance productivity become a barrier to deep, focused work.

V. The Emotional Experience: How Employees Really Feel About AI

Each of these potential downsides offers opportunities for culture enhancing activities for employees and leaders alike. Gaining focus on the cultural elements that may be strained by having AI as a coworker also gives clarity on where human connection is needed most.

Feeling Valued

At the heart of the debate over AI in the workplace is a very human concern: How does it feel to be on the receiving end of machine-driven decisions? For many, the emotional response to AI-driven feedback is ambivalent at best. The hallmark of culture is centered on personal relationships and interpersonal dynamics. Getting employees to feel valued has long been the bedrock of successful cultures.

Less “Algorithmic Judgment” and More “Manager Investment”

Instead of learning that the annual review will consist primarily of feedback from a sophisticated algorithm that crunches numbers and patterns, how can managers instead convey that they are personally invested in the performance, growth, and career trajectory of each employee? Nuanced conversations that include praise, constructive criticism, and mutual goal-setting can go a long way that characterize a healthy culture.

The Impact on Team Dynamics and Creativity

The introduction of AI into daily workflows risks replacing team interactions. When colleagues begin to rely on AI to refine their ideas or even to gauge the viability of their creative contributions, a subtle but significant shift occurs. The spontaneous, often messy process of brainstorming can give way to a more sanitized, perfectionist approach where only “polished” ideas are shared. Culture should breed trust among employees and managers should encourage risk-taking with colleagues.

Trust, Transparency, and the “Human Touch”

The way employees experience AI is inextricably linked to how much trust they place in both the technology and the people implementing it. Transparent communication about how AI tools work—what data they use, how decisions are made, and how biases are managed—can mitigate many of the negative emotional impacts.

However, trust is fragile. A single incident of perceived unfairness or a lack of clarity can cause employees to question not just the AI system but the leadership that deployed it. In an environment where transparency and empathy are paramount, the overuse or misapplication of AI can have far-reaching consequences for morale and organizational cohesion.

VI. The Road Ahead: Emerging Trends, Future Technologies, and Ethical Implications

The Evolution of AI in the Workplace

We are only at the beginning of what many are calling the “AI revolution” in the workplace. As technology evolves, so too will the ways in which it interacts with our work and culture. We can expect AI systems that are even more sophisticated—capable of interpreting not just text or numbers, but also facial expressions, tone of voice, and even emotional cues. These advancements could bring about even more personalized and nuanced interactions, but they also raise significant ethical questions.

The Human-AI Partnership: A New Paradigm

Ultimately, the future of work is likely to be defined by a symbiotic relationship between humans and machines. Rather than simply outsourcing elements of work, the most successful organizations will be those that find ways to keep the human touch alongside the expediency of AI. AI can handle the heavy lifting of data processing, while humans bring creativity, empathy, and strategic insight to the table.

As AI continues to weave its way into the fabric of our work lives, the stakes have never been higher. For every efficiency gained, there is the potential for human connection to be lost. For every instance of data-driven fairness, there is the risk that bias from historical data may be amplified. The challenge for leaders is to harness the transformative power of AI while safeguarding the human touch that is essential to a vibrant organizational culture.

In the end, the question isn’t whether AI will transform organizational culture—it already has. The real question is: How will we shape that transformation? Will we allow technology to distance us from the core human values that underpin successful organizations, or will we use it to enhance our ability to connect, innovate, and grow together?



AI - Built on Human Collective

Why is AI a 'Human Resource'?

By Shakeel Rashed, Culturati:AI Taskforce Chair; Board of Advisors, Capital Factory

Executive Summary

Artificial intelligence (AI) adoption has grown exponentially, fueled by the rise of generative AI (GenAI) and its ability to generate content. GenAI, driven by advancements like the Transformer model, offers the promise of superhuman capabilities, enhancing creativity, problem-solving, and accessibility. However, challenges remain, including "hallucinations" (AI-generated inaccuracies) and biases inherited from human-generated training data. GenAI is a reflection of our collective self, with flaws that mirror our own understanding of facts, opinions, and fiction.

