
AI for Productivity in the Age of Agentic AI



The AI4X Project Team




github.com/ai4x-org/AI4Prod ai4x-org.github.io ai4x.org@gmail.com

Abstract | Throughout history, transformative technologies have repeatedly reshaped production systems and economic activity. In the current wave of AI development, Agentic AI systems are emerging as a new class of productive instruments capable of goal perception, multi-step planning, environmental interaction, and autonomous workflow coordination. Beyond automating routine tasks, these systems have the potential to augment knowledge work, reorganize decision-making, and reshape the relationship between human labor, organizational routines, and machine intelligence. This development gives rise to the emerging research direction of AI for Productivity (AI4Productivity), which examines how advanced AI systems affect productivity across economic and social contexts.

In this paper, we conduct an in-depth investigation for explaining how Agentic AI creates productivity. We define AI4Productivity as the use of Agentic AI in real-world workflows to deliver economically valuable work with less human effort, and identify two core mechanisms of productivity gains: *Efficiency*, which reduces the time and labor required for existing tasks, and *Expansion*, which makes previously infeasible work economically viable. We further characterize the evolution of AI systems through four stages: conversational assistant, reactive operator, adaptive coordinator, and self-governing production system. Building on this framework, we review representative AI productivity tools and examine the sectoral diffusion of Agentic AI across three levels of adoption: early-stage exploration in manufacturing, agriculture, real estate, and government services; broader integration in trade, media, education, and legal services; deeper workflow orchestration in information technology, finance, and healthcare. Finally, we discuss key challenges including safety, long-horizon reliability, academia-industry gaps, adoption barriers, governance, and labor-market uncertainty. This paper provides a structured account of how Agentic AI is reshaping productivity and outlines future directions for reliable and socially beneficial AI-driven production systems.

We hope this paper will serve as a foundation for advancing AI4Productivity research and guiding the design of AI systems that deliver reliable, scalable, and broadly beneficial productivity gains.



Project Team

*Project Lead.

Project Leads and Core Contributors

Zhiheng Xi* (Fudan University)	Enyu Zhou* (Fudan University)	Xinyu Fang* (Zhejiang University)
Zhe Sun (Fudan University)	Baodai Huang (Fudan University)	Jiajun Sun (Fudan University)
Bicheng Deng (Fudan University)	Zhihao Zhang (Fudan University)	

Contributors and Technical Advisors

Wenxiang Chen (Fudan University)	Jiazheng Zhang (Fudan University)
Shichun Liu (Fudan University)	Xin Guo (Fudan University)
Zhikai Lei (ECNU)	Junke Wang (Fudan University)
Senjie Jin (Fudan University)	Yang Nan (Fudan University)
Yajie Yang (Fudan University)	Rui Zheng (Fudan University)
Hang Yan (Fudan University)	Yuchen Tian (HKBU)
Fang Wu (Stanford University)	Ling Yang (Princeton University)
Yingcheng Wu (Stanford University)	Yinghui He (Princeton University)
Kejun Ying (Stanford University)	Cheng Qian (UIUC)
Zhenfei Yin (University of Oxford)	Xuandong Zhao (UC Berkeley)
Mingchen Zhuge (KAUST)	Mengyue Yang (University of Bristol)

Senior Advisory Panel

Jun Wang (University College London)	Philip Torr (University of Oxford)
Zuxuan Wu (Fudan University)	Xipeng Qiu (Fudan University)
Qi Zhang (Fudan University)	Xuanjing Huang (Fudan University)
Tao Gui* (Fudan University)	Yu-Gang Jiang (Fudan University)



Series Preface

"Making machines think and act like humans" has long been a founding aspiration of AI. Technically, this aspiration has been repeatedly reinterpreted across paradigms: symbolic AI viewed intelligence as rules and deduction; statistical learning as pattern induction from data; task-oriented deep learning, such as CNNs and RNNs, transformed learned representations into practical systems; and recent scaling has advanced toward unified parametric models integrating language, vision, and broad knowledge.

For a long time, progress across these paradigms was assessed mainly through closed benchmarks and leaderboard comparisons. Yet from a broader public perspective, AI's most meaningful leaps are not merely higher scores, but the moments when it enters the real world, becomes embedded in everyday processes, and delivers sustained, measurable value. This requires systems that can operate under real constraints, use tools, collaborate with people, manage long-running tasks, and close the loop with concrete results.

This shift changes the central question. Viewed through a Wittgensteinian lens, *meaning is defined by use*. The question is no longer only "Can the model answer accurately?", but "Can it execute reliably?", that is, can it turn parametric knowledge into trustworthy actions within complex human systems and drive tangible value? Once AI becomes part of a value chain, it is judged not only by accuracy or fluency, but by throughput, quality, reliability, risk, accountability, and return on investment. The frontier is thus moving from model capability to real-world system performance.

This perspective motivates a series of studies on how AI creates durable value in real-world systems. Rather than treating AI progress as a catalog of model capabilities or application areas, we examine how technical capability enters real value chains, becomes practical, trustworthy, and accountable in deployment, and is transformed into durable value. We develop this unified inquiry into AI for value through three closely related settings.

AI for Productivity studies how AI is integrated into everyday work to generate economic and organizational value. It focuses on how AI accelerates routine tasks, coordinates multi-step workflows, supports decision-making, and enables work that was previously too costly, too slow, or too complex. The central concern is utility: whether AI can be embedded into real workflows in ways that improve productivity, quality, and operational efficiency.

AI for Science moves from labor to discovery. It examines how AI supports the scientific cycle, including hypothesis generation, experiment design, simulation, data analysis, interpretation, and literature synthesis. Here, value is not only economic but epistemic: AI must help produce knowledge that is rigorous, traceable, reproducible, and verifiable. The central concern is whether AI can expand the scope, speed, and reliability of scientific exploration while preserving scientific rigor.

AI for AI turns inward to the engine of AI progress itself. It assesses how AI systems are used to improve AI through stronger feedback loops, including agent-driven data generation, automated evaluation, synthetic supervision, model refinement, and iterative self-improvement. Here, value lies in bootstrapping capability under data, compute, and evaluation constraints. The central concern is whether AI can contribute to its own development in a controlled, measurable, and sustainable way.

Taken together, these studies offer a unified picture of AI for value. Across productivity, science, and AI for AI, the same questions recur: what can be delegated, what must remain under human control, how system behavior can be understood, how long-horizon reliability can be achieved, and how outcomes can be verified, audited, and improved. We aim to provide not only a map of methods, but also a framework for understanding how AI capability becomes durable real-world value.

Contents

Series Preface	1
1 Introduction	3
2 Conceptual Lens: Productivity and Productive Instruments	4
2.1 Conceptual Background: Productivity and Productive Instruments	4
2.2 Agentic AI: From Capability to Productivity	5
3 Framework: Mechanisms and Delegation Stages	8
3.1 Definition of AI for Productivity	8
3.2 Delegation Stages from a Productivity Perspective	9
4 Evidence: From Deployed Systems to Productivity Measurement	11
4.1 Deployed Systems as Evidence of Productivity Conversion.	11
4.2 Measuring Productivity: From Capability Scores to Economic Work	13
4.3 Research Scope to Industry Diffusion	13
5 Nascent-Diffusion Industries	15
5.1 Manufacturing: Bounded Automation in Safety-Critical Production	15
5.2 Agriculture: Decision Support under Environmental Uncertainty	16
5.3 Real Estate: Document Intelligence without Transaction Autonomy	18
5.4 Government and Public Administration: Public-Service Support	19
5.5 Agentic AI as Bounded Assistance in Early-Diffusion Domains	20
6 Growing-Diffusion Industries	22
6.1 Trade: Controlled Execution in Compliance-Heavy Workflows	22
6.2 Media: Workflow Acceleration under Editorial Accountability	23
6.3 Education: Scalable Personalization under Pedagogical Governance	25
6.4 Law: Procedural Automation under Professional Responsibility	26
6.5 Agentic AI as a Bounded Executor in Growing-Diffusion Domains	27
7 Widespread-Diffusion Industries	29
7.1 Information Technology: Executable Feedback and Autonomous Software Workflows .	29
7.2 Finance: Productivity Amplification under Risk Amplification	30
7.3 Medicine: Auditable Human-AI Collaboration in High-Stakes Care	32
7.4 Agentic AI as Scaled Orchestration in Widespread-Diffusion Domains	33
8 Cross-Industry Comparison: What Determines Productivity Diffusion?	36
9 Discussion and Challenge	36
9.1 Other Industries	36
9.2 Remaining Challenges on Productivity Diffusion	38
10 Conclusion and Further Outlook	41




Figure 1 | Overview of how agentic AI reshapes productivity through two mechanisms and four stages of delegated production responsibility. Using this framework, we assess how agentic AI is currently embedded across nine industries and extract cross-sector insights about where productivity gains are already visible, where they remain constrained, and what conditions are required for deeper diffusion.

1. Introduction

“What distinguishes the different economic epochs is not what is produced, but how it is produced, by what instruments of labour.”

— Karl Marx, *Capital: A Critique of Political Economy*

Across diverse traditions of economic thought—from classical political economy (Smith, 1937) and Marxian political economy (Marx, 1867) to modern innovation economics (Romer, 1989)—technological breakthroughs have long been recognized as defining the boundaries of economic eras. From the steam engine of the Industrial Revolution, to electrification enabling mass production, to computers extending the limits of cognitive labor, each technological leap has redefined the frontier of human capability.

With the evolution of Large Language Models (LLMs), AI agents are demonstrating a growing capacity to execute complex tasks and deeply engage with the real world (Chen et al., 2024c; Chhikara et al., 2025; Erdogan et al., 2025; Packer et al., 2023; Xiong et al., 2025b; Yan et al., 2025). Emerging AI agent systems such as OpenClaw (openclaw, 2026) and Claude Code (Anthropic, 2025b) already illustrate their growing impact on practical workflows and developer ecosystems. This shift not only introduces a new technological paradigm, but also establishes agents as a new class of instruments of production. Compared with tools of previous eras, AI agents exhibit two fundamentally transformative properties: **generality** and **autonomy**. These enable them to operate across task boundaries and complete an autonomous closed loop—from environment perception

and goal decomposition to decision-making and execution—without continuous human intervention (Hao et al., 2025; Xi et al., 2025), unlike traditional machines and software constrained by predefined rules and domain-specific programming (Russell and Norvig, 2020).

Agentic AI should therefore be understood as a new productive instrument. Its importance lies not only in improving isolated task performance, but in changing how work is executed, coordinated, and delegated. The central question of this paper is not whether agentic AI can improve isolated task performance, but under what conditions technical capability is converted into sustained productivity gains.

We argue that this conversion depends on two mechanisms and one delegation process. The **efficiency mechanism** reduces the human effort required for existing output by automating parts of complex workflows. The **expansion mechanism** makes valuable work feasible that would otherwise remain too costly, slow, or complex. These mechanisms become economically consequential only when production responsibility is gradually delegated from humans to AI systems. We therefore identify four stages of delegation: conversational assistant, reactive operator, adaptive coordinator, and self-governing production system.

Using this framework, we assess how agentic AI is currently embedded across nine industries and extract cross-sector insights about where productivity gains are already visible, where they remain constrained, and what conditions are required for deeper diffusion. Our contribution is not only to map deployed systems and benchmarks, but to explain why productivity gains are uneven across sectors. The evidence suggests that deeper diffusion depends on the digitization of production objects, feedback verifiability, error tolerance, workflow modularity, and accountability structures.

The remainder of the paper advances this argument as follows.

- **Section 2** establishes the conceptual lens by defining productivity as the outcome of how human labor is reorganized through productive instruments.
- **Section 3** defines the productivity mechanisms and delegation stages that connect technical capability to workflow-level productivity.
- **Section 4** uses deployed systems and benchmarks as evidence for how agentic AI begins to convert capability into economic work.
- **Sections 5–7** compare diffusion patterns across industries, distinguishing bounded assistance, bounded execution, and scaled orchestration.
- **Section 8** identifies structural constraints that explain why productivity diffusion remains uneven.
- **Section 9** synthesizes the implications for durable real-world value and delegated agency.

2. Conceptual Lens: Productivity and Productive Instruments

2.1. Conceptual Background: Productivity and Productive Instruments

This section introduces the conceptual lens used throughout the paper: productivity as the outcome of how human labor is reorganized through productive instruments. We draw on foundational concepts from economics and the history of technology not to advocate any particular political or ideological position, but to distill the notions of productivity and productive instruments into a structured lens for analyzing the economic impact of agentic AI. Productivity is commonly measured as output per unit of labor input, such as output per hour worked or GDP per hour worked (OECD, 2001). This definition is useful for empirical analysis, but it does not explain the mechanisms through which major productivity gains are generated. Economic growth theory has long emphasized technological

change as a central source of sustained productivity improvement (Solow, 1957). To understand where productivity growth comes from, we must therefore look not only at labor input, but also at the technologies and other non-human resources that shape the production process.

The concept of productive instruments provides a critical lens for understanding these mechanisms. In the classical political economy, production depends on the combination of human labor and the means of production (Marx, 1867). Within this broader category, productive instruments include the tools, machines, and technical infrastructures that directly shape the way work is carried out, coordinated, and scaled (Rosenberg, 1983). From this perspective, productivity depends not only on labor effort, but also on how labor is organized through these instruments. Therefore, changes in productive instruments often lead to broader changes in productivity.

In line with this framework, the major technological changes coincide with broader economic transformations (Bresnahan and Trajtenberg, 1995; Helpman, 1998). Throughout the history of industrial development, breakthroughs in key productive instruments have repeatedly expanded the frontier of what human societies could produce. The steam engine mechanized physical work and enabled the large-scale substitution of human and animal power (Helpman, 1998). Electrification and assembly-line systems reorganized production around standardized, continuous, and high-throughput manufacturing (Atkeson and Kehoe, 2001). Later, computers and digital networks transformed how information was processed, calculations were performed, and communication and coordination were organized, extending technological augmentation from manual labor to cognitive, administrative, and organizational work (David, 1990). In each case, the importance of the new technology did not lie only in its technical novelty. More importantly, it changed how labor, capital, and coordination were organized at scale.

Despite their significance, a common feature of earlier productive instruments is that they mainly amplified specific aspects of human work rather than supporting broad execution across tasks. Industrial machines primarily strengthened and accelerated physical execution, while information systems and conventional software improved the speed, accuracy, and scale of calculation, storage, and rule-based processing (Autor, 2015; Brynjolfsson and McAfee, 2014). These technologies generated substantial productivity gains, but their contribution was usually constrained by relatively fixed operating logics. Most were highly effective within predefined scopes and specific domains, while still depending on humans to set goals, coordinate workflows, and handle unexpected cases.

This distinction is central to understanding the potential economic role of agentic AI. Historically, the economic significance of novel technical systems lies not only in improving task performance, but also in reshaping the structure of work (Acemoglu and Restrepo, 2018; Bresnahan and Trajtenberg, 1995). In the next subsection, we apply this lens to agentic AI and examine how recent advances in capability, cost, benchmarking, and deployment have made it a potential driver of productivity.

2.2. Agentic AI: From Capability to Productivity

Within the framework mentioned above, agentic AI can be understood as a new productive instrument. Unlike earlier digital tools that mainly supported information retrieval, rule-based processing, or isolated task assistance, agentic AI can increasingly participate in longer and more complete workflows. Built on large language models and related foundation-model capabilities, these systems can interpret task contexts, reason about goals, decompose problems into executable steps, invoke external tools, and interact iteratively with dynamic environments to complete non-trivial objectives (Chen et al., 2024c; Chhikara et al., 2025; Erdogan et al., 2025; Packer et al., 2023; Xi et al., 2025; Xiong et al., 2025b; Yan et al., 2025; Zhou et al., 2026). Emerging systems such as OpenClaw (openclaw, 2026) and Claude Code (Anthropic, 2025b) already illustrate their growing role in practical

workflows and developer ecosystems.

Compared with tools of previous eras, agentic AI exhibits two fundamentally transformative properties: **generality** and **autonomy**. Generality refers to the ability to operate across diverse tasks, domains, and knowledge contexts. While traditional software is highly effective within predefined procedures, agentic AI can be repurposed across linguistic, analytical, and operational workflows, making it a reusable asset across the production pipeline. Currently, autonomy enables these systems to complete closed-loop workflows with minimal human intervention. Instead of merely executing isolated commands, agentic systems maintain context, decompose goals, invoke tools, and adapt to dynamic feedback. This shifts their role from assisting with fragmented micro-tasks to taking ownership of complete units of work. Together, these properties make agentic AI a dual engine of productivity: an **efficiency mechanism** that compresses complex workflows, and an **expansion mechanism** that enables previously unfeasible forms of human-machine coordination.

Two primary technological drivers enable this shift from capability to productivity. First, the capabilities of agents have improved substantially. Modern agents can perform long-horizon planning with feedback-driven revision, conduct multi-step reasoning, maintain structured memory across interactions, decompose complex goals into executable sub-tasks, and reliably invoke external tools within dynamic environments (Chen et al., 2024c; Erdogan et al., 2025; Park et al., 2023; Yan et al., 2025; Yao et al., 2023). Second, the cost of unit inference has decreased significantly (Dao et al., 2022; Maslej et al., 2025). Improved capabilities make it possible for agents to carry out more complex tasks in a stable manner, while lower inference costs make repeated and large-scale deployment economically feasible. Together, these changes allow agentic AI to move beyond experimental demonstrations and become part of real production workflows.

Recent macro-level signals also point to rapid expansion in investment and market expectations. The Stanford AI Index 2025 reports that U.S. private AI investment reached \$109.1B in 2024, while global private investment in generative AI totaled \$33.9B, representing an 18.7% year-over-year increase (Maslej et al., 2025). Looking ahead, International Data Corporation (IDC) projects that global AI spending will exceed \$632B by 2028 (International Data Corporation (IDC), 2024), and Bloomberg Intelligence estimates that the generative AI market could approach \$1.3T by 2032 under rapid growth assumptions (Singh et al., 2024). Although investment is not itself productivity, these signals indicate that firms increasingly expect AI systems to become economically consequential production technologies rather than marginal digital tools.

In parallel with these investment trends, the evaluation of AI systems is also shifting from narrow capability benchmarks toward productivity-oriented assessment. New benchmarks such as GDPval replace automatic matching scores with blind pairwise comparisons by domain experts, who evaluate AI and human outputs on multi-hour, multi-document tasks that resemble real office work (Patwardhan et al., 2025). Long-horizon agent environments test whether agents can sustain performance over extended plan-and-execute scenarios. For example, Vending-Bench evaluates agents running a small business over many simulated days (Backlund and Petersson, 2025), while EcoGym tests agents in interactive economies such as freelance work and business operations over long horizons (Hu et al., 2026). Other efforts directly examine the automation potential of high-value professional tasks. For instance, the AI Productivity Index (APEX) measures performance in consulting, finance, law, and medicine (Vidgen et al., 2025). At the same time, the Remote Labor Index and RE-Bench focus on remote knowledge work and research tasks (Mazeika et al., 2025; Wijk et al., 2025). Taken together, these evaluations show a clear shift from measuring whether AI can perform narrow tasks to measuring whether it can reliably deliver economically valuable work at or above professional quality.

Early enterprise-level evidence also suggests measurable performance gains. Based on time esti-

mates from approximately 100,000 AI-assisted conversational tasks, Anthropic reports that, for some tasks that take around 90 minutes without assistance, AI support can reduce completion time by approximately 80% (Tamkin and McCrory, 2025). In knowledge-intensive services, a joint experiment by BCG and Harvard Business School finds that consultants using AI completed tasks 25.1% faster while achieving a 40% improvement in quality scores (Dell’Acqua et al., 2023). At the macro level, extrapolations under widespread adoption scenarios project that U.S. labor productivity growth could increase by 1.8 percentage points annually over the coming decade, with corresponding increases in Total Factor Productivity (Tamkin and McCrory, 2025). Long-run GDP gains are estimated to range from 1.5% by 2035 to 3.7% by 2075 (Arnon, 2025).