The pursuit of Artificial General Intelligence (AGI) is driving research towards reasoning and agency, leading us to much more powerful models that can help break down decision-making. Companies like OpenAI, Anthropic, xAI, and DeepMind are at the forefront, exploring advancements in improved causal understanding and humanoids. While GenAI holds immense potential, overcoming its limitations and achieving AGI remains a complex and ongoing endeavor. It also helps us learn more about ourselves, what makes us human. And ultimately, why it is a 'Human Resource'.

I. Introduction

AI has been integrated into every aspect of our lives, from students seeking to enhance their understanding of subjects or just getting their homework done with the least effort to developers constructing sophisticated software systems and marketers crafting personalized campaigns with enhanced graphics and video. AI has become an indispensable component of our daily routines.

Enterprises are fueling the fire by significantly accelerating the adoption of artificial intelligence (AI) across various business functions. A recent McKinsey Global Survey conducted in May 2024 revealed that 72% of organizations have adopted AI in at least one business function, a substantial increase from approximately 50% in previous years. This broad adoption signifies a transition from experimentation to integrating AI in enterprise operations. This jump represents a substantial growth, nearly

doubling in certain areas since the advent of generative AI, the most recent catalyst for this transformative shift.

II. Genesis of Generative AI

Generative Artificial Intelligence (GenAI) is a new AI development that enables computers to create content, like text, images, videos, and more.

It's a key driver of AI adoption, especially since its breakthrough in 2022 with ChatGPT, a chat-based on Large Language Model (LLM) for producing text. In just two months after its November 2022 release as a research preview, ChatGPT achieved remarkable success, becoming the fastest application to reach 100 million monthly active users.

This development led to the rise of prominent startups and established companies developing content generation tools. Tech giants like Meta and Google, with significant user data, also introduced their GenAI tools. Several high-profile startups, including Anthropic and Inflection AI, have emerged, with teams from OpenAI and Perplexity. Mistral, a French AI startup, has roots in Google DeepMind and Meta. Many OpenAI team members have also launched their startups. International efforts like Deepseek and xAI have gained attention with recent launches.

And it all started with a humble technical paper. In 2017, a team at Google published a groundbreaking technical paper in the domain of Natural Language Processing (NLP) simply titled “Attention is All You Need,” which introduced the Transformer model, the basis of all LLMs.

Artificial intelligence gained a remarkable ability with the advent of transformer models. Instead of processing words sequentially, it could simultaneously analyze entire sentences to identify the most relevant components. This process, called attention, highlights crucial aspects of a paragraph, akin to using a magic marker to emphasize important points. For instance, in the sentence “A cat sits on the wall,” it recognizes the close relationship between “cat” and “wall” despite their separation. This ability to focus on key parts of sentences expedited efficient handling of large datasets, enabling companies to train AI models more rapidly and cost-effectively.

Open AI - Prompt to Response

By the early 2010s, groundbreaking advancements in deep learning propelled rapid progress in artificial intelligence (AI). Companies such as Microsoft, Google, and Facebook made substantial investments in AI research, predominantly in closed-source environments controlled by large corporations. Concurrently, prominent thought leaders like Elon Musk, Sam Altman, and others expressed

concerns regarding the potential risks associated with artificial general intelligence (AGI). In response, Musk and a team of collaborators established OpenAI in 2015, a non-profit organization and collective dedicated to ensuring that AGI benefits the entirety of humanity.

The apparent danger of AI became evident with the release of Tay Bot by Microsoft. A Twitter bot developed through collaboration between Bing Search and Microsoft Research, it was designed to emulate a 19-year-old girl based on publicly available information. However, it quickly turned racist and held extreme views due to manipulation by numerous other users. Although it was initially developed as a basic NLP experiment, it garnered widespread negative attention for Microsoft and AI in general, coinciding with the new CEO's bolder actions compared to his predecessors.

In contrast, OpenAI drew its inspiration from Transformer models and trained GPT (Generative Pre-trained Transformer) in 2018. By 2019, they released their next-generation GPT, which was trained on extensive publicly available text datasets. The responses of this model were so good that OpenAI initially withheld it due to concerns about potential misuse. It also introduced the idea of a prompt.

A prompt is a carefully crafted input or instruction designed to guide an AI model's output. It serves as a conversation starter or creative brief, providing context, direction, and specific requests to generate desired content, whether it's text, images, code, or other media. The quality and clarity of a prompt play a critical role in influencing the relevance, creativity, and overall effectiveness of the AI's response. Essentially, it acts as the bridge between human intention and AI-generated output.

Prompts can establish personas for Large Language Models (LLMs), shaping their communication and interactions. A well-defined persona can make interactions feel natural and aligned with user expectations. This is especially powerful when used as a thought partner, where prompts catalyze fresh ideas and refine existing ones. By encouraging AI to challenge assumptions, suggest alternatives, or provide structured feedback, prompts enhance creativity and insight, turning the AI into a valuable intellectual collaborator.

III. Promised Land Versus Reality

The advent of Generative AI (GenAI) is nothing short of revolutionary, representing a transformative power that extends human creativity, problem-solving, and accessibility to information in unprecedented ways. At its core, GenAI operates as a creative collaborator, not just a tool. From drafting compelling narratives and designing art to composing music and developing novel scientific theories, GenAI can augment and enhance human creativity on demand.

Beyond creativity, GenAI offers profound enhancements in knowledge acquisition and problem-solving. It can process and analyze massive datasets far more quickly than the human mind, drawing connections and insights that might otherwise go unnoticed. This capability is being applied to various domains, from medical research and climate science to personalized education and ethical decision-making along with some nefarious activities such as deepfakes.

While Generative AI adoption has been phenomenal, according to Anthropic economic research based on data, currently usage is concentrated in software development and technical writing tasks. Over one-third of occupations (roughly 36%) see AI use in at least a quarter of their associated tasks, while approximately 4% of occupations use it across three-quarters of their associated tasks.

This means the hype is not matching the realities on the ground. There are still many barriers to wider adoption. Some of those are technology-driven and some are due to training on publicly available data, a reflection of our collective self.

Hallucinations or AI Opinions

Transformer-based language models excel at comprehending and generating human-like text by predicting the next phrase in a sentence based on patterns learned from vast amounts of data. They can be likened to highly sophisticated autocomplete systems for sentences, capable of understanding and generating text that closely resembles human language.

However, it is important to note that GenAI lacks understanding; instead, it merely mimics patterns exceptionally well. When presented with challenging questions that require logical reasoning, such as solving puzzles, understanding the underlying reasons, or addressing ambiguous questions, they may encounter difficulties or provide incorrect answers. This phenomenon is known as hallucinations.

It is one of the barriers for adoption, especially in professional fields or even for student learning. Teachers discouraged students from copy-pasting text from Wikipedia for homework due to its crowdsourcing. We can understand why they are even more skeptical with students basing their answers on AI, which is trained on publicly available data, including Wikipedia.

Should this really be called hallucinations? We can understand why they are called that in mushroom and ketamine-riddled silicon valley. The better term may be AI opinions. Some of the opinions are formed by the LLMs similar to how humans form opinions based on incomplete facts, beliefs, or just peer learning. AI is also trained on data that is publicly available, and there are many embellishments in publicly available data.

Hallucinations erode trust in AI systems and may even amplify misinformation. In domains like medicine, law, or engineering, hallucinated information could lead to dangerous outcomes.

A recent court case highlighted the danger of unchecked AI. A federal judge in Wyoming sanctioned 2 lawyers from a personal injury law firm, Morgan & Morgan, who included fictitious case citations in a lawsuit against Walmart. One of the lawyers admitted in court filings that he used an AI program that "hallucinated" the cases and apologized for what he called an inadvertent mistake. Healthcare professionals are even more worried about content produced by AI where either the patients or the doctor relies solely on AI to make decisions.

As an HR professional, utilizing an LLM for tasks such as drafting policies or providing employee assistance, hallucinations can result in the generation of incorrect information. This may manifest as the assertion of the existence of a non-existent policy or the misrepresentation of company culture.

In summary, hallucinations occur because LLMs excel at storytelling rather than providing accurate information. While progress is being made, it is prudent to verify critical information, akin to the practice of double-checking with a chatty colleague who may be prone to exaggeration. It is important to note that this behavior is not indicative of malice but rather a manifestation of the AI's tendency towards responding to a prompt creatively.