At the same time, the translation from technical capability to sustained productivity is not automatic. The economic impact of agentic AI depends not only on model performance and inference cost, but also on how these systems are embedded in organizational structures, decision processes, and existing workflows (Bresnahan et al., 1999; Brynjolfsson et al., 2018). In many settings, adoption remains gradual, and human oversight still plays a central role (David, 1990; Raisch and Krakowski, 2021). Understanding how micro-level performance improvements become broader economic value therefore requires more than isolated case evidence. It requires a structured account of how responsibility shifts from humans to AI systems, how this transfer changes workflow design, and how these patterns differ across industries.

The key analytical problem is therefore a conversion problem: when and how do model capabilities become durable productivity gains in real workflows? This question links the conceptual lens above to the industry analysis that follows, where the same technical capability can produce very different productivity outcomes depending on workflow embedding, verifiability, error tolerance, and institutional accountability.

In the following sections, we formalize these productivity gains and introduce a four-stage framework to describe the transfer of agency from humans to AI systems. We then use this framework to examine the current pattern of industrial diffusion and explain why the level of AI autonomy differs across domains.

3. Framework: Mechanisms and Delegation Stages




Figure 2 | An illustration of the two mechanisms. Efficiency Mechanism: It enhances productivity by reducing time and cost for existing work. Expansion Mechanism: It expands the task scope and unlocks the unseen work to foster productivity.

3.1. Definition of AI for Productivity

In this paper, **AI for Productivity** refers to the use of agentic AI in real workflows to either reduce the human effort required for existing output or make valuable work feasible that would otherwise remain too costly, slow, or complex. Like earlier productivity technologies, Agentic AI can raise productivity through two mechanisms across industries. It can reduce the time and labor required for existing work. It can also make economically valuable work feasible that would otherwise remain undone because it is too costly or too complex, thereby further freeing human capacity, as we discuss next.

The Efficiency Mechanism. The Efficiency Mechanism is based on reducing the time and labor required for existing work. It applies when the cost of running an agent system for a task becomes lower than the cost of human time for the same task. In that case, organizations can automate parts of existing workflows and complete the same work faster and at lower cost (Acemoglu, 2024). This leads to immediate and measurable improvements in execution speed through task acceleration and the automation of repetitive work. Empirical evidence supports sizable time savings from Agentic AI assistance in knowledge work (Dell’Acqua et al., 2023; Peng et al., 2023; Tamkin and McCrory, 2025). Research from Anthropic indicates that Agentic AI can accelerate individual tasks by up to 80% (Tamkin and McCrory, 2025), particularly in knowledge-intensive sectors. At the macro level, Wharton’s Penn Wharton Budget Model (PWBM) projects that this efficiency boost will increase aggregate GDP and productivity levels by 3.7% by 2075 (Arnon, 2025), generating a transient but powerful wave of Total Factor Productivity (TFP) growth. Under this mechanism, productivity is gained by doing the same economic tasks, but with substantially lower human expenditure.

The Expansion Mechanism. While efficiency is about reducing the time and labor required for existing work, the Expansion Mechanism is about unlocking work that would otherwise remain undone. It relies on the autonomy and scalability of agentic AI to make previously costly or complex tasks feasible. Recent studies suggest that a significant share of AI-assisted work would not have been




Figure 3 | Four stages of AI reshaping productivity.

done at all without AI assistance (Kalliamvakou, 2022; Singla et al., 2025). For example, in Anthropic’s internal study, about 27% of AI-assisted tasks were described by users as work they “would not have done at all” without LLMs (Huang et al., 2025). Such work often includes exploratory iteration, large-scale engineering completion, and many long-tail improvements that are valuable but usually left undone because teams lack time or attention (Huang et al., 2025; Peng et al., 2023). In practice, reductions in time and labor costs can also make this broader scope of work easier to sustain and scale. Therefore, “AI for Productivity” should not be understood only as automation or substitution. Its more important long-run effect may come from reallocating saved human attention and labor toward higher-value innovation, decision-making, and service, which can drive continuous improvement in both products and workflows (Brynjolfsson et al., 2025).

3.2. Delegation Stages from a Productivity Perspective

To systematically examine how AI reshapes productivity across industries, we classify agentic AI deployment by the depth of production responsibility delegated from humans to AI systems. These stages do not simply describe how AI systems evolve; they describe how execution, coordination, monitoring, and eventually strategic control are redistributed within production workflows.

Stage 1: AI as Conversational Assistant. Stage 1 marks the baseline of AI integration, where AI systems function as responsive, informational instruments. At this stage, AI leverages its generality to assist with information retrieval, drafting, and preliminary analysis under single-turn instructions. The productivity gains are primarily driven by the *Efficiency Mechanism*, where AI accelerates the early phases of work by providing initial ideas.

However, as a conversational assistant, AI cannot interact with the external environment, maintain persistent state, or participate in the actual execution of tasks. Furthermore, its preliminary suggestions are often incomplete or unreliable. The productivity ceiling of this stage is thus defined by the absolute reliance on manual execution and output validation. Humans must fully take over to verify the accuracy of the AI’s advice, translate these rough ideas into concrete actions, and manually drive the entire production loop from start to finish.

Stage 2: AI as Reactive Operator. Stage 2 advances to operational automation, where AI transitions from a passive assistant to a reactive digital operator. At this stage, systems can autonomously execute narrowly scoped operations such as structured data manipulation, code generation, or text refinement. The *Efficiency Mechanism* deepens as productivity gains shift from preparatory assistance to the direct substitution of repetitive, low-level execution. By automating this operational labor, AI

deeply integrates into the individual steps of the workflow.

Yet, systems in Stage 2 remain severely constrained by their limited contextual scope and inability to navigate dynamic environments. They cannot decompose complex goals or seek alternative solutions when confronted with non-standard inputs and errors. The productivity ceiling of this stage is defined by the constant need for human oversight and error correction. While humans no longer perform the discrete tasks themselves, they must constantly intervene to provide highly specific instructions, check intermediate outputs, stitch together isolated AI outputs, and fix operational mistakes to prevent the workflow from breaking.

Stage 3: AI as Adaptive Coordinator. In Stage 3, AI acts as an adaptive coordinator. Beyond simple execution, these systems decompose complex goals into workflows, utilizing persistent memory and dynamic tool-integration to manage extended tasks and recover from failures. At this stage, the *Efficiency Mechanism* deepens as AI automates procedural coordination, reducing marginal costs and scaling individual capability to a team level. By removing high coordination overhead, AI unlocks complex economic activities that were previously unfeasible or uneconomical. This shift in efficiency thus enables the *Expansion Mechanism* take effect.

Despite these advancements, AI in Stage 3 remains bounded by explicitly provided contexts and predefined success criteria. While it can efficiently execute a well-structured workflow, it cannot understand implicit intentions or autonomously structure ambiguous concepts into solvable projects. The productivity ceiling of this stage is therefore defined by the inherent burden of problem scoping and task specification. Humans must still manually translate open-ended goals into detailed execution blueprints and define precise requirements before the AI can take over the coordination.

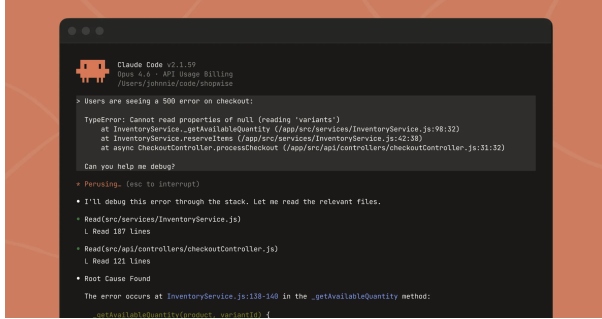
Stage 4: AI as Self-Governing Production System. In Stage 4, AI evolves into an autonomous self-governing production system. Moving beyond task execution, these systems translate abstract, outcome-oriented goals into structured subgoals through multi-dimensional trade-offs. They operate in open environments, maintain consistent internal state, adapt to dynamic constraints, and orchestrate multiple agents as a unified production ecosystem. Building on further improvements in the *Efficiency Mechanism*, the *Expansion Mechanism* reaches its full potential. By enabling scalable, asynchronous decision-making beyond human attention limits, AI autonomously organizes activities that were previously too complex or costly to manage, effectively expanding the production frontier.

However, giving AI this level of autonomy introduces new bottlenecks related to goal alignment and data limitations. Since the system relies solely on quantifiable information, it remains blind to complex human factors and unpredictable real-world events. The productivity ceiling of this stage is defined by the systemic overhead of strategic governance and alignment verification. While eliminating human execution and coordination provides massive productivity gains, it creates a new requirement for humans to monitor AI decisions, build safety measures, and ensure that high-level strategies remain aligned with human goals.


Taken together, these four stages trace a progressive delegation of production responsibility from humans to AI systems. Across this trajectory, the *Efficiency Mechanism* continuously deepens to optimize existing economic activities, which ultimately allows the *Expansion Mechanism* to take effect, enabling AI to actively expand the boundaries of industrial capability.

4. Evidence: From Deployed Systems to Productivity Measurement

4.1. Deployed Systems as Evidence of Productivity Conversion.



(a) Claude Code (Anthropic, 2026f)



(b) Claude Cowork (Anthropic, 2026b)

Figure 4 | Evidence from AI Productivity Tools. Left: Claude Code autonomously debugging a local repository by tracing error stacks and navigating the source code within a terminal environment. Right: Claude Cowork synthesizing fragmented data from local CSV and JSON files into a performance report with visualizations, executed as a scheduled task to ensure continuous, automated reporting on a predefined cadence.

To illustrate how theoretical stages of agentic AI are converted into real-world productivity gains, we examine three pioneering AI productivity tools: Claude Code (Anthropic, 2025b), Claude Cowork (Anthropic, 2026b), and OpenClaw (openclaw, 2026). These cases are not merely examples of product functionality. They show how workflow context determines whether agentic capability becomes sustained productivity: executable feedback in software engineering, data-workflow access for non-technical workers, and persistent background operation for individual users.

Claude Code. Claude Code shows why software is an early frontier for agentic productivity: the work is digital, feedback is executable, and iteration loops can be closed inside the development environment. As a highly capable autonomous programming agent designed for complex software engineering, it is integrated deeply into native development environments and functions not only as an assistant, but as a powerful, composable command-line engine. It uses persistent memory to actively master project-specific architectures and user conventions, with reusable skills enabling it to leverage prior experience for continuous self-evolution. It also supports multi-agent collaboration for robust, parallel task execution across large-scale systems and balances autonomy with strict permission controls to prevent unauthorized local file modifications.

These capabilities enable Claude Code to operate through both efficiency and expansion mechanisms. Under the *Efficiency Mechanism*, it drastically accelerates the software development lifecycle. For example, Rakuten reported that Claude Code performed 7 hours of continuous autonomous coding with 99.9% accuracy on complex refactoring tasks, reducing the time required to release new features by 79% (Anthropic, 2026c).

Furthermore, it triggers the *Expansion Mechanism* for both engineers and non-engineers. For engineers, it activates large-scale projects that might otherwise be delayed or abandoned due to resource constraints. At Wiz, a 50,000-line Python library was migrated to Go in 20 hours, a task originally estimated to take several months. According to the company’s leadership, this project would not have been pursued without Claude Code (Anthropic, 2026d). Additionally, the system lowers the barrier to software creation for non-technical roles, such as founders and product managers, by allowing

them to develop functional prototypes from natural language descriptions.

Consequently, Claude Code not only accelerates existing software development processes, but expands the boundaries of who can create software and what projects can realistically be undertaken.

Claude Cowork. Claude Cowork illustrates how agentic AI expands productivity by giving non-technical workers access to data workflows that previously required specialized technical labor. Claude Cowork ([Anthropic, 2026b](#)) applies the agentic architecture of Claude Code to non-technical knowledge workflows. Designed as a desktop-based assistant for enterprise productivity, it can produce professional deliverables such as spreadsheets and slides that can be further edited with Excel and PowerPoint. Users can also automate recurring duties via scheduled tasks and utilize project workspaces to provide isolated environments for distinct tasks.

These capabilities allow Cowork to automate high-effort and repeatable work, such as organizing local files, preparing meeting materials, and managing schedules, thereby improving productivity through the *Efficiency Mechanism*. At Airtree, Cowork reduced market and competitor research from two days to minutes and cut board meeting summaries from 30 minutes to two, allowing the team to spend more time analyzing insights and refining their strategy rather than creating reports ([Anthropic, 2026a](#)).

Furthermore, it triggers the *Expansion Mechanism* by empowering non-technical professionals to build systems that support their decision-making process. At Zapier, generating a unified visualization of influencer marketing data was previously unfeasible, as critical metrics like investment data and real-time engagement were scattered across multiple technical systems. The required SQL complexities and cross-platform integrations fell beyond the marketing team’s technical scope, forcing them to rely on delayed, retrospective reporting. With Cowork, a non-technical marketing leader independently built a live dashboard that synthesizes these fragmented data streams into a daily-updated visualization. This transition from retrospective analysis to forward-looking guidance allows the team to proactively reallocate budgets and prioritize high-performing channels ([Anthropic, 2026e](#)).

In this way, Cowork not only accelerates existing workflows but also creates new analytical practices and decision-making processes that were once precluded by technical barriers.

OpenClaw. OpenClaw highlights a shift from one-time assistance to persistent background operation, where productivity gains come from continuous monitoring and autonomous task initiation. Beyond professional scenarios, frameworks like OpenClaw ([openclaw, 2026](#)) exemplify a shift toward deeply personalized autonomous agents. Deployed directly on local machines, OpenClaw securely accesses private documents and interacts with native applications to establish a highly specific personal context. The system incorporates a persistent memory architecture, allowing the agent to adapt to user habits and historical data over time. Instead of reacting passively, OpenClaw monitors system states periodically via a heartbeat mechanism, initiates tasks autonomously, and supports recurring operations through cron-based scheduling. Users can monitor and direct this persistent local engine through integration with standard messaging platforms, including WhatsApp, Telegram, and Discord.

These capabilities enhance productivity through the *Efficiency Mechanism* by reducing the need for active human monitoring. For example, by leveraging the heartbeat mechanism, OpenClaw can continuously monitor deployed local services or long-running machine learning training tasks. When an anomaly or crash occurs, OpenClaw autonomously diagnoses the issue and executes necessary repairs, entirely without human intervention. The user simply receives a brief recovery notification

via Telegram, transforming tedious troubleshooting into seamless automation.

At the same time, OpenClaw drives the *Expansion Mechanism* by allowing solo professionals to deploy continuous, enterprise-grade operational systems from their personal computers. For example, a solo entrepreneur can establish a 24/7 market intelligence pipeline that autonomously scrapes industry databases overnight, filters massive datasets based on the user’s highly specific business context, and pushes actionable, synthesized leads directly into a WhatsApp chat by morning. By transforming a static local machine into an autonomous digital back-office, OpenClaw empowers individuals to execute data-intensive, continuous ventures that would previously require hiring a dedicated team of analysts.

Together, these case studies illustrate how agentic AI translates autonomy into real productivity gains. Across domains, these tools show that agentic AI not only accelerates existing workflows but also expands the scope of what individuals and organizations can achieve.

4.2. Measuring Productivity: From Capability Scores to Economic Work

The emergence of productivity-oriented benchmarks indicates a broader measurement shift. AI systems are increasingly evaluated not only by whether they answer correctly, but by whether they can complete economically valuable work over realistic time horizons. They tend to take one of two stances. The first targets the economic value an agent can produce. The second measures the length of task an agent can autonomously complete. Several benchmarks adopt the first stance. OpenAI’s GDPval scores models head-to-head against industry experts on 1,320 tasks drawn from 44 occupations across the nine largest sectors of U.S. GDP (Patwardhan et al., 2025). Mercor’s APEX uses rubric-scored expert tasks in investment banking, consulting, law, and medicine, with an agentic variant (APEX-Agents) for multi-step execution (Vidgen et al., 2025). \$OneMillion-Bench prices 400 expert-curated tasks at real market wages across five professional domains (Yang et al., 2026b). Xbench frames productivity as Technology-Market Fit, weighting performance by cost and market demand in verticals like recruitment and marketing (Chen et al., 2025a). Vending-Bench asks agents to run a simulated business for a year and scores them solely on dollars earned (Backlund and Petersson, 2025). The second stance is best represented by METR’s time horizon (Kwa et al., 2026), which collapses productivity to a single scalar: the length of human-expert work an agent can complete with 50% success. RE-Bench (Wijk et al., 2025) takes a similar approach for ML research engineering tasks, evaluating agents against 8-hour human-expert attempts.

Constructing such evaluations is resource-intensive. Each requires domain-expert knowledge. GDPval’s tasks come from professionals averaging 14 years of experience. APEX hires former Goldman analysts and big-law associates. \$OneMillion-Bench took over 2,000 hours of expert collaboration. Each also requires long task execution. A single Vending-Bench run consumes 60 to 100 million tokens, and METR tasks range up to 16 hours of human time. Indeed, for most industries, the role of agentic AI in driving productivity remains without formal evaluation. Benchmarks capture only part of productivity because they are better at measuring task performance, economic output proxies, and time horizons than workflow embedding, responsibility delegation, and institutional accountability. The industry analysis therefore complements benchmark evidence by assessing where agentic AI is actually embedded in workflows and how much production responsibility has been delegated.

4.3. Research Scope to Industry Diffusion

Our goal is not to exhaustively catalog every AI application in each sector. Instead, we use representative academic studies, deployed systems, benchmarks, and industry evidence to assess the depth of agentic AI integration.

We adopt the same industry list identified in the OpenAI GDPval benchmark (Patwardhan et al., 2025), which focuses on sectors with the largest contributions to U.S. GDP and high-earning, predominantly digital occupations. We then categorize these industries into three groups by assessing whether AI is embedded in real workflows, whether it can execute multi-step tasks, whether it uses external tools and feedback loops, whether human responsibility is partially transferred to AI systems, and whether observed productivity gains primarily reflect efficiency or expansion.

Specifically, in Section 5, we discuss **nascent-diffusion** industries, including *manufacturing, agriculture, real estate, and government*, which remain largely concentrated in Stage 1 and early Stage 2. In Section 6, we examine **growing-diffusion** industries such as *trade, media, education, and legal services*, which increasingly demonstrate Stage 2 and partial Stage 3 capabilities. In Section 7, we analyze **widespread-diffusion** industries including *information technology, financial services, and medicine*, which show the most active exploration of Stage 3 autonomy and early Stage 4 traits in constrained settings.

5. Nascent-Diffusion Industries

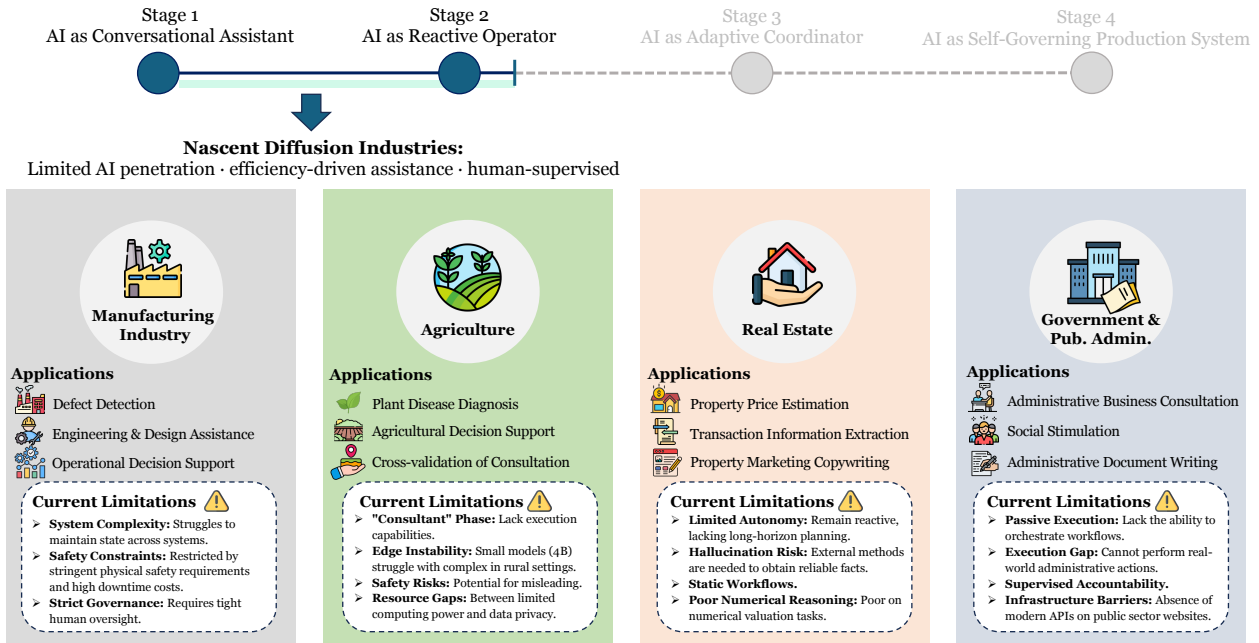


Figure 5 | Nascent Diffusion Industries at Stages 1–2 of AI Adoption. These sectors share traits of limited AI penetration, efficiency-driven assistance, and reliance on human supervision. For each, the figure identifies representative applications alongside key limitations, such as system complexity, hallucination risk, execution gaps, and infrastructure barriers, that hinder progression toward more autonomous Stage 3 and 4 deployments.