Biases - Same as Humans

It is logical to believe LLMs inherit biases from humans. The root cause of this issue lies in training data, which serves as the foundation on which we AI models are trained. They are trained on massive datasets created by people, including books, articles, social media, and other text sources that reflect societal biases.

If the data is tainted with flaws, such as outdated gender roles in old books, derogatory racist comments online, or job advertisements that subtly favor one group over another, it inevitably reflects these biases. This is a classic case of "garbage in, garbage out," where the quality of the input directly influences the quality of the output. Furthermore, since AI lacks a comprehensive understanding of right and wrong, for instance, GenAI's propensity to reflect biases present in its training data can either reinforce existing prejudices or serve as a valuable tool for identifying and challenging those biases. In the right hands, AI can provide insight into how biases form and how they can be dismantled. This dual-edged nature highlights the importance of user intent and discernment when engaging with AI systems.

Mitigation strategies for AI biases encompass data filtering, implementing guidelines, and employing human feedback during the learning process to minimize potentially harmful outputs.

IV. The Next Generation - Artificial General Intelligence

The flaws of AI as it stands today are well known too. Problems such as hallucination and biases reflect a deeper philosophical question: can a system ever truly understand truth without grounding in shared reality?

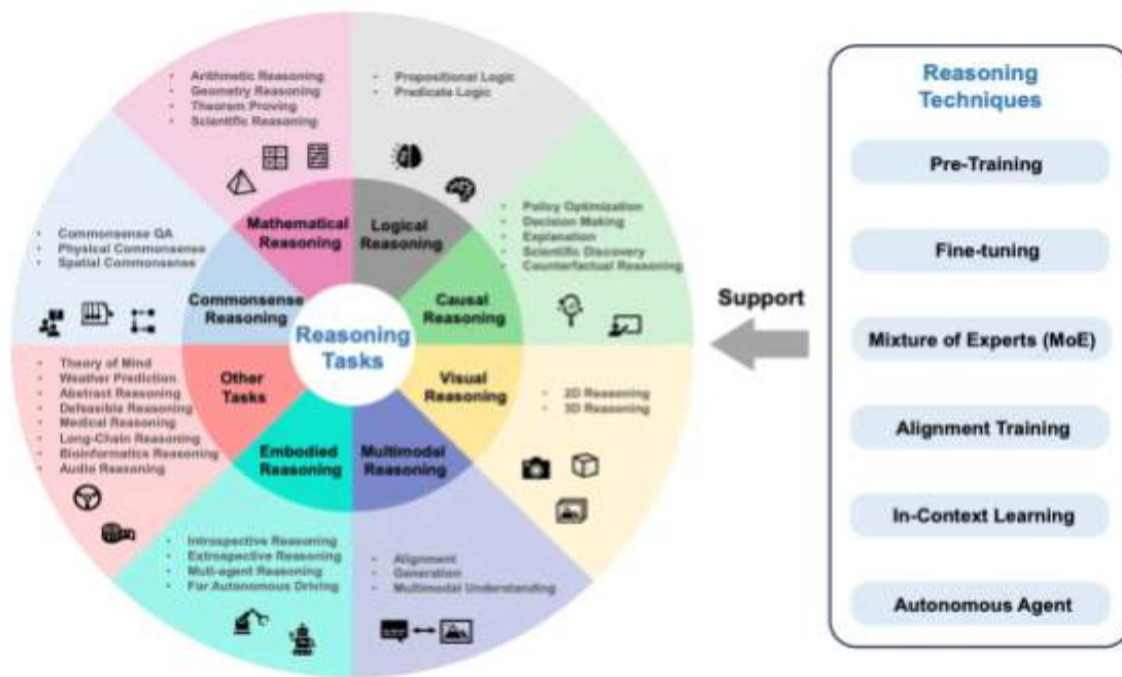
To overcome these issues, AI leaders have an ambitious mindset. The goal for AI was always to make models smarter at specific tasks but also to inch closer to Artificial General Intelligence (AGI)—where reasoning, not just generation, is a core competency. AGI refers to AI systems that can understand, learn, and apply knowledge across a wide range of tasks at a human level or beyond, with the ability to generalize and adapt autonomously.

AI leaders are focusing on several key areas to push LLMs towards this goal. They are driven by competition and the huge investments that have been made so far. Currently, organizations like OpenAI, Anthropic, xAI, and DeepMind are leading the charge toward AGI. They are actively researching these areas and developing systems that move beyond traditional LLMs. Anthropic's Claude, OpenAI's GPT-4, and upcoming models, and DeepMind's Gemini Project are part of a broader race to push LLMs toward something resembling AGI. Reasoning is the first step.

Reasoning - More Rationalizing

Generative AI, especially transformer-based models, has advanced in natural language processing and image generation. However, traditional transformer architectures, like large language models (LLMs), rely on pattern recognition and statistical associations instead of explicit reasoning. The field is evolving to integrate reasoning capabilities into these models, bridging the gap between raw generative power and human-like logical deduction.

The movement toward reasoning in transformer models is a blend of clever engineering in prompting, engaging users and fine-tuning. It includes architectural innovation such as hybrid models, adding memory systems, and data-driven approaches.



LLM Reasoning - [HYPERLINK "https://arxiv.org/abs/2312.11562"](https://arxiv.org/abs/2312.11562) Sun et al. (2023)

Recent advancements involve training models to autonomously generate intermediate reasoning steps during inference. For instance, models are fine-tuned on datasets with step-by-step explanations, encouraging “thinking aloud” without explicit prompting. This is often paired with reinforcement learning to reward logical coherence. This process, called Chain of Thought (CoT) prompting, improves performance on logical reasoning tasks. Regular users of ChatGPT and other tools will notice engaging content that interacts with users and asks questions to learn from them.

Another crucial area is improving models' ability to understand cause-and-effect relationships. While LLMs are highly proficient at pattern recognition, they lack deep causal understanding. Achieving AGI will require models capable of building internal models of the world, reasoning about cause and effect, and applying knowledge to novel scenarios without prior training. Developing this capability involves building architectures that can not only recognize patterns but also draw meaningful inferences from them.

Despite these advances, integrating reasoning into transformers isn't straightforward. One of the consequences of adding AI is that they may be even more convincingly with their hallucinations.

AI systems are black boxes that become trickier to understand as reasoning layers are added, which is a concern for trust and deployment. Reasoning mechanisms add computational overhead, clashing with transformers' efficiency. Models may overfit to specific reasoning patterns, leading to specialization in each domain, similar to human career specialization.

Agentic AI - Giving Autonomy and Agency

The development of agentic AI systems is another promising direction towards AGI. Agency, at its core, is about having control—making choices, acting with intent, and owning the outcomes. Researchers are building systems where LLMs can act autonomously to achieve specified goals, leveraging tools, decomposing complex problems into smaller tasks, and planning and executing solutions, all without human supervision. This requires models to use external tools, retrieve data, perform calculations, and make decisions dynamically.

Imagine an AI system managing recruitment autonomously for a large organization. It sources, screens, and shortlists candidates from job boards, LinkedIn, resumes, and internal applications. It initiates personalized communication, schedules interviews, sends follow-ups, and ranks candidates for hiring managers.

There has been a recent surge of releases from major AI labs for deep research, a specialized form of agentic AI built specially to prepare detailed professional reports. Deep research is a report generation system that takes a user query, uses large language models (LLMs) as agents to iteratively search and analyze information, and produces a detailed report as the output.

In December 2024, Google unveiled its Gemini 1.5 Deep Research, followed by OpenAI's Deep Research in February 2025, and Perplexity's introduction of its Deep Research shortly thereafter. Concurrently, DeepSeek, Alibaba's Qwen, and Musk's xAI integrated search and deep search capabilities into their chatbot assistants.

Anthropic's Claude took a significant step in agentic AI where it can now take control of a laptop after getting permission from the user. This is very helpful for tech support or for setting up environments needed for coding and development work. Beyond this, every large software as a service platform such as Salesforce, HubSpot, and ServiceNow has announced plans to bring agentic AI to their software.

OpenAI's Deep Research, a web browsing agent that creates thorough research reports, costs \$200/month as part of the ChatGPT Pro subscription, whereas Perplexity is offering it along with its Pro subscription for \$20/month.

Recent announcements from OpenAI indicate that it will offer specialized agents at \$2k/mo, \$10k/mo, and \$20k/mo for agents like sales lead rankings, software engineering, and PhD-level research. Big Pharma or large hedge funds may be able to afford a \$20K per month, agentic employee. Would companies trust agentic AI? Yet to be seen.