5.1. Manufacturing: Bounded Automation in Safety-Critical Production

Manufacturing represents a high-risk physical system characterized by strong production continuity, high downtime costs, and strict safety constraints. These conditions make it more feasible for Agentic AI to initially augment consultative and decision-support tasks rather than achieve end-to-end automation. Over the past decade, industrial AI has followed a clear trajectory: from localized perception and defect detection, to knowledge-enhanced systems via Retrieval-Augmented Generation (RAG), and further toward integration with engineering toolchains and early multi-agent architectures.

Diagnosis and Engineering Support as the Entry Point. Early industrial AI was dominated by single-task discriminative models such as anomaly detection and predictive maintenance, constrained by limited abnormal samples and high false-positive costs. Representative methods include PaDiM (Roth, 2021) and PatchCore (Defard et al., 2020) on the MVTec AD benchmark (Bergmann et al., 2021). While these approaches achieved perception-level automation, they lacked engineering context and generalization across production systems. With the accumulation of enterprise knowledge (e.g., manuals, logs, and specifications), industrial AI evolved toward RAG-based paradigms that connect models with private data sources, improving traceability and reducing hallucinations. Systems such as Schneider Electric’s EcoStruxure Copilot (Electric, 2024) exemplify this trend, enabling natural-language fault analysis and maintenance recommendations. However, these systems remained primarily limited to operational support.

Controlled Agents for Partial Operational Workflows. As automation complexity increased, generative AI expanded upstream into engineering design. “Engineering copilots” integrated with industrial platforms began supporting code generation, documentation retrieval, and control logic explanation. For example, Siemens Industrial Copilot ([Siemens, 2024](#)) enables natural language–driven automation code generation and system configuration. Academic efforts further explore generating PLC programs and simulation models from natural language via retrieval-augmented and preference-aligned training. Despite these advances, such systems largely function as cognitive assistants, lacking continuous workflow awareness and the ability to execute cross-system tasks autonomously.

Recent developments extend AI from engineering to operations by integrating real-time industrial data. Systems such as ABB Genix Copilot ([ABB, 2024](#)) combine sensor data and historical records to interpret alarms and support decision-making. More advanced architectures, explored by companies like Honeywell ([Honeywell, 2025](#)) and Siemens ([Siemens, 2025](#)), introduce multi-agent coordination and tool invocation for partial task automation. However, due to strict requirements on safety, auditability, and control, these systems operate under human governance and remain at the stage of “interpretation–recommendation–partial automation,” lacking full autonomy.

Overall, from 2016 to 2025, manufacturing AI has evolved from perception-level automation to early-stage Agentic systems, progressively expanding its functional scope from defect detection and predictive maintenance to knowledge-grounded reasoning, engineering assistance, and partial workflow orchestration. Modern systems are increasingly capable of integrating multi-source industrial data, retrieving enterprise knowledge via RAG, assisting in code generation and system configuration, and providing interpretable recommendations for operational decision-making. Market forecasts further reflect this trajectory, with MarketsandMarkets projecting growth from USD 34.18B in 2025 to USD 155.04B by 2030 (CAGR 35.3%) ([MarketsandMarkets, 2025](#)), and Technavio estimating a 37.7% CAGR for generative AI in manufacturing ([Technavio, 2025](#)).

Why Production Control Remains Hard to Delegate. However, despite these advances, the autonomy boundary of industrial AI remains fundamentally constrained by physical safety requirements, production continuity, and the high cost of errors. Current systems still lack persistent state awareness, robust cross-system execution capability, and deterministic reliability in long-horizon tasks. As a result, most applications remain confined to roles such as interpretation, diagnosis, and recommendation, with only limited forms of tool-mediated automation. From the perspective of autonomy, manufacturing AI is still predominantly situated in AI-as-Conversational-Assistant and AI-as-Reactive-Operator regimes, where AI enhances human decision-making rather than independently controlling end-to-end production processes.

5.2. Agriculture: Decision Support under Environmental Uncertainty

Agriculture is a highly variable physical production domain shaped by environmental uncertainty, biological heterogeneity, and dispersed field conditions. These characteristics make it difficult for Agentic AI to move directly toward end-to-end autonomous execution in real farming scenarios. As a result, LLMs are increasingly applied in agriculture first for tasks such as yield prediction, disease diagnosis, and farm management, where they can support decision-making and selected operational steps. To overcome the domain knowledge limitations of general-purpose models, agricultural AI has evolved along a clear technical trajectory: from specialized unimodal LLMs to multimodal integrations and currently toward agentic systems augmented by tools. The following analysis assesses these advances as a transition from passive consultation to more capable but still bounded decision-support systems.

Agronomic Knowledge Support as the Entry Point. Specialized agricultural LLMs are transitioning from parametric adaptation toward dynamic retrieval and data synthesis. While early models like AgriBERT (Rezayi et al., 2022) and AgroGPT (Awais et al., 2025) leveraged domain-specific fine-tuning to improve terminology alignment, subsequent systems such as AgroLLM (Samuel et al., 2025) and ShizishanGPT (Yang et al., 2025d) introduced RAG to provide reliable, localized pesticide recommendations and precise plant disease diagnoses by dynamically retrieving the latest agricultural guidelines. More recently, IPM-AgriGPT (Zhang et al., 2025c) and AgriGPT (Yang et al., 2025c) have explored data-centric strategies, utilizing adversarial generation and multi-agent pipelines to synthesize high-quality agronomic corpora, addressing the critical scarcity of expert agricultural data. These advancements suggest that coupling data synthesis with dynamic retrieval is an efficient pathway for adapting general-purpose LLMs to the agricultural domain.

Multimodal and Tool-Augmented Agents in Field Workflows. Substantial research efforts have been directed toward extending large models to multimodal inputs, a critical requirement for visual and hands-free agricultural operations. Early vision-language initiatives, such as Agri-LLaVA (Wang et al., 2024a), focused on image-based crop disease diagnosis, while subsequent efforts like the AgriGPT-VL Suite (Yang et al., 2025a) further refined multimodal alignment via curriculum learning and policy optimization. To address the demands of real-world field environments, research has expanded into tri-modal architectures. For instance, AgriGPT-Omni (Yang et al., 2025b) integrates speech, text, and vision to handle complex field acoustics and dialects. These collective advancements represent a pivotal leap toward integrated multimodal interaction frameworks, significantly unlocking the latent productivity of digital agriculture.

To bridge the gap between passive information retrieval and actionable decision support, recent research efforts have integrated LLMs into specialized tool-use and multi-agent frameworks. Systems such as AgriDoctor (Zhang et al., 2025a) exemplify this trend by dynamically orchestrating vision models and knowledge graphs for precision management. Further research efforts, such as the AgMMU framework (Ke et al., 2025), utilize collaborative agent pipelines—comprising specialized Retriever, Reflector, and Answerer agents—to maintain reasoning consistency across diverse visual inputs. Moreover, this research trajectory is currently culminating in physical execution. Specifically, hierarchical planners such as AgriAgent (Yang et al., 2026a) and sensor-integrated systems like CottonBot (Kandamali et al., 2025) enable the automation of robotic workflows and irrigation, thereby signaling a definitive paradigm shift toward autonomous precision agriculture.

Why Farm-Level Autonomy Remains Fragile. Overall, agricultural AI has evolved from specialized LLMs to retrieval-enhanced systems, multimodal models, and early agentic architectures with tool-use capabilities. These systems increasingly combine agronomic knowledge, multimodal perception, and selected forms of planning or robotic support to improve diagnosis, recommendation, and partial workflow execution. However, agricultural generative AI still remains largely in *Stage 1 (Conversational Assistant)*, functioning primarily as advisory and diagnostic systems that require human intervention to translate model outputs into physical actions. Recent agentic frameworks suggest an emerging transition toward *Stage 2 (Reactive Operator)*, but current models still lack the robustness, long-horizon planning, and reliable decision-making needed for dynamic agricultural environments. Key bottlenecks include the reasoning limitations of small language models deployed at the edge and the lack of tightly integrated software and hardware infrastructures required for complex physical operations. As a result, most agricultural applications remain assistant-like or narrowly bounded operator systems rather than autonomous managers of end-to-end farm operations.

5.3. Real Estate: Document Intelligence without Transaction Autonomy

Real estate is a document-intensive and coordination-heavy domain that relies on fragmented information, localized market knowledge, client communication, and multi-party transactions. These characteristics make it more feasible for Large Language Models to first augment information processing and decision-support tasks rather than move directly toward end-to-end autonomous execution. In this context, research efforts have focused on integrating LLMs into traditional workflows to handle unstructured data and improve selected stages of real estate analysis, contract processing, and marketing. From a technical perspective, the evidence broadly reflects a progression from in-context learning, to supervised finetuning, and further toward Retrieval-Augmented Generation.

Documents and Market Information as the Entry Point. The initial application of LLMs in the real estate domain primarily leverages their intrinsic semantic understanding capabilities. As an early exploration, researchers have applied LLMs with in-context learning for housing price prediction. Although empirical evaluations demonstrate that isolated LLMs underperform traditional machine learning models on price prediction tasks (Geerts et al., 2025), these efforts establish a baseline for domain-specific deployment. Similar approaches have also been utilized to evaluate the attribute extraction performance of LLMs across various prompting configurations (eg. one-shot, few-shot, etc.) (Kvet et al., 2025). At this stage, however, LLMs function mainly as flexible advisors for narrow tasks.

Grounded Generation for Client-Facing Workflows. To handle domain-specific data processing needs, research has progressed toward utilizing fine-tuned LLMs for automated information extraction. These systems are designed to read unstructured real estate contracts and answer certain questions accordingly (Zhao and Gao, 2024). Compared to rule-based approaches, instruction-tuned models demonstrate superior robustness and reduce the need for manual feature engineering. This development expands the role of LLMs from general semantic assistance to more structured document understanding, but their function remains largely supportive, centered on contract reading and information processing rather than autonomous transaction execution.

Recent applications of Retrieval-Augmented Generation in real estate leverage retrieved data to enhance the persuasiveness of marketing content. For instance, the AI Realtor framework (Wu et al., 2025a) utilizes RAG to extract unexpected features associated with a property. This targeted grounding strategy effectively captures user interest and stimulates purchase desire, demonstrating RAG’s capability to generate highly engaging and persuasive marketing content. Even so, these systems remain concentrated in search, description, and marketing support, rather than in higher-stakes functions such as negotiation, compliance management, or multi-party coordination.

Why Transaction Responsibility Remains Human-Centered. Overall, real estate AI has evolved from early prompt-based semantic assistance to fine-tuned models for contract understanding and retrieval-grounded systems for marketing and client interaction. However, despite this technical progression from zero-shot inference to RAG-based systems, most applications still remain within *Stage 1 (AI-as-Conversational-Assistant)* or early *Stage 2 (AI-as-Reactive-Operator)*. Current systems can support tasks such as attribute extraction, document interpretation, price-related analysis, and grounded property description, but they remain largely reactive and lack persistent state, long-horizon planning, and autonomous task decomposition capabilities. In practice, the LLM functions mainly as an advanced automation component embedded within a fixed workflow, while core decisions such as pricing strategies, negotiation approaches, legal judgment, and transaction execution remain the responsibility of human actors. As a result, real estate AI still operates primarily in assistant-like roles focused on information processing, document understanding, and client-facing support, rather than

as an autonomous manager of end-to-end real estate transactions.

5.4. Government and Public Administration: Public-Service Support

Government and public administration constitute a high-stakes institutional domain shaped by legal accountability, procedural rigidity, public trust, and heterogeneous legacy systems. These characteristics make it difficult for large language model-based agents to move directly toward autonomous decision-making or delegated public authority. As a result, current research has more often introduced such systems first in roles related to regulatory retrieval, scenario simulation, and bounded administrative assistance. From a technical perspective, existing work broadly reflects a progression from retrieval-augmented systems for regulatory compliance, to multi-agent social simulation systems for policy modeling, and further toward operational agentic systems for selected administrative workflows.

Regulatory Retrieval and Administrative Drafting as the Entry Point. A prevalent line of research in this domain focuses on mitigating hallucinations and enhancing regulatory alignment through Retrieval-Augmented Generation (RAG) and knowledge graphs (KGs). For instance, systems such as GraphCompliance (Chung et al., 2025) and PrivComp (Garza et al., 2024) integrate unstructured regulatory text with structured logic, allowing the model to validate its outputs against specific legal clauses. Technically, these systems prioritize precision over autonomy, acting as advanced semantic search engines with limited generative agency.

Simulation and Browser Agents for Bounded Public Workflows. A second stream employs multi-agent systems to model social dynamics. Frameworks like GOVSIM (van Erven et al., 2025) utilize game-theoretic scenarios to simulate collective behavior and resource governance, offering insights for resource governance strategies under varying communication constraints. Similarly, Park et al. (2024) demonstrate that agents grounded on human interviews can better reflect public opinion for policy evaluation. Despite these improvements, existing research questions whether such agents are capable of reliable real-world decision-making, as they exhibit severe internal inconsistencies and misalignment with established evaluation frameworks (Pokharel et al., 2025). At this stage, therefore, multi-agent systems are more useful as exploratory or analytical components than as dependable substitutes for public judgment.

A third category encompasses agentic systems engineered to execute specific administrative workflows. For administrative document generation, Musumeci et al. (2024) employ a multi-agent framework that decomposes complex drafting tasks into collaborative sub-tasks for the production of administrative documents. Meanwhile, preliminary studies such as Agents4Gov (de Castro Oliveira et al., 2025) adapt conventional browser agents to accommodate the low-resource characteristics and absence of API functionality inherent to government website usage. These efforts extend agentic AI from policy support and simulation toward selected forms of workflow execution. However, the systems remain bounded to narrow administrative settings and may still require human oversight to ensure procedural correctness and operational reliability.

Why Public Authority Remains Hard to Delegate. Overall, government-facing AI has evolved from retrieval-grounded regulatory assistance to multi-agent policy simulation and early agentic systems for selected administrative workflows. Yet despite this architectural diversity, most applications still remain within *Stage 1 (AI-as-Conversational-Assistant)* and early *Stage 2 (AI-as-Reactive-Operator)*. Retrieval-augmented systems function mainly as fixed compliance-support pipelines rather than autonomous agents, while simulation agents serve as analytical components. Operational systems, such

as browser navigation and administrative document drafting, support only localized sub-task execution and remain unable to perform workflow orchestration, or delegated public administration. As a result, current government AI systems still function primarily as specialized tools under human supervision rather than as autonomous entities entrusted with public administrative processes.

5.5. Agentic AI as Bounded Assistance in Early-Diffusion Domains

Across manufacturing, agriculture, real estate, and government, Agentic AI has begun to create visible value, but its role in production remains primarily that of an *assistant* rather than an autonomous executor. Recent progress has introduced domain-specialized fine-tuned models, RAG-enhanced advisory assistants, and early tool-using or multi-agent systems into real workflows. Yet in most cases, these systems still operate by supporting human decision-making, providing recommendations, and participating only in a small number of narrowly bounded execution steps. Adoption is therefore already visible across a relatively wide range of tasks, but the *depth* of delegation remains shallow. These sectors are thus representative *nascent-diffusion industries*: Agentic AI is no longer confined to isolated predictive or assistive use, but its integration into production remains uneven, highly supervised, and largely concentrated in **Stage 1 (AI-as-Conversational-Assistant)** and early **Stage 2 (AI-as-Reactive-Operator)**.

At this stage, productivity gains arise overwhelmingly through the **Efficiency Mechanism**. By reducing human time in interpretation, drafting, diagnosis, and routine coordination, Agentic AI improves existing workflows without fundamentally reshaping their underlying control structure. By contrast, the **Expansion Mechanism** remains largely muted, because many tasks beyond routine human execution are constrained less by cognitive cost than by physical conditions, natural uncertainty, execution risk, and the need for direct intervention in real environments. This also explains why these sectors have not yet moved to the next stage of diffusion. Compared with digital-native domains, they are more tightly coupled to the physical world and less deeply embedded in fully digital operating environments, while placing much stricter demands on safety, low-risk execution, and verifiable and governable action. Since current agent still face limitations in persistent state, long-horizon reliability, and robust interaction with complex real-world production settings, broader delegation remains difficult to justify in practice. As a result, Agentic AI in these industries continues to function mainly as a supervised productive instrument for bounded assistance rather than as a fully autonomous production actor.

Table 1 | Representative domain-specific applications in Nascent-diffusion industries.