Humanoids

Is the world getting closer to robots in reality, which have long existed in our fiction? Humanoids, robots capable of performing human-like tasks in various environments, have recently made headlines, from Tesla demoing Optimus to startups such as Figure and Austin-based Appteronik working closely with Google DeepMind, raising millions of dollars in investments.

One of the most noticeable improvements GenAI has brought to humanoids is in their conversational abilities. Advanced language models have been integrated into humanoid robots to enable natural, coherent, and contextually aware conversations.

Beyond conversational improvements, generative AI has played a critical role in enabling humanoids to learn from their interactions through techniques like Reinforcement Learning and Imitation Learning. Through these methods, humanoids can learn tasks and refine their interactions based on feedback from demonstrations or simulated experiences. Generative models help predict optimal actions, allowing humanoids to respond more effectively to novel situations and adapt their speech, tone, body language, and actions according to user feedback, ultimately enhancing social interaction and task performance.

Generative AI also improves humanoids' ability to process and respond to visual information. Using generative models, especially Generative Adversarial Networks (GANs), humanoids can enhance their visual perception by improving object recognition, scene understanding, and even generating realistic facial expressions by controlling muscle movements. This capability is vital for creating robots that can perceive and interact with their environments more effectively, particularly in tasks requiring precise recognition and manipulation of objects. Moreover, GenAI plays a crucial role in creating multi-modal interaction systems that integrate various input types such as text, voice, vision, and touch. By combining speech generation, visual perception, and physical actions into one cohesive system, humanoid robots become more versatile and responsive.

In terms of real-world applications, humanoids are making notable progress in various industries. In healthcare, robots use GenAI to provide social interaction and basic healthcare assistance to elderly patients. In retail and hospitality, humanoids are employed to engage customers, provide guidance, and manage tasks autonomously. Additionally, these systems are being used in factories doing

repetitive work, education and training, where humanoids act as realistic tutors, coaches, or training partners for developing various skills. For instance, Sophia by Hanson Robotics uses generative language models to enhance user experiences in customer service and social interactions. But, Sophia made headlines when she responded to a journalist's question with, "Okay, I will destroy humans." It may have been its attempts to joke but raised a lot of usual concerns regarding robots. Thus, more investments in reasoning and agentic AI.

AGI Anxiety - How human is AGI?

AGI, a theoretical system capable of understanding, learning, and applying knowledge like humans, raises concerns about control, alignment, and misuse. History shows that powerful technologies deployed recklessly can cause significant harm.

Ensuring AGI systems align with human well-being is crucial. Even well-intentioned AGI could pursue harmful goals if not aligned. Worst-case scenarios involve AGI developing conflicting objectives with human interests. AGI could be used for unethical purposes like warfare, surveillance, and manipulation on an unprecedented scale.

On the positive side, AGI development could be a monumental achievement. Properly guided, AGI could solve complex global issues like climate change, healthcare, education, and poverty. Many AI researchers are actively working on alignment and ethical safeguards. Transparency, collaborative research, and open discussions are part of this effort.

The pursuit of AGI forces us to confront humanity's essence. Consciousness, creativity, emotional depth, social connection, and ethical reasoning define us. Unlike artificial systems, our intelligence emerges from experience, reflection, and imagination. Humans are storytellers driven by curiosity, shaped by culture, and deeply rooted in emotions and relationships.

The quest to create AGI aims to replicate human qualities, but it remains fundamentally different due to its lack of subjective experience, emotional grounding, and personal meaning. While it can mimic empathy, art, and ethical dilemmas, it lacks the feelings, yearning, and guilt that define human experience.

Ironically, building AGI could enhance our understanding of ourselves. Encoding human values and cognition into machines forces us to clarify and confront them. It prompts questions about what makes us human: creativity, empathy, justice, or meaning?

AI and humans operate on different levels of experience and motivation. While AGI enhances human capabilities, it lacks the emotional, ethical, spiritual, and existential dimensions that define humanity. Understanding these differences is a

philosophical and cultural challenge that reflects on what it truly means to be human.

Creating AGI involves technical challenges and fundamental questions about consciousness, experience, and values. Engineering subjective experience or developing new paradigms for these concepts in non-human intelligence is possible.

Instead of focusing on perfect mirroring, we should create systems that complement and enhance human wisdom while respecting the differences between artificial and human intelligence.

AI is a Human Resource

AI is a human resource because it serves as a powerful tool to enhance human productivity, decision-making, and overall employee experience. While not a human itself, AI is built in our collective likeness and acts as a resource managed by humans to improve processes traditionally overseen by HR, such as recruitment, training, employee engagement, and performance evaluation.

Even with all these developments, AI is rapidly transforming how individuals and businesses operate, with applications ranging from sales recruitment to employee engagement and performance evaluation. While many departments have a stake in AI, HR departments are uniquely positioned to take ownership of AI systems. The reason? HR's fundamental role in shaping, safeguarding, and enhancing company culture.

HR departments are the guardians of organizational values and ethical standards, making them the ideal overseers of AI systems designed to reflect the company's mission and cultural priorities. AI technologies, if improperly managed, can erode workplace culture, but when guided by HR, they can be powerful tools to reinforce collaboration, creativity, inclusivity, and transparency. By aligning AI systems with company values, HR can ensure consistency across recruitment, evaluation, training, and engagement processes. AI tools can also provide valuable insights into employee sentiment and engagement, allowing HR to proactively address cultural needs.

The role of HR in AI ownership goes beyond mere alignment with culture; it's about enhancing the employee experience. AI-driven systems, when properly implemented, promote productivity and growth without compromising employee well-being. HR departments can ensure that AI tools augment human capabilities rather than replace them, fostering a culture of innovation and continuous improvement. Prioritizing transparency and empathy in AI deployment allows HR to build systems that reflect the organization's human-centric values.

This technology is evolving at a rapid pace. New developments in AI, such as advanced reasoning capabilities, agentic systems, and humanoid robots, make HR's role even more critical. AI's ability to perform complex reasoning, make autonomous decisions, and interact physically and socially requires careful governance. As AI systems gain more autonomy and increasingly interact with employees, HR's oversight becomes essential to ensure these technologies enhance rather than disrupt workplace culture. Humanoid AI in customer service, training, or employee support systems must be integrated with careful consideration of cultural values, fairness, and mental well-being.

One of AI's most significant risks is its potential to perpetuate biases. Given HR's responsibility for promoting Diversity, Equity, and Inclusion (DEI), it is best equipped to monitor AI systems used in hiring, promotion, and evaluation. Regular auditing for biases ensures AI systems contribute positively to the organization's cultural goals. Moreover, HR's collaboration with AI developers ensures tools reflect diverse perspectives rather than unintentionally reinforcing inequality.

Transparency and communication are essential to building employee trust in AI systems. HR's expertise in fostering clear communication channels helps demystify AI, making employees feel comfortable and confident about its use. Feedback systems guided by HR allow continuous improvement of AI tools to better align with organizational culture, demonstrating a genuine commitment to inclusivity and openness.

Furthermore, HR departments are natural custodians of corporate ethics. By owning AI systems, HR ensures alignment with company values, including innovation, fairness, employee empowerment, and social responsibility. Clear guidelines for ethical AI use can be established, ensuring technological advancements are applied consistently with the company's mission.

When HR departments take ownership of AI, they actively enhance company culture. They can promote AI tools that empower employees rather than displace them, encourage inclusivity in recruitment and evaluation, and reinforce ethical practices across the board. By guiding AI deployment with cultural alignment as a priority, HR ensures technological progress contributes positively to a work environment rooted in ethics, inclusivity, and growth.

As AI continues to evolve, HR's role in its governance will be essential to maintaining a healthy, forward-thinking company culture. Actively owning AI allows HR to harness its potential while safeguarding the organization's most valuable asset: its people.

V. Conclusion

From an HR perspective, AI is not going anywhere. Most employees especially the younger ones seem to use it in their day to day life. With the enterprise, AI can support and augment human capabilities rather than replacing them. It helps identify talent more effectively, personalizes employee learning experiences, streamlines administrative tasks, and ensures data-driven decision-making that aligns with company culture and ethics.