Paper	Task	Technique	Stage
<i>Manufacturing</i>			
PaDiM (Roth, 2021)	Defect detection	Anomaly detection	Stage 1
PatchCore (Defard et al., 2020)	Defect detection	Anomaly detection	Stage 1
EcoStruxure Copilot (Electric, 2024)	Fault diagnosis	RAG	Stage 1–2
Siemens Industrial Copilot (Siemens, 2024)	Engineering assistance	Tool-augmented LLM	Stage 2
ABB Genix Copilot (ABB, 2024)	Decision support	Industrial copilot	Stage 2
Honeywell AI Assistant (Honeywell, 2025)	Task automation	MAS	Stage 2–3
Siemens AI Automation System (Siemens, 2025)	Workflow automation	MAS	Stage 2–3
<i>Agriculture</i>			
AgroLLM (Samuel et al., 2025)	Management Q&A	RAG	Stage 1

Table 1 continued

Paper	Task	Technique	Stage
ShizishanGPT (Yang et al., 2025d)	Agricultural assistance	Tool-augmented LLM	Stage 1–2
Agri-LLaVA (Wang et al., 2024a)	Pest / disease diagnosis	Multimodal LLM	Stage 1
IPM-AgriGPT (Zhang et al., 2025c)	Pest / disease management	Contextual reasoning	Stage 1
CottonBot (Kandamali et al., 2025)	Irrigation advice	RAG + agentic LLM tools	Stage 2
AgriAgent (Yang et al., 2026a)	Task planning	Tool orchestration	Stage 2–3
<i>Real Estate</i>			
Utilizing Large Language Models for Information Extraction from Real Estate Transactions (Zhao and Gao, 2024)	Information extraction	SFT	Stage 1
Real Estate Attribute Value Extraction Using Large Language Models (Kvet et al., 2025)	Information extraction	Prompt engineering	Stage 1
AI Realtor (Wu et al., 2025a)	Automated copywriting	RAG	Stage 1
On the Performance of Large Language Models for Real Estate Appraisal (Geerts et al., 2025)	Price prediction	In-context learning	Stage 1
<i>Government</i>			
Cooperate or Collapse: Emergence of Sustainable Cooperation in a Society of LLM Agents (van Erven et al., 2025)	Social simulation	Simulation MAS	Stage 1
Generative Agent Simulations of 1,000 People (Park et al., 2024)	AI for social impact / social simulation	Simulation MAS	Stage 1
GraphCompliance (Chung et al., 2025)	Compliance	KG	Stage 1–2
PrivComp-KG (Garza et al., 2024)	Compliance	RAG + KG	Stage 1–2
LLM Based Multi-Agent Generation of Semi-structured Documents (Musumeci et al., 2024)	Semi-structured document generation	MAS	Stage 1–2
Agents4Gov (de Castro Oliveira et al., 2025)	Browser use	Single agent	Stage 1–2

6. Growing-Diffusion Industries

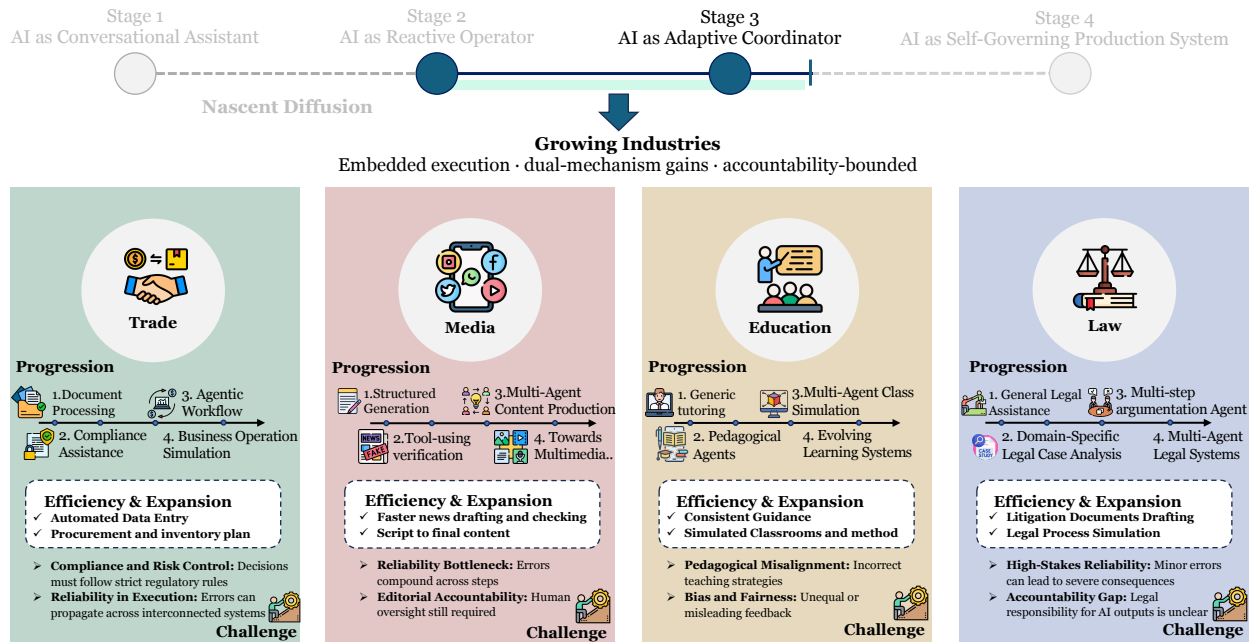


Figure 6 | Growing Industries at Stages 2–3 of AI Adoption. These sectors are characterized by embedded execution, dual-mechanism gains in both efficiency and capability expansion, and accountability-bounded deployment. For each industry, the figure traces a four-step progression from basic task assistance toward multi-agent and simulation-based workflows, identifies efficiency and expansion gains, and highlights the key challenges—including reliability bottlenecks, editorial and pedagogical accountability, bias and fairness concerns, and high-stakes responsibility gaps—that constrain progression toward deeper delegation.

6.1. Trade: Controlled Execution in Compliance-Heavy Workflows

Trade operations span multiple stages, including international commerce, customs clearance, logistics, and compliance. Core information is carried through documents such as invoices, packing lists, and bills of lading, which circulate across customs systems, ERP platforms, and regulatory infrastructures. Unlike content-generation scenarios, trade environments are highly standardized and low-fault-tolerant. Systems must therefore not only generate text but also extract structured fields, provide traceable evidence, and safely interface with downstream processes such as declaration, risk screening, and inspection handling. With the rise of digital trade, AI systems have further expanded toward operational tasks such as inventory management, procurement, and supply-chain coordination, requiring participation in business processes beyond document understanding.

Document Processing as the First Productivity Layer. Early systems were centered on OCR and intelligent document processing (IDP), transforming unstructured trade documents into structured data for entry, reconciliation, and validation (Microsoft Corporation, 2025). With the introduction of LLMs, systems evolved toward assisted decision-making with limited semantic understanding and evidence linkage. Multimodal document understanding models integrate text, layout, and visual features to support clause localization, field relation modeling, and cross-field consistency checking, reducing reliance on OCR text alone (Xu et al., 2020). Meanwhile, Retrieval-Augmented Generation (RAG) retrieves regulations and historical cases to provide traceable evidence, mitigating hallucina-

tion and knowledge drift. In practice, such systems remain positioned as expert-assistance tools in low-fault-tolerance environments. WCO discussions (Clicteur and Bertrand, 2023) show that AI-assisted classification reduces routine workload while supporting human review through confidence scores and candidate outputs. Industry frameworks increasingly integrate document processing, risk assessment, and classification into unified workflows spanning pre-declaration, clearance, and post-audit stages (PricewaterhouseCoopers (PwC), 2024). However, LLM outputs are typically constrained by evidence retrieval, rule validation, and human review, and remain limited in performing responsibility-bearing operations such as compliance decisions or system submissions (World Economic Forum, 2024).

Tool-Orchestrated Workflows for Trade Operations. AI systems are evolving toward agentic workflows that participate in business execution. Tasks such as retrieval, calculation, form filling, submission, and notification are decomposed into traceable tool-invocation chains forming closed-loop workflows (Mely.ai, 2025). Within this architecture, LLMs handle task understanding and orchestration, while high-risk decisions—such as tariff calculation and sanctions screening—are executed by deterministic modules. This ensures auditability and accountability. Examples include AWS supply-chain systems, where multiple agents coordinate order queries, inventory checks, and rule-based decisions (Pazak and Baeuml, 2025). In customs scenarios, IMF guidance emphasizes that GenAI should support explainable risk analysis while adhering to strict governance frameworks, with RAG as a core component (Aslett et al., 2025).

Why Compliance Responsibility Limits Full Automation. Recent benchmarks suggest improving agent capabilities in long-horizon trade operations. Vending-Bench (Backlund and Petersson, 2025) and its extensions evaluate inventory management, pricing, and profitability over extended horizons, while Anthropic’s Project Vend (Anthropic, 2025) demonstrates early automation in real environments. Multi-agent supply-chain simulations further show efficiency gains in coordination and cost reduction. However, these experiments also reveal instability in strategy and risk control, such as poor inventory management or irrational pricing. Consequently, real-world deployment remains constrained by rule engines, evidence retrieval, and human oversight to mitigate economic and compliance risks.

Overall, AI for Trade is transitioning from document automation to workflow orchestration. While LLM agents can execute multi-step tasks and participate in business processes under controlled settings, real-world systems still require human supervision due to regulatory and financial risks. Current systems are therefore positioned between Stage 2 and Stage 3, with a substantial gap remaining toward unrestricted trade automation.

6.2. Media: Workflow Acceleration under Editorial Accountability

Media production involves multiple linked stages, including reporting, verification, editing, packaging, and distribution. Unlike emerging domains that remain focused on isolated automation, recent work in media increasingly explores systems capable of producing long-form, structured outputs such as full news reports, complete films, or non-linear editing sequences. Research in this field has branched into three distinct paradigms: static workflows for structured generation, autonomous single agents for subtask completion, and multi-agent ecosystems for complex production management.

Structured Generation for Reporting and Content Production. One line of research focuses on utilizing large language models as modular processors within fixed, human-defined workflows, especially in text-based domains. For instance, news generation is often decomposed into fixed stages

(e.g. retrieval, semantic scoring, summarization) to produce high-quality long-form reports (Lin and Tsai, 2025). Similarly, fact-checking systems like SANCTUARY (Dharmavaram and Hakak, 2025) enforce a strict pipeline of claim generation, evidence retrieval, and ordered reasoning to ensure factual accuracy. In media analysis, the JOA-ICL framework (Lee et al., 2025a) first splits and evaluates structural segments of a full article, and aggregates them to infer the overall position, to enable more accurate article-level stance detection.

Agentic Verification, Editing, and Simulation. A second line of research equips agents with the autonomy to interact with external environments, such as search engines or knowledge bases, addressing the need for real-time accuracy and dynamic investigation. This approach is prominent in automated fact-checking, where static retrieval is insufficient. Agents are equipped with various tools, such as web search and source credibility assessment mechanisms, to verify dynamic and complex factual claims (Cui et al., 2025). As these autonomous systems become more sophisticated, evaluating their reliability is increasingly critical. To this end, unified frameworks like OpenFactCheck (Iqbal et al., 2024) provide comprehensive evaluation pipelines, allowing researchers to customize fact-checking agents and systematically benchmark their verification efficacy against standardized datasets.

Why Editorial Judgment Remains Difficult to Delegate. Multi-agent systems have emerged as a powerful paradigm for organizing complex media generation and evaluation. Within this paradigm, research primarily follows two complementary architectural logics: feedback-driven collaboration which refines outputs through repeated feedback loops between reviewer agents and generator agents, and role-based collaboration where multi-step tasks are completed through role-based division and joint contribution among all agents.

Feedback-driven Collaboration. In this paradigm, agents simulate user or expert responses to provide iterative feedback for generator agents or human writers to optimize content quality. In audiovisual media, video editing is modeled as an adversarial dialogue between an editor and a critic agent, allowing for continuous self-correction based on aesthetic feedback (Sandoval-Castaneda et al., 2025). In text-based media, the JRE-L framework (Jiang et al., 2025) iteratively generates and refines the generated news through simulated editorial and reader feedback loops, while the MADES framework deploys memory-augmented agents to debate perspectives and construct tailored supplementary contexts to enhance news understanding (Ouyang, 2025). Beyond content creation, this evaluative simulation also extends to media distribution, where personalized recommendation policies are optimized via simulated user interactions, eliminating the need for real-world A/B testing (Wang et al., 2024c; Zhang et al., 2024a).

Role-based Coordination. In this paradigm, agents collaborate through hierarchical, role-based task decomposition to divide complex tasks into executable sub-tasks. For instance, in audiovisual media generation, approaches like FilmAgent (Xu et al., 2025b) and Kubrick (He et al., 2024a) assign agents to specific film crew roles to generate and refine structured scripts and 3D rendering commands for production engines like Unity3D or Blender. Extending these preliminary movie generation approaches, MovieAgent (Wu et al., 2025b) first explores automated long-form video generation by introducing multi-agent Chain of Thought (CoT) planning, while AutoDirector (Ni et al., 2024) integrates multi-sensory elements such as long takes, special effects, soundtracks, voice dubbing, and lip-synchronization into automatic movie generation, which further enriches the expressiveness of the generated movies.

Collectively, media-oriented LLM agents remain between Stage 2 and the early phase of Stage 3. They can already support multi-step planning, role specialization, and bounded workflow coordina-

tion across journalism, recommendation, and audiovisual production, but their autonomy remains limited in real-world use. At present, these systems function primarily as sophisticated tools for accelerating specific sub-tasks, while final judgment, quality control, and release decisions remain under human supervision.

6.3. Education: Scalable Personalization under Pedagogical Governance

In education, AI agents are required to do more than retrieve facts: they must understand how to teach, guide students step-by-step, and support knowledge transmission through ongoing interaction, while also meeting stricter pedagogical and ethical requirements because of their potential influence on students' development and values. Accordingly, the following discussion unfolds across three dimensions: the transition from general-purpose models to specialized pedagogical agents, the enhancement of educational interaction through multi-agent architectures, and the evolving assessment and governance of agent autonomy.

Tutoring and Feedback as the Initial Productivity Gains. Early applications of generative AI in education often relied on general-purpose conversational models as virtual tutors (Swacha and Gracel, 2025). However, the educational domain presents unique and demanding challenges. Beyond information retrieval, effective teaching requires pedagogical intuition, goal decomposition, factuality, and socio-emotional support. General models often struggled with these domain-specific difficulties, frequently providing generic feedback or directly giving away answers rather than appropriately guiding the student. To address these fundamental gaps, the field shifted towards the development of specialized educational models, such as EduChat (Dan et al., 2023), Google LearnLM (LearnLM Team, Google, 2024), and Pedagogy-R1 (Lee et al., 2025b). By training extensively on pedagogical texts and real-world teaching interactions, these specialized agents are better equipped to handle nuanced tasks like Socratic questioning and complex grading. This shift allows human educators to focus on personalized interventions at a higher-level while verifying the accuracy, tone, and pedagogical alignment of AI-generated feedback (Chu et al., 2025a; Fajardo-Ramos et al., 2025).

Multi-Agent Interaction for Adaptive Learning Workflows. Beyond improving single models for discrete actions, researchers are leveraging multi-agent architectures to facilitate goal-oriented, multi-step reasoning. A primary application is simulation for safe training and strategy evaluation. For example, SocraticLM simulates complete teaching workflows to generate fine-tuning data (Liu et al., 2024), while Agent4Edu and Classroom Simulacra model virtual students to test pedagogy without human subjects (Gao et al., 2025; Xu et al., 2025a). To maintain output quality during extended conversations, systems such as AutoFeedback introduce inspector agents that improve feedback accuracy and reduce common quality issues (Guo et al., 2024). Furthermore, the increasing complexity of agent interactions requires more specialized evaluation frameworks. To address this specific demand, MathTutorBench introduces a reward mechanism that penalizes providing direct answers to encourage the model to offer effective guidance (Macina et al., 2025). Collectively, these multi-agent advances enable scalable, personalized learning pathways (Chu et al., 2025b; Nedungadi et al., 2024) and drive the transition of educational agents towards the Stage 3 (AI-as-Adaptive-Coordinator) paradigm.

Why Teaching Autonomy Requires Value Alignment. Educational agent deployment diverges from purely task-oriented systems by necessitating strict alignment with pedagogical paradigms, long-term learning objectives, and human-centric values. Consequently, research increasingly focuses on designing pedagogically aligned and trustworthy architectures. This involves integrating instructional frameworks—such as Socratic scaffolding—into agent workflows while mitigating critical risks like

answer leakage and bias amplification (Delikoura et al., 2025; Wang et al., 2024b). Future advancements demand the co-design of cognitive reasoning, interaction protocols, and ethical governance to ensure scalable automation remains anchored in fundamental educational principles (Chu et al., 2025b).

Overall, AI in education is moving from general conversational assistance toward more specialized and interactive support embedded in teaching and learning workflows. Current systems increasingly automate routine pedagogical tasks and enable more adaptive feedback, simulation, and personalized interaction, showing a gradual transition from Stage 2 toward early Stage 3. However, their role remains that of a supervised pedagogical assistant rather than an autonomous substitute for teachers, since real-world deployment is still constrained by the need for pedagogical alignment, value-sensitive guidance, and sustained human oversight.

6.4. Law: Procedural Automation under Professional Responsibility

Legal work is highly reasoning-intensive, procedurally structured, and closely tied to professional responsibility and institutional accountability. Early legal intelligence systems operated in a Conversational Assistant mode, where AI supported retrieval and drafting but remained fully dependent on user control. As legal workflows require reasoning across multiple procedural steps and related laws, systems began to incorporate structured reasoning and external knowledge grounding, reducing cognitive burden but still requiring human supervision. Recent developments in multi-agent coordination and procedural automation indicate a gradual shift toward bounded autonomy. This progression reflects not only technical refinement but also a change in how legal work is organized and delegated between humans and machines.

Legal Retrieval and Drafting as Efficiency Gains. Early work relied on adapting general-purpose language models to legal text. BERT-based models were mainly used for extraction and classification. Later, instruction-tuned generative models enabled drafting and questioning (Ouyang et al., 2022). However, these general models showed clear limitations in legal reasoning and raised concerns about privacy and deployment control. The research focus then moved to training models directly on large-scale legal corpora. Instead of treating legal tasks as simple downstream adaptation, researchers aimed to encode domain knowledge into the model itself. Models such as LawGPT, Lawyer LLaMA, and ChatLaw were developed under this direction (Cui et al., 2023; Huang et al., 2023; Zhou et al., 2024b). SaulLM adopted continual pre-training to maintain general language ability while incorporating legal knowledge (Colombo et al., 2024). Alignment was also treated more seriously in legal settings. In InternLM-Law, reinforcement learning from human feedback was applied not only to improve output quality, but to enforce professional standards and reduce the risk of generating improper legal strategies (Fei et al., 2025).

Multi-Step Reasoning and Tool Use in Legal Workflows. Fluency does not guarantee factual accuracy (Rawte et al., 2023). To reduce hallucination, researchers moved from relying only on parametric memory to using retrieval-augmented generation (Lewis et al., 2020). This approach allows models to access external legal documents during inference. However, standard vector retrieval is often insufficient for precise legal terminology and complex statute relations. To improve retrieval quality, some systems integrate knowledge graphs to capture structured relations between legal concepts. For example, legal knowledge graph enhanced LLM frameworks combine structured statute relations with language modeling to support multi-step reasoning across related provisions (Shi et al., 2024). This allows the model to trace logical connections between statutes instead of relying only on surface similarity. At the same time, structured reasoning methods are introduced to improve

logical consistency in complex cases (Wei et al., 2022). Frameworks such as Disc-LawLLM apply explicit reasoning formats, including syllogistic reasoning and the IRAC (Issue, Rule, Application, and Conclusion) structure, to make the inference process more transparent and grounded in legal rules (Yue et al., 2023). Single law-agents are continuously evolving by integrating instruction decomposition, multimodal visual perception, and advanced reward optimization algorithms like GRPO to reinforce logical reasoning. Despite these advances, single-agent systems remain bounded within legal workflows, as they can support retrieval, reasoning, and drafting but still require human initiation, verification, and approval for consequential legal work.

Why Legal Agency Remains Professionally Bounded. As systems began to tackle more complex legal tasks, researchers began to adopt multi-agent architectures. These systems reflect the division of labor found in real legal practice. Early designs focused on task decomposition. Legal workflows were divided into stages such as intake, retrieval, reasoning, and drafting, and each stage was assigned to a specific agent (Wang and Yuan, 2025). Multi-agent systems also improve reasoning through structured interaction. In the transactional domain, PAKTON ensures traceable grounding by using a collaborative multi-agent workflow for contract review and question answering (Raptopoulos et al., 2025). Similarly, LAW uses an orchestrator to connect multiple text agents and domain tools to collaboratively handle complex contract queries like calculating termination dates (Watson et al., 2025). Beyond task support, some frameworks explore legal simulation. AgentCourt models trial procedures to study how judicial decisions evolve under structured interaction (He et al., 2024b). Law in Silico constructs simulated digital societies populated by legal agents to evaluate how statutes may produce unintended effects before real-world implementation (Wang et al., 2025d). These systems move from assistive reasoning toward structured institutional modeling, enabling the analysis of legal processes at a systemic level.

Taken together, legal AI is moving gradually from Stage 2 toward Stage 3. Current systems can decompose goals, coordinate tools, and complete parts of legal workflows through multi-step reasoning, thereby strengthening the *Efficiency Mechanism* mainly through procedural automation and reduced coordination costs. However, their autonomy remains bounded, and human professionals must remain on the loop to connect abstract legal objectives with concrete tasks. Further progress depends on improving citation reliability, robustness under prompt variation, and clearer standards for legal alignment and delegated agency.

6.5. Agentic AI as a Bounded Executor in Growing-Diffusion Domains

Across trade, media, education, and law, Agentic AI is increasingly functioning as a *bounded executor* embedded in real production workflows. As a new productive instrument, its role in these domains is to absorb procedural cognitive labor across multiple intermediate stages of production, including scenario simulation, partial workflow management, iterative revision, and cross-system orchestration under human supervision. Productivity gains in these domains arise mainly through the *Efficiency Mechanism*, with Agentic AI reducing the time, coordination, and verification costs of complex intermediate tasks. A more limited *Expansion Mechanism* is also emerging, enabling broader workflow coverage and larger-scale coordination or simulation.

What limits further progress is the difficulty of making execution both reliable and accountable under the constraints of each domain. Trade requires compliance that can be checked, media requires factual accuracy and editorial oversight, education requires alignment with teaching goals, and law requires clear reasoning supported by traceable evidence. The key bottleneck is that Agentic AI can participate in workflows, but still cannot fully take on responsibility for outcomes. It can already handle many structured substeps, yet it is still weaker at maintaining consistency over long

tasks, responding to exceptions, and operating safely when errors carry serious consequences. The main safety risks are also fairly clear, including hallucinated content, unstable multi-step reasoning, bias in sensitive contexts, privacy and intellectual-property leakage, and over-reliance by human users. Taken together, Agentic AI in growing-diffusion domains is best understood as a workflow-embedded productive instrument that raises productivity by taking over increasingly broad ranges of intermediate tasks and by deepening execution across more complex stages of production, while leaving important execution and final judgment in human hands.

Table 2 | Representative domain-specific applications in Growing-diffusion industries.

Paper	Task	Technique	Stage
<i>Trade</i>			
Azure Microsoft Foundry Document Intelligence (Microsoft Corporation, 2025)	Information extraction	OCR / IDP	Stage 1
Layoutlmv2 (Xu et al., 2020)	Document understanding	Pretrain & SFT	Stage 1
Mely.ai (Mely.ai, 2025)	Workflow execution	Tool orchestration	Stage 2
AWS supply-chain systems (Pazak and Baeuml, 2025)	Supply-chain tasks	MAS	Stage 2–3
GenAI compliance and risk-support systems (Aslett et al., 2025)	Risk analysis	RAG	Stage 1–2
Vending-Bench (Backlund and Petersson, 2025)	Business simulation	Long-horizon benchmark	Stage 2–3
Project Vend (Anthropic, 2025)	Business automation	Agent-based Simulation	Stage 2–3
<i>Education</i>			
EduChat (Dan et al., 2023)	Intelligent tutoring	LLM chatbot	Stage 1
LearnLM (LearnLM Team, Google, 2024)	Tutoring	Instruction tuning	Stage 1
Pedagogy-R1 (Lee et al., 2025b)	Educational reasoning	Reasoning post-training	Stage 1
SocraticLM (Liu et al., 2024)	Personalized tutoring	Socratic prompting	Stage 2
LLM-based Automated Grading with Human-in-the-Loop (Chu et al., 2025a)	Automated grading	HITL	Stage 2
AutoFeedback (Guo et al., 2024)	Automatic feedback	MAS	Stage 2
Agent4Edu (Gao et al., 2025)	Learner simulation	Generative agents	Stage 2–3
Classroom Simulacra (?)	Behavior simulation	Generative agents	Stage 2–3
<i>Law</i>			
LawGPT (Zhou et al., 2024b)	Legal drafting	Pretrain & SFT	Stage 1
Lawyer LLaMA (Huang et al., 2023)	Legal consultation	Pretrain & SFT & Retrieval	Stage 1
InternLM-Law (Fei et al., 2025)	Legal QA & case analysis	Two-stage SFT	Stage 1
Disc-LawLLM (Yue et al., 2023)	Legal services & reasoning	SFT & Retrieval	Stage 1
SaulLM (Colombo et al., 2024)	Legal text generation	Pretrain & SFT & DPO	Stage 1
ChatLaw (Cui et al., 2023)	Legal consultation	MoE & MAS & KG	Stage 2
L-MARS (Wang and Yuan, 2025)	Legal QA	MAS & Agentic Search	Stage 2
PAKTON (Raptopoulos et al., 2025)	Contract QA & review	MAS & RAG	Stage 2–3
AgentCourt (He et al., 2024b)	Judicial decision-making	MAS & Retrieval	Stage 2–3
LAW (Watson et al., 2025)	Contract query answering	MAS & Tool Orchestration	Stage 3
Law in Silico (Wang et al., 2025d)	Legal society simulation	Agent-based Simulation	Stage 3–4

7. Widespread-Diffusion Industries




Figure 7 | Widespread Industries at Stage 3 and early Stage 4 traits in constrained settings. These sectors are characterized by deep workflow integration, dual-mechanism productivity gains, and governed autonomy. For each industry, the figure identifies representative applications, spanning auto coding, autonomous trading, and clinical documentation, alongside core expansion mechanisms such as workflow orchestration, role specialization, cross-source data fusion, and multi-agent orchestration across complex multistep tasks.

7.1. Information Technology: Executable Feedback and Autonomous Software Workflows

As a domain built around highly structured digital artifacts such as code, repositories, and execution logs, information technology provides especially favorable conditions for Agentic AI. Tasks are modular, feedback is often executable, and intermediate outputs can be checked and revised within relatively standardized environments. For this reason, information technology shows one of the clearest and most continuous trajectories from model-level code generation to workflow-level autonomous execution. This evolution can be understood in three stages: post-training to improve single-round code generation, agentic frameworks to manage complex software workflows, and agentic reinforcement learning to strengthen performance in real-world engineering environments.

Code Generation as the Initial Automation Layer. Early work in this domain focused mainly on post-training code large language models to improve single-round generation quality, especially in instruction following and local program correctness. Supervised fine-tuning approaches such as WizardCoder (Luo et al., 2024), together with reinforcement learning methods that incorporate execution feedback such as CodeRL (Le et al., 2022), PPOCoder (Shojaee et al., 2023), StepCoder (Dou et al., 2024), and DeepSeek-Coder-V2 (DeepSeek-AI et al., 2024) demonstrated substantial gains on standalone code generation tasks. However, this line of work largely treated programming as a single-step prediction problem, making it difficult to scale to advanced coding challenges which require workflow orchestration, multi-perspective testing and iterative debugging.

Repository-Level Agents for Iterative Engineering Work. To overcome the limits of single-round code generation, subsequent work shifted from improving isolated outputs to structuring programming as a multi-step workflow. A first step in this direction came from self-debugging methods, which showed that language models could iteratively detect and correct errors through designed prompts or self-generated tests (Adnan et al., 2025; Chen et al., 2025b, 2024b). Meanwhile, fixed-pipeline frameworks such as CodePlan (Wen et al., 2025) showed that manually orchestrated workflows enable LLMs to tackle complex, repository-level coding tasks. Building on this insight, multi-agent systems such as ChatDev (Qian et al., 2024) introduced explicit role decomposition and standardized operating procedures, recasting software development as a structured pipeline of coordinated sub-tasks. Expanding this multi-agent paradigm beyond software engineering, MapCoder (Islam et al., 2024) assigns specialized roles like example retrieval, planning, code generation, and debugging to different agents to solve complex algorithm problems. Together, these frameworks marked a shift from code generation toward bounded workflow execution, but their effectiveness still depended heavily on hand-designed processes and fixed coordination structures.

Why Software Becomes the Frontier of Agentic Productivity. As agent-based approaches became more complex, the literature increasingly recognized that procedural orchestration alone could not fully compensate for model limitations in long-horizon reasoning. Real software environments involve changing repository states, hidden dependencies, and delayed feedback, which make fixed workflows brittle. Recent research has therefore returned to strengthening agents’ internal reasoning and generalization abilities through reinforcement learning and direct interaction with executable environments. Several works established execution and feedback environments grounded in real-world code repositories, including SWE-Gym (Pan et al., 2025), AgentGen (Hu et al., 2025), and R2E-Gym (Jain et al., 2025), which together form key infrastructure for training and evaluating software engineering agents. Leveraging these environments, DeepSWE (Luo et al., 2025) demonstrated that large-scale reinforcement learning can significantly improve autonomous exploration and bug-fixing performance under complex dependency structures. RepoForge (Chen et al., 2025c) further advanced this direction by optimizing the end-to-end execution and data pipeline, enabling higher-throughput training at scale. This stage marked a shift from manually designed workflow scaffolds toward agents that can adapt more effectively within real software engineering loops.

Taken together, existing research in information technology shows that AI coding systems have evolved from local developer assistance toward increasingly autonomous execution in software engineering workflows. Agents can already handle goals that are clearly bounded in scope yet require complex, multi-step reasoning, such as resolving real-world GitHub issues through iterative planning, file editing, environment interaction, and testing. In this sense, information technology provides one of the clearest domains in which Agentic AI has moved beyond passive support and toward sustained workflow execution. Based on these capabilities, current systems are best understood as operating at Stage 3 and showing early Stage 4 traits in constrained settings. At the same time, their broader deployment remains constrained by long-horizon reliability and the risks of failure in real development environments.

7.2. Finance: Productivity Amplification under Risk Amplification

As a high-risk domain characterized by data intensity, complex decision logic, and strict regulation, the financial industry requires agents with strong domain specialization, reliable numerical reasoning, and robust compliance awareness. Early systems mainly served as auxiliary tools for text processing and information retrieval, but recent advances have empowered financial agents to participate in core processes such as investment research, trading, and risk management. This pattern indicates that finance can amplify productivity through agentic execution, but it can also amplify

risk when autonomous decisions interact with markets, compliance systems, and customer assets.

Financial Reasoning and Research Automation. To address the inherent gaps of general LLMs in financial scenarios, a critical step is to adapt them into specialized financial foundation models with enhanced reasoning and strict compliance. For example, researchers have explored domain-specific pre-training on financial corpora (Lee et al., 2024), hybrid paradigms that combine LLMs with traditional predictive models (Sanz-Guerrero and Arroyo, 2024), and reinforcement learning-based post-training to improve reasoning stability (Liu et al., 2025). Compliance capabilities are also integrated into model inference through supervised fine-tuning alongside external post-hoc checks (Geng, 2026; Zhu et al., 2025). However, these efforts mainly improve reasoning or compliance in isolation, rather than holistic trade-offs among accuracy, traceability, and efficiency in real-world deployment, leaving a clear gap between research progress and practical demands (Dong et al., 2025; Evite et al., 2026). To bridge this gap and fully leverage the improved foundational capabilities, financial AI is evolving toward an agentic stage of proactive task execution, where agents autonomously plan workflows, invoke specialized tools, and iteratively refine strategies around business objectives.

Agentic Execution across Trading, Banking, and Insurance. In practice, current systems increasingly adopt both single-agent specialization and multi-agent collaboration across financial sub-domains.

Securities and Asset Management: In investment research, multi-agent collaboration has become a dominant paradigm. Systems such as FinRobot decompose research workflows into sub-tasks handled by specialized agents, while debate- or verification-based frameworks improve report reliability (Cai et al., 2025; Yang et al., 2024). Some approaches further integrate code agents and vision models to automate data acquisition and visualization pipelines (Jin et al., 2025b). Beyond research generation, agents are increasingly applied to trade execution. For example, FinAgent autonomously invokes technical analysis tools and iteratively refines strategies through reflection and memory mechanisms (Zhang et al., 2024b). Multi-agent trading frameworks further simulate institutional decision structures to improve the robustness of strategies, while some systems introduce internal competitive and ranking mechanisms among agents to filter noisy signals in dynamic markets (Xiao et al., 2024; Zhao et al., 2025). Additionally, reinforcement-learning-based systems such as QuantAgent enable adaptive strategy optimization to capture excess returns in simulated markets (Xiong et al., 2025a).

Banking Sector: In banking, intelligent agents are increasingly being integrated into core operational processes. In risk management, LLM-based agents combine structured transaction data with unstructured information such as market signals and public sentiment to support adaptive risk assessment and real-time decision-making (Cao et al., 2024; Suryavanshi and Kumar, 2025). Meanwhile, conversational execution agents are emerging as intelligent interfaces for banking systems. Frameworks such as the Ryt AI Banking Agent enable users to perform account operations, payments, and service inquiries through natural language by invoking internal banking APIs (McKinsey & Company, 2025). Together, these systems illustrate the transition of banking agents from isolated task automation toward integrated workflow execution.

Insurance Sector: In insurance services, intelligent agents are increasingly used for automated claims processing and fraud detection. Multimodal systems combining computer vision and LLM reasoning can jointly analyze policy documents and damage images to support automated claim assessment (Miao et al., 2026). This multimodal fusion also improves fraud detection by integrating visual, textual, and structured claims data (Asgarian et al., 2023). When fraud detection requires data from multiple institutions, but centralized sharing is restricted, federated learning provides a

privacy-preserving framework for collaborative modeling across institutions (Kairouz et al., 2021).

Why Financial Autonomy Must Remain Governable. Despite rapid progress in agentic financial systems, real-world deployment remains fundamentally constrained by the intrinsic properties of financial markets, including high uncertainty, strong interdependence, and strict regulatory requirements. As a result, current research is increasingly shifting from purely improving task performance to building robust and governable agent architectures, including the integration of risk-aware reasoning, explicit constraint modeling, and compliance-aligned decision processes directly into agent workflows. In parallel, there is a growing emphasis on designing feedback-rich environments and evaluation protocols that reflect real-world financial dynamics, enabling agents to learn and be assessed under more realistic market conditions.

7.3. Medicine: Auditable Human-AI Collaboration in High-Stakes Care

Medical Foundations and Clinical Reasoning Alignment. Early medical AI research focused on integrating the linguistic reasoning abilities of general large language models with specialized medical knowledge. While general models exhibit strong language capabilities, they lack sufficient clinical knowledge coverage and rigor in medical reasoning. To address this gap, early efforts relied on large-scale medical knowledge injection, transforming general models into domain-specific medical foundation models. For instance, PMC-LLaMA (Wu et al., 2024) integrates biomedical literature and textbooks with medical instruction tuning to enable domain-specific analysis and question answering. Beyond text, generative models have also been extended to biomedical tasks. The PanGu Drug Model (Lin et al., 2022), trained on large-scale molecular data, supports molecular generation and property prediction, demonstrating that generative capabilities can transfer to drug discovery workflows. As the knowledge base matured, research shifted toward capability alignment. Supervised fine-tuning (SFT) transformed static knowledge into clinical reasoning abilities, as seen in Huatuo-26M (Wang et al., 2025c) and BenTsao (Wang et al., 2025a). Further work improved interpretability of reasoning, such as ClinicalGPT-R1 (Lan et al., 2025). Meanwhile, reinforcement learning and verifiable alignment were introduced to enforce medical safety and reduce bias. For example, HuatuoGPT-o1 (Chen et al., 2024a) designs reward mechanisms for reasoning quality, while approaches inspired by Constitutional AI (Bai et al., 2022) incorporate ethical constraints via RLAIIF.

Agentic Workflows for Documentation, Diagnosis, and Discovery. Medical AI is now transitioning from question-answering systems to executable agentic systems. In clinical settings, single-agent systems integrate deeply with workflows. Systems such as Google MedLM (Yossi Matias, Aashima Gupta, 2023) and Microsoft Dragon Copilot (Kees Hertogh, 2025) connect to EHRs, generate structured clinical notes, retrieve evidence, and assist coding. Similarly, transcription systems like Nuance Copilot (Microsoft, 2025) and AWS HealthScribe (Amazon Web Services, 2023) automate clinical documentation, significantly reducing physician workload. Multi-agent systems further extend capabilities by simulating collaborative medical teams. Agent Hospital (Li et al., 2024) enables iterative learning through simulated cases, while MATEC (Cho et al., 2025) coordinates specialized agents in critical care scenarios. MedAgent-Pro (Wang et al., 2025e) introduces evidence-based workflows combining planning, tool use, and decision aggregation, reflecting a shift toward controllable and auditable execution. Beyond hospitals, agentization extends to pharmaceutical and public domains. Prompt-to-Pill (Vichentijevikj et al., 2025) integrates drug discovery pipelines into end-to-end multi-agent systems, while virtual laboratory platforms (Zhang et al., 2026a) enable large-scale parallel hypothesis testing. On the public side, systems such as ChatGPT Health (OpenAI, 2026) provide continuous health consultation, supporting users in understanding symptoms and navigating healthcare processes. These developments show that AI agents can execute tasks beyond human temporal limits,

including large-scale data analysis, real-time transcription, and continuous consultation, becoming increasingly important executors in healthcare systems.

Why Healthcare Productivity Depends on Governed Delegation. The development of medical AI has evolved from knowledge injection and capability alignment to agent-centered task execution. Rather than relying solely on scaling, progress increasingly depends on controllability, compliance, and workflow integration. AI improves healthcare through both efficiency mechanisms—automating workflows and assisting decision-making—and expansion mechanisms, such as large-scale molecular screening (Vichentijevikj et al., 2025) and continuous public health consultation (OpenAI, 2026). Accordingly, evaluation criteria have shifted from response quality to controlled execution and governance. Real-world systems enforce safety through tool constraints, evidence binding, permission control, and human oversight. Most current systems operate at Stage 3 (Adaptive Coordinator), while only early Stage 4 traits appear in constrained settings because safety and regulatory requirements remain binding. Overall, the core value of medical AI lies in transforming complex healthcare processes into scalable and auditable human–AI collaborative systems.

7.4. Agentic AI as Scaled Orchestration in Widespread-Diffusion Domains

Across information technology, finance, and medicine, Agentic AI is an increasingly central *productive instrument* embedded in high-frequency, high-value workflows. In these domains, production already depends heavily on structured digital materials such as code repositories, financial data, transaction records, clinical notes, medical images, and standardized evidence. This makes it easier for agents to take part in continuous workflow orchestration. As a result, Agentic AI is increasingly used as an authorized executor of tasks. It can make detailed plans, retrieve supportive information, and revise intermediate results within rules. Its contribution to productivity works in both width and depth. In width, it now covers a broad set of end-to-end processes, including repository-level code maintenance and debugging, market analysis and claims processing, and clinical documentation, evidence synthesis, and drug discovery support. In depth, delegation has moved to sustained multi-step execution. This makes these domains the clearest cases in which Agentic AI is becoming deeply integrated into production.

This productivity gain is driven by both mechanisms discussed earlier. First, the *Efficiency Mechanism* remains the foundation. Agents reduce the time and verification costs of complex cognitive work by shortening many intermediate steps. Second, the *Expansion Mechanism* is much more visible here than in earlier stages. Agents make it possible to perform tasks at a scale, speed, and level of continuity that human teams alone would struggle to sustain. Examples include parallel code debugging, continuous market monitoring, large-scale evidence synthesis, and always-on public or enterprise services. However, this does not mean that these domains have fully moved into unrestricted autonomy. The main bottleneck is no longer whether agents can produce useful outputs. It is whether they can act reliably under uncertainty while meeting demands for accountability.

This challenge appears in different forms across domains. In Information Technology Industry, the main problem is long-horizon reliability in changing environments with hidden dependencies and real deployment risks. In finance, the key concern is that autonomous action may turn model error into market, compliance, or even systemic risk. In medicine, the barrier is even higher because systems must meet strict requirements for safety, explainability, privacy, and legal responsibility when outputs affect patient care. As deployment becomes more widespread, safety problems also become more concrete. Hallucinated reasoning, prompt injection, unauthorized tool use and privacy leakage are no longer just model-quality issues. They become direct production risks. For this reason, the key task at this stage is not simply to push autonomy further, but to ensure that deeper delegation

remains safe, controllable, and accountable.

Table 3 | Representative domain-specific applications in Widespread-diffusion industries.