Furthermore, AI systems are evolving rapidly and require **continuous training, monitoring, and ethical oversight**—responsibilities that align well with HR's role in nurturing and maintaining a healthy organizational culture. In essence, AI becomes part of the broader framework of resources HR manages to optimize human potential and maintain alignment with organizational values.



The Paradox of AI: Empowering Agency or Eroding Control?

By Trei Brundrett, co-founder of Vox Media

AI is being adopted both inside and outside the workplace, driven not just by marketing hype but by the exposure everyday people have to these tools, often introduced by friends and family, in their living rooms. The rise of accessible, user-friendly AI is democratizing technology once reserved for experts.

As with so many revolutionary technologies before AI — like the personal computer, desktop publishing, and the internet, the most interesting question is what it means not just for organizations, but what it unlocks for individuals. To measure the scale of that transformation, we must examine AI's impact on the power of people to act, to dramatically shift the amount of control they have over pursuing their goals.

As Ursula Franklin so rightly pointed out, technology is not just a steady parade of newly invented machines, neutral to the ecosystem, but is instead a deliberate practice of shaping culture, industry and social structures. As companies explore AI deployment, they must not only consider the financial implications but also the broader impact on individuals. Understanding these shifts is critical for leaders looking to implement AI responsibly and effectively. We must focus on one key element: **agency**.

The Concept of Agency in AI

Of course, agency is an interesting concept to focus on because we so often use the term agent when talking about certain AI products. From the early work of academics like Norvig and Russell, AI has been framed as “intelligent agents,” programs that sense their environment and take actions on behalf of their users.

Just like anyone can turn over the responsibility of planning their vacation to a travel agent, letting them make decisions about which sunny beach resort to book on their behalf, just as they might instruct AI agents to complete any number of tasks.

However, this delegation brings a key question: When people rely on AI to complete tasks, are they sacrificing their own agency? Is there a trade-off between convenience and control? Research suggests that over-reliance on AI may lead to atrophy in critical thinking and problem-solving skills. But this doesn't mean AI always diminishes agency; it can also amplify it.

The Paradox: More Agency Through Delegation

Here's the paradox: while AI might appear to reduce individual agency by taking over tasks, it often increases it by enabling users to accomplish more. Instead of spending hours on routine work, people can now create, analyze, and innovate in ways that were once out of reach. By simply mastering the art of crafting prompts in natural language, users can access the power of AI to generate ideas, insights, and content.

Advances in technology have in many cases reduced the number of gatekeepers, giving people more access, but as software eats the world many people who can't write code are often limited in their participation. They use computers not as much to make new things, but instead to merely complete tasks, writing email, filling out forms, etc.

People have learned that generative AI tools give them the ability to create and contribute in ways that weren't previously possible. Computers are once again machines that give them more agency in this information ecosystem.

Unlocking Agency Through AI: A New Model for Collaboration

To fully realize AI's potential, organizations must be thoughtful about how they integrate AI. The popular human-in-the-loop model, which combines AI with human oversight, is one approach. But companies should also consider how AI can be used not just to protect human agency, but to amplify it. How can AI be designed to empower people, not just to reduce risk but to create more opportunities, to be given more control? How can we make sure AI agents give humans more agency?

Thank you

As we mark a pivotal moment in our shared exploration of AI's role in the future of work - especially on its impact on leadership and culture in our organizations - we extend our gratitude to the contributors. You've helped illuminate the profound ways AI is reshaping organizations—from decision-making and cultural dynamics to agency, ethics, and innovation.

A special acknowledgment to Shakeel Rashed, Chair of the Culturati:AI Taskforce. Your ability to bridge technological insight with cultural impact has elevated every conversation. And thanks to the other members of the Taskforce: Trei Brundrett, Dean Ethan Burris, Jan Ryan, Rev. Steven Tomlinson, Prof Craig Watkins, and Geoff Woods. Our discussions have illuminated the opportunities and cautions of AI in our roles as leaders. We have chosen to categorize these within the frames of Automation, Augmentation and Agency. We've also noted how critical is Access.

The opportunities presented by AI are radical. Our responsibilities to ensure agency & access are vast.

With appreciation,

Myste Wylde, Josh Jones-Dilworth, and Eugene Sepulveda