Paper	Task	Technique	Stage
<i>Information Technology</i>			
WizardCoder (Luo et al., 2024)	Single-round Coding	SFT	Stage 2
CodeRL (Le et al., 2022)	Single-round Coding	RL	Stage 2
PPOCoder (Shojaee et al., 2023)	Single-round Coding	RL	Stage 2
StepCoder (Dou et al., 2024)	Single-round Coding	RL	Stage 2
DeepSeek-Coder-V2 (DeepSeek-AI et al., 2024)	Single-round Coding	RL	Stage 2
Self-Debug (Chen et al., 2024b)	Self-debugging	Agentic Framework	Stage 2–3
Revisit Self-Debugging (Chen et al., 2025b)	Self-debugging	Agentic Framework	Stage 2–3
Large Language Model Guided Self-Debugging Code Generation (Adnan et al., 2025)	Self-debugging	Agentic Framework	Stage 2–3
ChatDev (Qian et al., 2024)	Repository-level Coding	Agentic Framework	Stage 3
MapCoder (Islam et al., 2024)	Single-round Coding	Agentic Framework	Stage 3
CodePlan (Wen et al., 2025)	Repository-level Coding	Agentic Framework	Stage 3
SWE-Gym (Pan et al., 2025)	Repository-level Coding	Agentic RL	Stage 3–4
AgentGen (Hu et al., 2025)	Repository-level Coding	Agentic RL	Stage 3–4
R2E-Gym (Jain et al., 2025)	Repository-level Coding	Agentic RL	Stage 3–4
DeepSWE (Luo et al., 2025)	Repository-level Coding	Agentic RL	Stage 3–4
RepoForge (Chen et al., 2025c)	Repository-level Coding	Agentic RL	Stage 3–4
<i>Finance</i>			
Alpha-GPT (Wang et al., 2023)	Alpha mining	Human-AI interaction	Stage 2
Fin-R1 (Liu et al., 2025)	Financial reasoning	RL	Stage 2
DianJin-R1 (Zhu et al., 2025)	Financial reasoning	RL	Stage 2
FinRobot (Yang et al., 2024)	Financial analysis	MAS	Stage 3
FinAgent (Zhang et al., 2024b)	Trading	Tool augmentation	Stage 3–4
TradingAgents (Xiao et al., 2024)	Trading	MAS	Stage 3
FinDebate (Cai et al., 2025)	Financial analysis	MAS	Stage 3
FinSight (Jin et al., 2025b)	Financial research	Agentic search	Stage 3–4
RiskLabs (Cao et al., 2024)	Risk prediction	Multimodal LLM	Stage 2
Credit Risk Meets Large Language Models (Sanz-Guerrero and Arroyo, 2024)	Credit scoring	LLM scoring	Stage 2
MMLM-CA (Miao et al., 2026)	Claims adjudication	Multimodal LLM	Stage 2
<i>Medical</i>			
PMC-LLaMA (Wu et al., 2024)	Medical QA	SFT	Stage 1
PanGu Drug (Lin et al., 2022)	Drug discovery	SFT	Stage 1
Huatuo-26M (Wang et al., 2025c)	Clinical reasoning	SFT	Stage 2
BenTsao (Wang et al., 2025a)	Clinical reasoning	SFT	Stage 2
ClinicalGPT-R1 (Lan et al., 2025)	Medical reasoning	RL	Stage 2
HuatuoGPT-o1 (Chen et al., 2024a)	Medical reasoning	RL	Stage 2
MedLM (Yossi Matias, Aashima Gupta, 2023)	Clinical assistance	Agentic RL	Stage 3
Microsoft Dragon Copilot (Kees Hertogh, 2025)	Clinical assistance	Agentic RL	Stage 3

Table 3 continued

Paper	Task	Technique	Stage
AWS HealthScribe (Amazon Web Services, 2023)	Documentation	Agentic RL	Stage 3
Agent Hospital (Li et al., 2024)	Diagnosis	MAS	Stage 3
MATEC (Cho et al., 2025)	Critical care	MAS	Stage 3
MedAgent-Pro (Wang et al., 2025e)	Decision support	MAS	Stage 3
Prompt-to-Pill (Vichentijevikj et al., 2025)	Drug discovery	MAS	Stage 3
Virtual Biotech (Zhang et al., 2026a)	Experimentation	MAS	Stage 3

8. Cross-Industry Comparison: What Determines Productivity Diffusion?

The cross-industry comparison shows that the uneven diffusion of agentic AI is not simply a matter of model capability. The same technical system can generate substantial productivity gains in one sector while remaining confined to bounded assistance in another. This pattern indicates that productivity diffusion depends on whether a sector can convert agentic capability into reliable workflow execution and accountable delegation.

Digitization of production objects. Agentic AI diffuses more deeply when the objects of production are already digital. Code, documents, structured data, transaction records, and clinical notes can be read, modified, checked, and revised inside software environments. This helps explain why information technology, finance, and parts of medicine show deeper integration than agriculture, manufacturing, and public administration, where physical conditions or institutional procedures limit direct execution.

Feedback verifiability. Productivity gains are more visible when outputs can be verified through executable tests, structured evidence, expert rubrics, or operational metrics. Software has unusually strong feedback because code can be run and tested. By contrast, education, media, government, and healthcare often require qualitative judgment, social context, or professional interpretation, which makes closed-loop autonomy harder to justify.

Error tolerance. The cost of mistakes strongly shapes the depth of delegation. In low-tolerance domains such as finance, healthcare, manufacturing, law, and government, errors can create legal, physical, financial, or public-trust consequences. In these settings, agentic AI may still improve efficiency, but autonomy remains bounded by review, permission control, audit trails, and fallback procedures.

Workflow modularity. Agent orchestration works best when production can be decomposed into explicit subtasks with stable interfaces. Trade documentation, legal drafting, software maintenance, and clinical documentation all contain modular workflow components that agents can execute or coordinate. Where work depends on tacit knowledge, ambiguous goals, or constantly changing environmental conditions, agents remain closer to decision support than delegated execution.

Accountability structure. The hardest barrier is often not whether AI can act, but who remains responsible when it acts. Sectors with clear supervisory structures, auditable outputs, and permission boundaries can delegate more operational responsibility. Sectors where authority, liability, or professional judgment cannot be cleanly transferred keep humans at the center of final decisions. The evidence therefore suggests that productivity diffusion is ultimately governed by the interaction between technical capability and institutional accountability.

9. Discussion and Challenge

9.1. Other Industries

To complement the previous analysis, the following discussion assesses the progress and challenges of AI applications in recent years across several additional sectors (e.g., power and gas, transportation and warehousing, accommodation and catering, and culture and entertainment).

9.1.1. Power and Gas.

The power sector is transitioning toward LLM- and agent-driven automated operations, primarily categorized into automated system analysis and intelligent operational control. In system analysis, frameworks like X-GridAgent and GridMind (Jin et al., 2025a; Wen and Chen, 2025) automate scientific computing (e.g., optimal power flow) by coupling high-level reasoning with deterministic engineering tools. Expanding on this, recent studies introduce graph and tabular embeddings to help LLMs directly process complex grid constraints for power flow optimization, as demonstrated by PowerGraph-LLM (Bernier et al., 2025). For operational control, some approaches like Grid-Agent (Zhang et al., 2025b) combine semantic planning with domain solvers to coordinate real-time remediation tasks like topology reconfiguration. Specialized models such as GAIA (Cheng et al., 2024) are developed specifically to handle advanced power dispatch operations. Similarly, LLM4DistReconfig (Christou et al., 2025) utilizes fine-tuned LLMs to rapidly predict optimal network configurations during dynamic topology changes. In summary, while these AI applications reduce operational complexity, challenges in safety, controllability, and real-world deployment remain.

9.1.2. Traffic and Mobility.

The transportation sector is transitioning from static, rule-based management to dynamic, LLM-driven autonomous coordination. This evolution addresses operational challenges progressively, from macro-level traffic simulation to micro-level control and routing execution. Initially, evaluating city-scale transportation policies was hindered by the complex manual setup of traffic simulators. To address this, some studies utilize LLM-guided agentic frameworks to automate traffic scenario generation and simulation through natural language, exemplified by AgentSUMO (Jeong et al., 2025) and TrafficSimAgent (Du et al., 2025). To overcome the traditional reliance on rigid heuristics for intersection traffic control, certain approaches employ LLM-based agents equipped with common sense and reasoning capabilities to dynamically optimize signal timings, as demonstrated by LLMLight (Lai et al., 2025). Regarding physical fleet execution, solving dynamic Vehicle Routing Problems traditionally demanded extensive human expertise to design routing algorithms. Recent frameworks like VRPAgent (Hottung et al., 2025) and AFL (Zhang et al., 2026b) leverage LLMs to autonomously discover and refine heuristic operators for complex fleet coordination. Overall, while these progressive applications highlight the immense potential of AI agents across transportation systems, challenges in real-time latency and multi-agent safety remain.

9.1.3. Accommodation and Catering.

The accommodation and catering sectors are transitioning from labor-intensive service models to AI-integrated autonomous operations. To overcome early constraints imposed by manual staff availability and static pricing rules, some studies employ LLM-driven conversational agents to automate complex front-desk booking and inquiry workflows, as demonstrated by recent systems (Aransyah et al., 2026). For backend optimization, certain approaches utilize predictive AI to dynamically forecast demand and adjust pricing, replacing traditional empirical models (Henriques and Pereira, 2024; Talón-Ballesteros et al., 2022). Moving to the catering domain, to overcome efficiency and consistency bottlenecks in manual multi-dish preparation, researchers develop customized robotic systems to automate physical cooking tasks, such as the dual-arm robot YORI (Noh et al., 2025). To further enhance kitchen autonomy, some studies have leveraged large language models to translate natural-language recipes into structured task trees and executable robotic cooking plans for embodied kitchen systems (Kanazawa et al., 2024; Sakib and Sun, 2024).

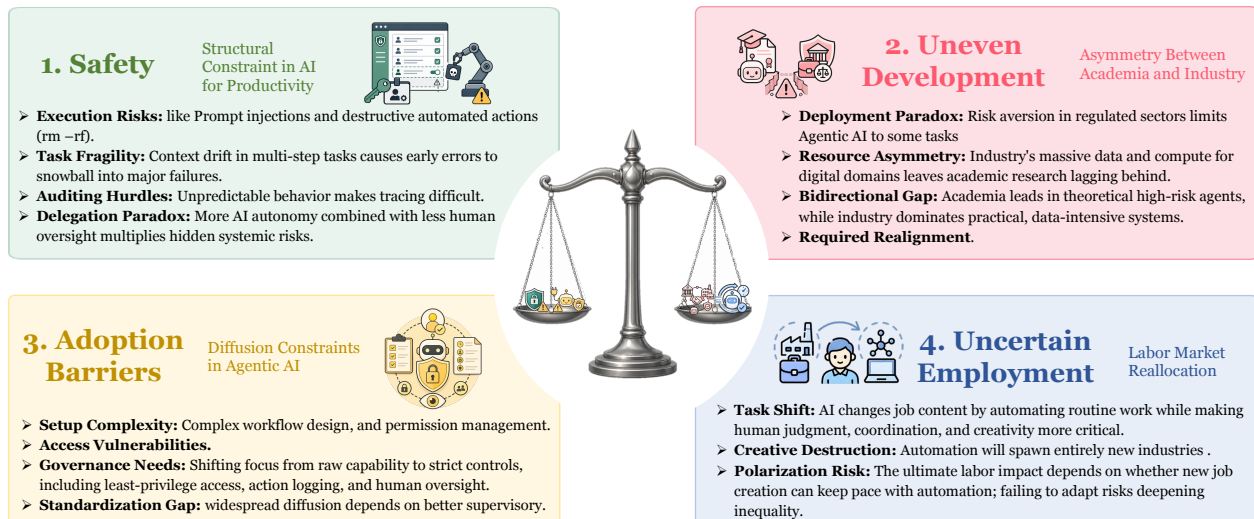


Figure 8 | Structural Constraints on Productivity Diffusion. This figure outlines four constraints on the broader diffusion of agentic AI: (1) safety and long-horizon reliability risks from execution errors, task fragility, and reduced oversight; (2) the capability-deployment gap between academia and industry; (3) adoption friction from setup complexity, governance needs, and standardization gaps; and (4) labor reallocation and employment uncertainty from task shifts and labor market polarization. Together, they indicate that the future of agentic AI depends on balanced progress in safety, governance, and socioeconomic adaptation.

9.1.4. Culture and Entertainment.

The cultural and entertainment sectors are transitioning from manual, resource-intensive production to generative AI and agent-driven interactive ecosystems. Traditional manual asset creation carries high temporal and financial costs. In response, multi-agent frameworks such as FilmAgent (Xu et al., 2025b) and MAViS (Wang et al., 2026), together with modern generative pipelines (Tsiavos and Kitsios, 2025), have emerged to automate scriptwriting, cinematography, and long-sequence video storytelling workflows. Within gaming, to overcome traditional limitations, recent workflows integrate generative AI to accelerate character creation pipelines (Wu et al., 2025c), while LLM-based NPC systems are enabling more interactive and active in-game characters (Christiansen et al., 2024; Kim et al., 2025). Similarly, AI-enhanced conversational systems are now deployed in public venues like museums to replace static, one-way exhibit interpretations with dynamically enriched visitor experiences (Wen and Ma, 2024). However, alongside these productivity gains, the expansion of generative AI continues to intensify critical debates over copyright, authorship, and creative integrity (Al-Busaidi et al., 2024; Amankwah-Amoah et al., 2024).

9.2. Remaining Challenges on Productivity Diffusion

These challenges suggest that the next stage of AI for Productivity will be determined less by model capability alone than by workflow integration, verification infrastructure, permission design, and institutional accountability.

9.2.1. Safety and Long-Horizon Reliability

Safety is an essential constraint in AI for productivity. As generative AI and high-permission agentic systems are increasingly integrated into enterprise production workflows, system safety concerns are

evolving from traditional issues of software vulnerabilities and access control into a more complex risk structure involving model behavior, data flows, and decision execution. Unlike earlier information systems that primarily functioned as passive tools, AI systems in production environments actively participate in reasoning, decision-making, and even cross-system operations. As a result, risks extend beyond information leakage to higher-level issues such as behavioral deviation and loss of execution control.

In practice, these risks have already manifested in diverse forms. First, at the data level, existing studies and industry reports indicate that a large number of enterprise employees routinely input customer data, financial records, and even proprietary code into external generative models, thereby bypassing established data governance and access control mechanisms and significantly increasing the risk of sensitive information leakage (Cyberhaven Labs, 2025; Menlo Security, 2025). Second, at the model level, prompt injection has been widely recognized as one of the primary attack vectors in generative AI systems. Through carefully crafted inputs, attackers can induce models to deviate from their intended objectives, potentially exposing internal knowledge or performing unauthorized actions (A10 Networks, 2025; Practical DevSecOps, 2026). Third, at the system level, as agents become deeply integrated with toolchains, high-permission agents may trigger destructive behaviors during automated execution due to policy misalignment or contextual misunderstanding, such as deleting critical data or executing operations that violate compliance constraints. These issues are particularly pronounced in low-tolerance domains such as finance, healthcare, and public administration, where even minor errors can lead to systemic consequences or legal liabilities.

Long-Horizon Reliability as a Structural Source of Risk. A key source of safety risk lies in the structural mismatch between long-horizon task requirements and the capability profile of current models. Real-world production tasks typically take the form of multi-stage, tool-integrated continuous control processes, where correctness is defined at the trajectory level rather than by the accuracy of individual steps (Abhyankar et al., 2025; Wang et al., 2025b). However, current LLMs, trained under next-token prediction objectives, lack native mechanisms for modeling long-term state and constraint consistency, making them prone to intention drift, constraint violations, and breakdowns in intermediate reasoning during execution. This structural fragility further manifests as error compounding in long sequences: small early deviations can propagate and amplify by altering subsequent decision contexts, ultimately leading to unpredictable system behavior. Meanwhile, context drift weakens the persistence of critical constraints, causing models to gradually deviate from initial objectives over extended interactions. In high-permission settings, these issues directly translate into safety risks, including unauthorized actions, incorrect decision execution, and unintended data exposure. Therefore, insufficient long-horizon reliability is not only a performance bottleneck but also a foundational constraint on system safety.

Sustainability Constraints and Persistent Security Vulnerabilities. Another class of safety challenges arises from sustainability constraints in the long-term operation of AI productivity systems. First, the non-deterministic behavior of models and ongoing version evolution reduce reproducibility, complicating security auditing, fault tracing, and compliance verification (Atil et al., 2024; Ma et al., 2024). Second, sustained capability improvement relies heavily on external data and environment interaction, while high-quality data—particularly realistic long-horizon task trajectories—remains scarce in practice, thereby limiting robust improvements in complex decision-making scenarios. This constraint on data availability and capability growth makes systems more likely to be deployed or scaled without sufficient validation, leading to the accumulation of latent risks (Dulac-Arnold et al., 2021). In addition, continuous deployment environments introduce new attack surfaces, such as prompt injection and toolchain-based attacks, transforming safety from a one-time design concern

into an ongoing operational challenge (Reddy and Gujral, 2025). At the same time, as AI systems take on more responsibilities, human involvement in supervision and review may diminish, weakening the ability to detect anomalous behavior. This dynamic of “increased delegation and reduced oversight” further amplifies systemic safety risks (Huang et al., 2025).

9.2.2. *The Capability-Deployment Gap between Academia and Industry*

Another challenge of the current AI landscape is a bidirectional asymmetry between academic research and industrial reality. While academia often leads in exploring high-risk agents, industry leads in building data-intensive systems.

In high-stakes, regulated sectors such as finance and law, academic research often explores Stage 3 autonomous agents capable of complex portfolio management or legal reasoning. However, real-world deployment lags significantly behind these theoretical capabilities. Industry reports indicate that adoption is restricted to low-risk, auxiliary functions such as customer support and document drafting, rather than core decision-making processes (Elad, 2025). This cautious approach creates a paradox of limited economic returns. Since organizations limit AI to low-impact tasks to reduce legal and regulatory risk, potential productivity gains are also restricted, which further discourages the deeper integration needed to fully realize the value of autonomous AI systems (AI, 2025).

In contrast, this dynamic is reversed in digital-native domains such as software engineering. Leveraging large proprietary datasets, expanded GPU resources, and user feedback loops that remain inaccessible to open-source researchers, proprietary coding agents (e.g. Claude Code (Anthropic, 2025a), Codex (OpenAI, 2025), Cursor (Anysphere, 2023), OpenClaw (openclaw, 2026), WisPaper (Infinite, 2025)) consistently outperform academic and open-source systems across a range of benchmarks requiring repository-level reasoning and end-to-end task completion (Jimenez et al., 2024). In this context, the uneven distribution of resources results in slower progress in academic research compared to industrial development.

Consequently, a substantial gap exists between the theoretical potential of AI agents and their practical implementation across industries. Addressing this gap is critical, since uneven progress may leave high-risk industries without effective automation capabilities, while also leading to the concentration of core functionalities within closed, proprietary systems. To align these developments, future research efforts necessitate enhanced coordination between academia and industry. Academic research should shift its focus from sole emphasis on performance metrics to prioritizing reliability and auditability for regulated sectors. Additionally, the open research community should develop data-efficient methodologies to maintain parity with industrial-scale systems.

9.2.3. *Adoption Friction in Real Organizations*

A third challenge is that agentic AI still faces high barriers to broad adoption. Recent systems show that AI can move beyond text-based assistance toward tool use and multi-step task execution, but this has not yet translated into wide deployment across ordinary users and firms (Singla et al., 2025; Zhou et al., 2024a). The gap reflects not only model limits, but also the practical difficulty of deploying, supervising, and safely governing agents in real settings (Hughes, 2026).

These barriers come from several sources. Many agent systems still require complex setup, including tool connection, workflow design, and permission management, which keeps effective use concentrated among technically capable users and organizations (Hughes, 2026). At the same time, stronger execution ability often requires broader access to files, browsers, APIs, and enterprise systems, which makes mistakes more costly and expands the attack surface. This is especially important

because agent systems blur the line between data and instructions and remain vulnerable to prompt-injection style attacks (Greshake et al., 2023; Lazer et al., 2026).

Current responses therefore focus less on raw capability alone and more on control, identity, and governance. Recent standards and governance work places growing weight on authentication, authorization, least-privilege access, action logging, and human approval for high-risk actions (Booth et al., 2026; Infocomm Media Development Authority, 2026). This suggests that broader diffusion will depend not only on making agents stronger, but also on making them easier to monitor, constrain, and integrate into existing institutions (Staufer et al., 2026).

This pattern is similar to the early development of personal computers. Early PC systems also had high barriers and diffused slowly at first. Wider adoption came later through better interfaces, standardized platforms, and practical applications that made the technology useful to ordinary users and firms (Bresnahan and Greenstein, 1999; Jovanovic and Rousseau, 2005). Agentic AI may now be in a similar middle stage between early experimentation and wider diffusion. Looking ahead, progress will likely depend on narrow high-value use cases, stronger supervisory interfaces, and clearer standards for permissions, tool use, and auditing (Booth et al., 2026; Helpman and Trajtenberg, 1994).

9.2.4. Labor Reallocation and Employment Uncertainty

A fourth challenge is that the employment effects of agentic AI remain highly uncertain. Agents may automate some tasks and raise productivity, but they may also change job structure, reshape labor demand, and create adjustment pressures before new employment opportunities fully emerge.

One way to understand this is through the task-based framework (Acemoglu and Autor, 2011). This view treats production as a set of tasks performed by labor, capital, or algorithms, rather than as a fixed set of occupations. Agentic AI therefore matters because it changes which tasks can be automated, which still rely on human labor, and how companies reorganize work internally (Acemoglu, 2024; Acemoglu and Restrepo, 2022; Autor et al., 2003). As a result, the labor-market effects of Agentic AI are likely to appear first through changes in job content. Routine and standardized tasks may weaken, while judgment, coordination, and creative tasks become relatively more important.

A broader perspective comes from endogenous growth theory, which emphasizes innovation, knowledge accumulation, and the creation of new products and activities as sources of long-run growth (Aghion and Howitt, 1992; Babina et al., 2024; Romer, 1989). In this view, AI affects employment through a process of creative destruction. Some jobs may contract, but new industries and new forms of work may also emerge. The overall outcome therefore depends on whether new task creation, organizational adaptation, and worker reallocation can keep pace with automation. If they do not, AI may deepen polarization and inequality. If they do, AI may raise productivity while also supporting new employment opportunities.

10. Conclusion and Further Outlook

The central argument of this paper is that Agentic AI should be understood as a new productive instrument that is beginning to reorganize how work is executed, coordinated, and scaled. Across industries, its economic value comes from two closely connected mechanisms. First, it reduces the time and labor required for existing tasks and therefore improves productive **efficiency**. Second, it makes a broader range of economically valuable tasks feasible at acceptable cost and thereby **expands** the scope of production. Yet the diffusion of this new productive instrument remains highly uneven across sectors. In some industries, AI still plays a bounded assistive role. In others, it has already been embedded in multi-step workflows as a supervised executor. In a small number of digital-native

domains, it is moving toward more sustained autonomous agency. Taken together, the importance of Agentic AI lies not only in improving efficiency at the task level, but also in its growing capacity to reshape the mode of production itself.

This also means that the impact of Agentic AI increasingly goes beyond benchmarks and cannot be captured by numbers alone. As agents take over more execution, coordination, monitoring, and routine decision-making, production relations themselves are beginning to change. The role of human labor is gradually shifting away from directly completing large numbers of intermediate tasks and toward goal-setting, handling unexpected exceptions, and verifying final outputs. At the organizational level, firms may increasingly redesign workflows around hybrid human and agent systems, reallocate control across layers of execution, and create new divisions of labor among workers, managers, and machine agents. For this reason, the key long-run question is not simply how much output AI can generate, but how it changes who performs which tasks, who controls the workflow, and who remains accountable for the outcome. The future of AI for Productivity therefore depends not only on further gains in capability, but also on progress in reliability, safety, evaluation, and institutional governance.

Looking ahead, **AI for Productivity** can be seen as the first major effect of Agentic AI in a broader developmental trajectory. After directly transforming production, the next step is likely to be **AI for Science**. At this stage, Agentic AI will begin to accelerate scientific discovery itself through literature review, hypothesis generation, simulation, experimentation, and cross-disciplinary knowledge integration. Beyond that lies **AI for AI**. As human capability approaches its limits in some parts of AI development, Agentic AI will increasingly participate in improving AI itself, including data, architectures, harnesses, and evaluation systems. At the macro level, AI for Productivity, AI for Science, and AI for AI may form a mutually reinforcing loop. AI reorganizes production, stronger productive capacity supports faster scientific discovery, and advances in science and engineering feed back into more capable AI systems, which then further reshape production. The deepest impact of Agentic AI may therefore lie not in automating isolated tasks, but in driving the emergence of a new socio-technical production system in which economic activity, knowledge creation, and AI development increasingly coevolve. The central insight is that AI for Productivity is not merely a story of faster task completion. It is a story of delegated agency: as AI systems take over execution, coordination, and monitoring, productivity gains increasingly depend on how organizations redesign workflows, assign responsibility, and govern autonomous action.

References

- A10 Networks. Generative ai security risks, 2025. URL <https://www.a10networks.com/blog/generative-ai-security-risks/>. Accessed: 2026-04-12.
- ABB. Helping industries do better with generative ai: Abb launches genix copilot with microsoft. Press Release, December 2024. URL <https://new.abb.com/news/detail/122021>. ABB Ability™ Genix Copilot integrates generative AI to boost productivity, efficiency and sustainability in industrial operations in collaboration with Microsoft.
- Reyna Abhyankar, Qi Qi, and Yiyang Zhang. Osworld-human: Benchmarking the efficiency of computer-use agents. *arXiv (Cornell University)*, abs/2506.16042, 2025. doi: 10.48550/arxiv.2506.16042.
- Daron Acemoglu. The simple macroeconomics of ai*. *Economic Policy*, 40(121):13–58, 08 2024. ISSN 0266-4658. doi: 10.1093/epolic/eiae042. URL <https://doi.org/10.1093/epolic/eiae042>.

- Daron Acemoglu and David Autor. Chapter 12 - skills, tasks and technologies: Implications for employment and earnings. volume 4 of *Handbook of Labor Economics*, pages 1043–1171. Elsevier, 2011. doi: [https://doi.org/10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5). URL <https://www.sciencedirect.com/science/article/pii/S0169721811024105>.
- Daron Acemoglu and Pascual Restrepo. Artificial intelligence, automation and work. NBER Working Papers 24196, National Bureau of Economic Research, Inc, Jan 2018. URL <https://ideas.repec.org/p/nbr/nberwo/24196.html>.
- Daron Acemoglu and Pascual Restrepo. Tasks, automation, and the rise in u.s. wage inequality. *Econometrica*, 90(5):1973–2016, 2022. doi: <https://doi.org/10.3982/ECTA19815>. URL <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA19815>.
- Muntasir Adnan, Zhiwei Xu, and Carlos C. N. Kuhn. Large language model guided self-debugging code generation. *CoRR*, abs/2502.02928, 2025. doi: 10.48550/ARXIV.2502.02928. URL <https://doi.org/10.48550/arXiv.2502.02928>.
- Philippe Aghion and Peter Howitt. A model of growth through creative destruction. *Econometrica*, 60(2):323–351, 1992. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/2951599>.
- Mavvrik AI. Achieving ai roi: Key findings from the 2025 forbes ai study. <https://www.mavvrik.ai/forbes-ai-study-2025/>, 2025.
- Adil S. Al-Busaidi, Raghu Raman, Laurie Hughes, Mousa Ahmed Albashrawi, Tegwen Malik, Yogesh K. Dwivedi, Thuraiya Al-Alawi, Mohammed AlRizeiqi, Gareth Davies, Mark Fenwick, Parul Gupta, Shashikala Gurple, Apeksha Hooda, Paulius Jurcys, Daryl Lim, Nicola Lucchi, Tanvi Misra, Ramakrishnan Raman, Anuragini Shirish, and Paul Walton. Redefining boundaries in innovation and knowledge domains: Investigating the impact of generative artificial intelligence on copyright and intellectual property rights. *Journal of Innovation & Knowledge*, 9(4):100630, 2024. doi: 10.1016/j.jik.2024.100630. URL <https://doi.org/10.1016/j.jik.2024.100630>.
- Joseph Amankwah-Amoah, Samar Abdalla, Emmanuel Mogaji, Amany Elbanna, and Yogesh K. Dwivedi. The impending disruption of creative industries by generative ai: Opportunities, challenges, and research agenda. *International Journal of Information Management*, 79:102759, 2024. doi: 10.1016/j.ijinfomgt.2024.102759. URL <https://doi.org/10.1016/j.ijinfomgt.2024.102759>.
- Amazon Web Services. Introducing aws healthscribe – automatically generate clinical notes from patient-clinician conversations using aws healthscribe. <https://aws.amazon.com/blogs/industries/industries-introducing-aws-healthscribe/>, 2023. AWS Industries Blog, Accessed: 2026-03-09.
- Anthropic. Project vend: Can claude run a small shop? (and why does that matter?). <https://www.anthropic.com/research/project-vend-1>, June 2025. Accessed: 2026-04-25.
- Anthropic. Claude Code by Anthropic | AI Coding Agent, Terminal, IDE, 2025a. URL <https://claude.com/product/claude-code>.
- Anthropic. Claude code, 2025b. URL <https://docs.anthropic.com/en/docs/claude-code>.
- Anthropic. Customer story | Airtree, 2026a. URL <https://claude.com/customers/airtree>.
- Anthropic. Cowork: Claude Code power for knowledge work | Claude by Anthropic, 2026b. URL <https://claude.com/product/cowork>.

- Anthropic. Customer story | Rakuten, 2026c. URL <https://claude.com/customers/rakuten>.
- Anthropic. Customer story | Wiz, 2026d. URL <https://claude.com/customers/wiz>.
- Anthropic. Customer story | Zapier Q&A, 2026e. URL <https://claude.com/customers/zapier-cowork-qa>.
- Anthropic. Claude Code for Enterprise | Claude by Anthropic, 2026f. URL <https://claude.com/product/claude-code/enterprise>.
- Anyphere. Cursor: The best way to code with AI, 2023. URL <https://cursor.com/>.
- Muhammad Fikry Aransyah, Bambang Hermanto, Anang Muftiadi, and Hera Oktadiana. Conversational AI in hospitality and tourism: a bibliometric–systematic review. *Cogent Business & Management*, 13(1):2613599, 2026. doi: 10.1080/23311975.2026.2613599. URL <https://doi.org/10.1080/23311975.2026.2613599>.
- Alexander Arnon. The projected impact of generative ai on future productivity growth, September 2025. URL <https://budgetmodel.wharton.upenn.edu/issues/2025/9/8/projected-impact-of-generative-ai-on-future-productivity-growth>.
- Azin Asgarian, Rohit Saha, Daniel Jakubovitz, and Julia Peyre. Autofraudnet: A multimodal network to detect fraud in the auto insurance industry. *arXiv preprint arXiv:2301.07526*, January 2023. doi: 10.48550/arXiv.2301.07526. URL <https://arxiv.org/abs/2301.07526>.
- Joshua Aslett, Thomas Cantens, François Chastel, Emmanuel A. Crown, and Stuart Hamilton. Generative artificial intelligence for compliance risk analysis: Applications in tax and customs administration. Technical notes and manuals no. 2025/013, International Monetary Fund, August 2025. URL <https://www.imf.org/-/media/files/publications/tnm/2025/english/tnmea2025013.pdf>. Technical guidance on using generative artificial intelligence (GenAI) for compliance risk analysis in tax and customs administration.
- Berk Atil, Alexa Chittams, Liseng Fu, Ferhan Ture, Lixinyu Xu, and Breck Baldwin. Non-determinism of "deterministic" llm settings. *arXiv (Cornell University)*, abs/2408.04667, 2024. doi: 10.48550/arxiv.2408.04667.
- Andrew Atkeson and Patrick J Kehoe. The transition to a new economy after the second industrial revolution. Working Paper 8676, National Bureau of Economic Research, December 2001. URL <http://www.nber.org/papers/w8676>.
- David H. Autor. Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30, September 2015. doi: 10.1257/jep.29.3.3. URL <https://www.aeaweb.org/articles?id=10.1257/jep.29.3.3>.
- David H. Autor, Frank Levy, and Richard J. Murnane. The skill content of recent technological change: An empirical exploration*. *The Quarterly Journal of Economics*, 118(4):1279–1333, 11 2003. ISSN 0033-5533. doi: 10.1162/003355303322552801. URL <https://doi.org/10.1162/003355303322552801>.
- Muhammad Awais, Ali Husain Salem Abdulla Alharthi, Amandeep Kumar, Hisham Cholakkal, and Rao Muhammad Anwer. Agrogpt: Efficient agricultural vision-language model with expert tuning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 5687–5696, 2025. doi: 10.1109/WACV61041.2025.00555. URL <https://doi.org/10.1109/WACV61041.2025.00555>.

- Tania Babina, Anastassia Fedyk, Alex He, and James Hodson. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745, 2024. ISSN 0304-405X. doi: <https://doi.org/10.1016/j.jfineco.2023.103745>. URL <https://www.sciencedirect.com/science/article/pii/S0304405X2300185X>.
- Axel Backlund and Lukas Petersson. Vending-bench: A benchmark for long-term coherence of autonomous agents. *CoRR*, abs/2502.15840, 2025. doi: 10.48550/ARXIV.2502.15840. URL <https://doi.org/10.48550/arXiv.2502.15840>.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- Paul Bergmann, Kilian Batzner, Michael Fauser, David Sattlegger, and Carsten Steger. The mvtec anomaly detection dataset: A comprehensive real-world dataset for unsupervised anomaly detection. *International Journal of Computer Vision*, 129(4):1038–1059, 2021. doi: 10.1007/s11263-020-01400-4.
- Fabien Bernier, Jun Cao, Maxime Cordy, and Salah Ghamizi. Powergraph-llm: Novel power grid graph embedding and optimization with large language models. *IEEE Transactions on Power Systems*, 40(6):5483–5486, November 2025. doi: 10.1109/TPWRS.2025.3596774. URL <https://doi.org/10.1109/TPWRS.2025.3596774>.
- Harold Booth, William Fisher, Ryan Galluzzo, and Joshua Roberts. Accelerating the adoption of software and artificial intelligence agent identity and authorization. Concept paper, National Institute of Standards and Technology, February 2026.
- Timothy F. Bresnahan and Shane Greenstein. Technological competition and the structure of the computer industry. *The Journal of Industrial Economics*, 47(1):1–40, 1999. doi: <https://doi.org/10.1111/1467-6451.00088>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-6451.00088>.
- Timothy F. Bresnahan and M. Trajtenberg. General purpose technologies ‘engines of growth’? *Journal of Econometrics*, 65(1):83–108, 1995. ISSN 0304-4076. doi: [https://doi.org/10.1016/0304-4076\(94\)01598-T](https://doi.org/10.1016/0304-4076(94)01598-T). URL <https://www.sciencedirect.com/science/article/pii/S030440769401598T>.
- Timothy F. Bresnahan, Erik Brynjolfsson, and Lorin M. Hitt. Information technology, workplace organization and the demand for skilled labor: Firm-level evidence. NBER Working Papers 7136, National Bureau of Economic Research, Inc, May 1999. URL <https://ideas.repec.org/p/nbr/nberwo/7136.html>.
- Erik Brynjolfsson and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company, New York, 2014.
- Erik Brynjolfsson, Daniel Rock, and Chad Syverson. Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics, February 2018. URL <https://ideas.repec.org/h/nbr/nberch/14007.html>.
- Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. Generative ai at work*. *The Quarterly Journal of Economics*, 140(2):889–942, 02 2025. ISSN 0033-5533. doi: 10.1093/qje/qjae044. URL <https://doi.org/10.1093/qje/qjae044>.

- Tianshi Cai, Guanxu Li, Nijia Han, Ce Huang, Zimu Wang, Changyu Zeng, Yuqi Wang, Jingshi Zhou, Haiyang Zhang, Qi Chen, et al. Findebate: Multi-agent collaborative intelligence for financial analysis. *arXiv preprint arXiv:2509.17395*, 2025. doi: 10.48550/arXiv.2509.17395. URL <https://arxiv.org/abs/2509.17395>.
- Yupeng Cao, Zhi Chen, Qingyun Pei, Fabrizio Dimino, Lorenzo Ausiello, Prashant Kumar, K. P. Subbalakshmi, and Papa Momar Ndiaye. Risklabs: Predicting financial risk using large language model based on multi-sources data. *arXiv preprint arXiv:2404.07452*, 2024. doi: 10.48550/arXiv.2404.07452. URL <https://arxiv.org/abs/2404.07452>.
- Junying Chen, Zhenyang Cai, Ke Ji, Xidong Wang, Wanlong Liu, Rongsheng Wang, Jianye Hou, and Benyou Wang. Huatuogpt-o1, towards medical complex reasoning with llms. *arXiv (Cornell University)*, abs/2412.18925, 2024a. doi: 10.48550/arxiv.2412.18925.
- Kaiyuan Chen, Yixin Ren, Yang Liu, Xiaobo Hu, Haotong Tian, Tianbao Xie, Fangfu Liu, Haoye Zhang, Hongzhang Liu, Yuan Gong, et al. xbench: Tracking agents productivity scaling with profession-aligned real-world evaluations. *arXiv preprint arXiv:2506.13651*, 2025a.
- Xiancai Chen, Zhengwei Tao, Kechi Zhang, Changzhi Zhou, Xinyu Zhang, Wanli Gu, Yuanpeng He, Mengdi Zhang, Xunliang Cai, Haiyan Zhao, and Zhi Jin. Revisit self-debugging with self-generated tests for code generation. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar, editors, *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, pages 18003–18023. Association for Computational Linguistics, 2025b. URL <https://aclanthology.org/2025.acl-long.881/>.
- Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. Teaching large language models to self-debug. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024b. URL <https://openreview.net/forum?id=KuPixIqPiq>.
- Zhilong Chen, Chengzong Zhao, Boyuan Chen, Dayi Lin, Yihao Chen, Arthur Leung, Gopi Krishnan Rajbahadur, Gustavo Ansaldi Oliva, Haoxiang Zhang, Aaditya Bhatia, Chong Chun Yong, and Ahmed E. Hassan. Repoforge: Training a SOTA fast-thinking SWE agent with an end-to-end data curation pipeline synergizing SFT and RL at scale. *CoRR*, abs/2508.01550, 2025c. doi: 10.48550/ARXIV.2508.01550. URL <https://doi.org/10.48550/arXiv.2508.01550>.
- Zhiyuan Chen, Shiqi Shen, Guangyao Shen, Gong Zhi, Xu Chen, and Yankai Lin. Towards tool use alignment of large language models. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 1382–1400. Association for Computational Linguistics, 2024c. doi: 10.18653/V1/2024.EMNLP-MAIN.82. URL <https://doi.org/10.18653/v1/2024.emnlp-main.82>.
- Yuheng Cheng, Huan Zhao, Xiyuan Zhou, Junhua Zhao, Yuji Cao, and Chao Yang. Gaia – a large language model for advanced power dispatch. *arXiv preprint arXiv:2408.03847*, 2024. doi: 10.48550/arXiv.2408.03847. URL <https://arxiv.org/abs/2408.03847>.
- Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. Mem0: Building production-ready ai agents with scalable long-term memory. *arXiv preprint arXiv:2504.19413*, 2025.

- Andrew Cho, Jason Woo, Brian Shi, Aishwaryaa Udeshi, and Jonathan Woo. The application of matec (multi-ai agent team care) framework in sepsis care. *arXiv (Cornell University)*, abs/2503.16433, 2025. doi: 10.48550/arxiv.2503.16433.
- Frederik Roland Christiansen, Linus Nørgaard Hollensberg, Nikoline Bach Fleng Jensen, Kristian Julsgaard, Kristian Nyborg Jespersen, and Ivan Adriyanov Nikolov. Exploring presence in interactions with LLM-driven NPCs: A comparative study of speech recognition and dialogue options. In *Proceedings of the 30th ACM Symposium on Virtual Reality Software and Technology*, pages 1–11, 2024. doi: 10.1145/3641825.3687716. URL <https://doi.org/10.1145/3641825.3687716>.
- Panayiotis Christou, Md. Zahidul Islam, Yuzhang Lin, and Jingwei Xiong. LLM4DistReconfig: A fine-tuned large language model for power distribution network reconfiguration. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4136–4155, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.naacl-long.208. URL <https://aclanthology.org/2025.naacl-long.208/>.
- Yucheng Chu, Hang Li, Kaiqi Yang, Yasemin Copur-Gencturk, and Jiliang Tang. Llm-based automated grading with human-in-the-loop. *arXiv preprint arXiv:2504.05239*, 2025a. doi: 10.48550/arXiv.2504.05239. URL <https://arxiv.org/abs/2504.05239>. Accepted to IEEE TALE 2025.
- Zhendong Chu, Shen Wang, Jian Xie, Tinghui Zhu, Yibo Yan, Jinheng Ye, Aoxiao Zhong, Xuming Hu, Jing Liang, Philip S. Yu, and Qingsong Wen. Llm agents for education: Advances and applications. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 13782–13810, Suzhou, China, 2025b. Association for Computational Linguistics. URL <https://aclanthology.org/2025.findings-emnlp.743/>.
- Jiseong Chung, Ronny Ko, Wonchul Yoo, Makoto Onizuka, Sungmok Kim, Tae-Wan Kim, and Won-Yong Shin. Graphcompliance: Aligning policy and context graphs for llm-based regulatory compliance. *arXiv preprint arXiv:2510.26309*, 2025.
- Thibo Cliteur and Dries Bertrand. Leveraging ai for customs classification purposes. WCO News 102 – Issue 3 / 2023 (Point of View article), October 2023. URL <https://mag.wcoomd.org/magazine/wco-news-102-issue-3-2023/leveraging-ai-for-customs-classification-purposes/>. Article on how artificial intelligence can support Customs classification of products under the Harmonized System (HS) and improve trade compliance.
- Pierre Colombo, Telmo Pessoa Pires, Malik Boudiaf, Dominic Culver, Rui Melo, Caio Corro, André F. T. Martins, Fabrizio Esposito, Vera Lúcia Raposo, Sofia Morgado, and Michael Desa. Saullm-7b: A pioneering large language model for law. *CoRR*, abs/2403.03883, 2024. doi: 10.48550/ARXIV.2403.03883. URL <https://doi.org/10.48550/arXiv.2403.03883>.
- Jiaxi Cui, Zongjian Li, Yang Yan, Bohua Chen, and Li Yuan. Chatlaw: Open-source legal large language model with integrated external knowledge bases. *CoRR*, abs/2306.16092, 2023. doi: 10.48550/ARXIV.2306.16092. URL <https://doi.org/10.48550/arXiv.2306.16092>.
- Zikun Cui, Tianyi Huang, Chia-En Chiang, and Cuiqianhe Du. Toward verifiable misinformation detection: A multi-tool llm agent framework. In *Proceedings of the 2025 International Conference on Generative Artificial Intelligence for Business*, pages 179–185, 2025.

- Cyberhaven Labs. 2025 ai adoption and risk report. Technical report, Cyberhaven, 2025. URL <https://info.cyberhaven.com/hubfs/Content%20PDF/Cyberhaven%20Labs%20-%202025%20AI%20Adoption%20%26%20Risk%20Report.pdf>. Accessed: 2026-04-12.
- Yuhao Dan, Zhikai Lei, Yiyang Gu, Yong Li, Jianghao Yin, Jiaju Lin, Linhao Ye, Zhiyan Tie, Yougen Zhou, Yilei Wang, Aimin Zhou, Ze Zhou, Qin Chen, Jie Zhou, Liang He, and Xipeng Qiu. Educhat: A large-scale language model-based chatbot system for intelligent education. *arXiv*, 2023.
- Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/67d57c32e20fd0a7a302cb81d36e40d5-Abstract-Conference.html.
- Paul David. The dynamo and the computer: An historical perspective on the modern productivity paradox. *American Economic Review*, 80:355–61, 01 1990.
- Joao Victor de Castro Oliveira, Nathan Rufino, Silvio Levcovitz, Matheus Yasuo Ribeiro Utino, Fábio Lobato, Marcelo Turine, Solange Rezende, Antonio Fernando Lavareda Jacob Junior, and Ricardo Marcondes Marcacini. Agents4gov: Privacy-preserving web agents using open llms for public sector tasks. In *Encontro Nacional de Inteligência Artificial e Computacional (ENIAC)*, pages 580–591. SBC, 2025.
- DeepSeek-AI, Qihao Zhu, Daya Guo, Zhihong Shao, Dejian Yang, Peiyi Wang, Runxin Xu, Y. Wu, Yukun Li, Huazuo Gao, Shirong Ma, Wangding Zeng, Xiao Bi, Zihui Gu, Hanwei Xu, Damai Dai, Kai Dong, Liyue Zhang, Yishi Piao, Zhibin Gou, Zhenda Xie, Zhewen Hao, Bingxuan Wang, Junxiao Song, Deli Chen, Xin Xie, Kang Guan, Yuxiang You, Aixin Liu, Qiushi Du, Wenjun Gao, Xuan Lu, Qinyu Chen, Yaohui Wang, Chengqi Deng, Jiashi Li, Chenggang Zhao, Chong Ruan, Fuli Luo, and Wenfeng Liang. Deepseek-coder-v2: Breaking the barrier of closed-source models in code intelligence. *CoRR*, abs/2406.11931, 2024. doi: 10.48550/ARXIV.2406.11931. URL <https://doi.org/10.48550/arXiv.2406.11931>.
- Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier. Padim: A patch distribution modeling framework for anomaly detection and localization. *Lecture notes in computer science*, pages 475–489, 2020. doi: 10.1007/978-3-030-68799-1_35.
- Iris Delikoura, Yi R. (May) Fung, and Pan Hui. From superficial outputs to superficial learning: Risks of large language models in education. *arXiv preprint arXiv:2509.21972*, 2025. doi: 10.48550/arXiv.2509.21972. URL <https://arxiv.org/abs/2509.21972>.
- Fabrizio Dell’Acqua, Edward McFowland III, and Ethan Mollick. Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. Working Paper 24-013, Harvard Business School, 2023. URL <https://www.hbs.edu/faculty/Pages/item.aspx?num=64700>.
- Arbaaz Dharmavaram and Saqib Hakak. SANCTUARY: An efficient evidence-based automated fact checking system. In Mubashara Akhtar, Rami Aly, Christos Christodoulopoulos, Oana Cocarascu, Zhijiang Guo, Arpit Mittal, Michael Schlichtkrull, James Thorne, and Andreas Vlachos, editors, *Proceedings of the Eighth Fact Extraction and VERification Workshop (FEVER)*, pages 247–257, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 978-1-959429-53-1. doi: 10.18653/v1/2025.fever-1.19. URL <https://aclanthology.org/2025.fever-1.19/>.

- Yifei Dong, Fengyi Wu, Kunlin Zhang, Yilong Dai, Sanjian Zhang, Wanghao Ye, Sihan Chen, and Zhi-Qi Cheng. Large language model agents in finance: A survey bridging research, practice, and real-world deployment. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 17889–17907, Suzhou, China, 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.findings-emnlp.972. URL <https://aclanthology.org/2025.findings-emnlp.972/>.
- Shihan Dou, Yan Liu, Haoxiang Jia, Limao Xiong, Enyu Zhou, Wei Shen, Junjie Shan, Caishuang Huang, Xiao Wang, Xiaoran Fan, Zhiheng Xi, Yuhao Zhou, Tao Ji, Rui Zheng, Qi Zhang, Xuanjing Huang, and Tao Gui. StepCoder: Improve code generation with reinforcement learning from compiler feedback. *CoRR*, abs/2402.01391, 2024. doi: 10.48550/ARXIV.2402.01391. URL <https://doi.org/10.48550/arXiv.2402.01391>.
- Yuwei Du, Jun Zhang, Jie Feng, Zhicheng Liu, Jian Yuan, and Yong Li. TrafficSimAgent: A hierarchical agent framework for autonomous traffic simulation with MCP control. *arXiv preprint arXiv:2512.20996*, 2025. doi: 10.48550/arXiv.2512.20996. URL <https://arxiv.org/abs/2512.20996>.
- Gabriel Dulac-Arnold, Nir Levine, Daniel Mankowitz, Jerry Li, Cosmin Paduraru, Sven Gowal, and Todd Hester. Challenges of real-world reinforcement learning: definitions, benchmarks and analysis. *Machine Learning*, 110(9):2419–2468, 2021. doi: 10.1007/s10994-021-05961-4.
- Barry Elad. Ai in finance statistics by market size, adoption, use cases, trends and facts (2025). <https://electroiq.com/stats/ai-in-finance-statistics/>, 11 2025.
- Schneider Electric. Building sustainability’s digital future with ecostruxure™ resource advisor copilot: Schneider electric’s latest ai advancement, April 2024. URL <https://www.se.com/eg/en/about-us/newsroom/news/press-releases/building-sustainability%E2%80%99s-digital-future-with-ecostruxure%E2%84%A2-resource-advisor-copilot-schneider-electric%E2%80%99s-latest-ai-advancement-6613e0bafbe2d764a03d066/>. Press release on the launch of a conversational AI tool embedded in EcoStruxure Resource Advisor Copilot for sustainability and energy data management.
- Lutfi Eren Erdogan, Nicholas Lee, Sehoon Kim, Suhong Moon, Hiroki Furuta, Gopala Anumanchipalli, Kurt Keutzer, and Amir Gholami. Plan-and-act: Improving planning of agents for long-horizon tasks. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu, editors, *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*, volume 267 of *Proceedings of Machine Learning Research*. PMLR / OpenReview.net, 2025. URL <https://proceedings.mlr.press/v267/erdogan25a.html>.
- Patricia Marcella Evite, Ekaterina Svetlova, and Doina Bucur. Trade-offs in financial ai: Explainability in a trilemma with accuracy and compliance. *arXiv preprint arXiv:2602.01368*, 2026. doi: 10.48550/arXiv.2602.01368. URL <https://arxiv.org/abs/2602.01368>.
- Diana Carolina Fajardo-Ramos, Andrés Chiappe, and Javier Mella-Norambuena. Human-in-the-loop assessment with ai: implications for teacher education in ibero-american universities. In *Frontiers in Education*, volume 10, page 1710992. Frontiers Media SA, 2025.
- Zhiwei Fei, Songyang Zhang, Xiaoyu Shen, Dawei Zhu, Xiao Wang, Jidong Ge, and Vincent Ng. Internlm-law: An open-sourced chinese legal large language model. In Owen Rambow, Leo Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, and Steven Schockaert, editors,

- Proceedings of the 31st International Conference on Computational Linguistics, COLING 2025, Abu Dhabi, UAE, January 19-24, 2025*, pages 9376–9392. Association for Computational Linguistics, 2025. URL <https://aclanthology.org/2025.coling-main.629/>.
- Weibo Gao, Qi Liu, Linan Yue, Fangzhou Yao, Rui Lv, Zheng Zhang, Hao Wang, and Zhenya Huang. Agent4edu: Generating learner response data by generative agents for intelligent education systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 23923–23932, April 2025. doi: 10.1609/aaai.v39i22.34565. URL <https://doi.org/10.1609/aaai.v39i22.34565>.
- Leon Garza, Lavanya Elluri, Anantaa Kotal, Aritran Piplai, Deepti Gupta, and Anupam Joshi. Privcomp-kg: Leveraging knowledge graph and large language models for privacy policy compliance verification. *arXiv preprint arXiv:2404.19744*, 2024.
- Margot Geerts, Manon Reusens, Bart Baesens, Seppe vanden Broucke, and Jochen De Weerd. On the performance of LLMs for real estate appraisal. *CoRR*, abs/2506.11812, 2025. URL <https://arxiv.org/abs/2506.11812>.
- Wenting Geng. Could large language models work as post-hoc explainability tools in credit risk model governance? *arXiv preprint arXiv:2602.18895*, 2026. doi: 10.48550/arXiv.2602.18895. URL <https://arxiv.org/abs/2602.18895>.
- Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. More than you’ve asked for: A comprehensive analysis of novel prompt injection threats to application-integrated large language models. *CoRR*, abs/2302.12173, 2023. doi: 10.48550/ARXIV.2302.12173. URL <https://doi.org/10.48550/arXiv.2302.12173>.
- Shuchen Guo, Ehsan Latif, Yifan Zhou, Xuan Huang, and Xiaoming Zhai. Using generative ai and multi-agents to provide automatic feedback. *arXiv preprint arXiv:2411.07407*, November 2024. doi: 10.48550/arXiv.2411.07407. URL <https://arxiv.org/abs/2411.07407>.
- Yilun Hao, Yongchao Chen, Yang Zhang, and Chuchu Fan. Large language models can solve real-world planning rigorously with formal verification tools. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3434–3483, 2025.
- Liu He, Yizhi Song, Hejun Huang, Pinxin Liu, Yunlong Tang, Daniel Aliaga, and Xin Zhou. Kubrick: Multimodal agent collaborations for synthetic video generation. *arXiv preprint arXiv:2408.10453*, 2024a.
- Zhitao He, Pengfei Cao, Chenhao Wang, Zhuoran Jin, Yubo Chen, Jiexin Xu, Huaijun Li, Kang Liu, and Jun Zhao. Agentscourt: Building judicial decision-making agents with court debate simulation and legal knowledge augmentation. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Findings of the Association for Computational Linguistics: EMNLP 2024, Miami, Florida, USA, November 12-16, 2024*, volume EMNLP 2024 of *Findings of ACL*, pages 9399–9416. Association for Computational Linguistics, 2024b. doi: 10.18653/V1/2024.FINDINGS-EMNLP.549. URL <https://doi.org/10.18653/v1/2024.findings-emnlp.549>.
- Elhanan Helpman. *General purpose technologies and economic growth*. MIT press, 1998.
- Elhanan Helpman and Manuel Trajtenberg. A time to sow and a time to reap: Growth based on general purpose technologies. NBER Working Paper 4854, National Bureau of Economic Research, September 1994. Available at SSRN: <https://ssrn.com/abstract=226530>.

- Henrique Henriques and Luis Nobre Pereira. Hotel demand forecasting models and methods using artificial intelligence: A systematic literature review. *Tourism & Management Studies*, 20(3):39–51, 2024. doi: 10.18089/tms.20240304. URL <https://doi.org/10.18089/tms.20240304>.
- Honeywell. Honeywell unveils innovative ai assistant for industrial operators in honeywell forge production intelligence. Press Release, February 2025. URL <https://www.honeywell.com/us/en/press/2025/02/honeywell-unveils-innovative-ai-assistant>. Honeywell announced a new generative AI assistant embedded in Honeywell Forge Production Intelligence to help industrial operators streamline tasks, automate troubleshooting and enhance production insights.
- André Hottung, Federico Berto, Chuanbo Hua, Nayeli Gast Zepeda, Daniel Wetzels, Michael Römer, Haoran Ye, Davide Zago, Michael Poli, Stefano Massaroli, Jinkyoo Park, and Kevin Tierney. Vrapagent: LLM-driven discovery of heuristic operators for vehicle routing problems. *arXiv preprint arXiv:2510.07073*, 2025. doi: 10.48550/arXiv.2510.07073. URL <https://arxiv.org/abs/2510.07073>.
- Mengkang Hu, Pu Zhao, Can Xu, Qingfeng Sun, Jian-Guang Lou, Qingwei Lin, Ping Luo, and Saravan Rajmohan. Agentgen: Enhancing planning abilities for large language model based agent via environment and task generation. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, pages 496–507, 2025.
- Xavier Hu, Jinxiang Xia, Shengze Xu, Kangqi Song, Yishuo Yuan, Guibin Zhang, JinCheng Ren, Boyu Feng, Li Lu, Tiejong Zeng, Jiaheng Liu, Minghao Liu, He Zhu, Yuchen Eleanor Jiang, Wei Wang, and Wangchunshu Zhou. Ecogym: Evaluating llms for long-horizon plan-and-execute in interactive economies, 2026. URL <https://arxiv.org/abs/2602.09514>.
- Quzhe Huang, Mingxu Tao, Zhenwei An, Chen Zhang, Cong Jiang, Zhibin Chen, Zirui Wu, and Yansong Feng. Lawyer llama technical report. *CoRR*, abs/2305.15062, 2023. doi: 10.48550/ARXIV.2305.15062. URL <https://doi.org/10.48550/arXiv.2305.15062>.
- Saffron Huang, Bryan Seethor, Esin Durmus, Kunal Handa, Miles McCain, Michael Stern, and Deep Ganguli. How ai is transforming work at anthropic, 2025. URL <https://anthropic.com/research/how-ai-is-transforming-work-at-anthropic/>.
- Chris Hughes. The agentic ai governance blind spot. <https://www.resilientcyber.io/p/the-agentic-ai-governance-blind-spot>, 2026.
- Atom Infinite. WisPaper - Your AI Scholar Search Engine | AI Literature Review, 2025. URL <https://www.wispaper.ai/en>.
- Infocomm Media Development Authority. Model ai governance framework for agentic ai. Technical report, Infocomm Media Development Authority, Singapore, January 2026.
- International Data Corporation (IDC). Worldwide spending on artificial intelligence forecast to reach \$632 billion in 2028, according to a new idc spending guide. Press Release, August 2024. URL <https://www.businesswire.com/news/home/20240819177906/en>. IDC forecasts global artificial intelligence (AI) spending, including generative AI, will more than double by 2028 when it is expected to reach \$632 billion with strong growth across industries and technology categories.
- Hasan Iqbal, Yuxia Wang, Minghan Wang, Georgi Nenkov Georgiev, Jiahui Geng, Iryna Gurevych, and Preslav Nakov. Openfactcheck: A unified framework for factuality evaluation of llms. In Delia Irazú Hernández Farías, Tom Hope, and Manling Li, editors, *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: EMNLP 2024 - System Demonstrations*,

- Miami, Florida, USA, November 12-16, 2024, pages 219–229. Association for Computational Linguistics, 2024. doi: 10.18653/V1/2024.EMNLP-DEMO.23. URL <https://doi.org/10.18653/v1/2024.emnlp-demo.23>.
- Md. Ashraful Islam, Mohammed Eunus Ali, and Md. Rizwan Parvez. Mapcoder: Multi-agent code generation for competitive problem solving. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 4912–4944. Association for Computational Linguistics, 2024. doi: 10.18653/V1/2024.ACL-LONG.269. URL <https://doi.org/10.18653/v1/2024.acl-long.269>.
- Naman Jain, Jaskirat Singh, Manish Shetty, Liang Zheng, Koushik Sen, and Ion Stoica. R2egym: Procedural environments and hybrid verifiers for scaling open-weights SWE agents. *CoRR*, abs/2504.07164, 2025. doi: 10.48550/ARXIV.2504.07164. URL <https://doi.org/10.48550/arXiv.2504.07164>.
- Minwoo Jeong, Jeeyun Chang, and Yoonjin Yoon. Speak to simulate: An LLM-guided agentic framework for traffic simulation in SUMO. In *Proceedings of the 8th ACM SIGSPATIAL International Workshop on GeoSpatial Simulation (GeoSIM '25)*, pages 45–48. Association for Computing Machinery, 2025. doi: 10.1145/3764921.3770151. URL <https://doi.org/10.1145/3764921.3770151>.
- Gongyao Jiang, Xinran Shi, and Qiong Luo. Jre-l: Journalist, reader, and editor llms in the loop for science journalism for the general audience. In *North American Chapter of the Association for Computational Linguistics*, 2025. URL <https://api.semanticscholar.org/CorpusID:275931878>.
- Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R Narasimhan. SWE-bench: Can language models resolve real-world github issues? In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=VTF8yNQm66>.
- Hongwei Jin, Kibaek Kim, and Jonghwan Kwon. Gridmind: LLMs-powered agents for power system analysis and operations. In *SC Workshops '25: Proceedings of the SC '25 Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 560–568, St. Louis, MO, USA, November 2025a. Association for Computing Machinery. doi: 10.1145/3731599.3767409. URL <https://doi.org/10.1145/3731599.3767409>.
- Jiajie Jin, Yuyao Zhang, Yimeng Xu, Hongjin Qian, Yutao Zhu, and Zhicheng Dou. Finsight: Towards real-world financial deep research. *arXiv preprint arXiv:2510.16844*, 2025b. doi: 10.48550/arXiv.2510.16844. URL <https://arxiv.org/abs/2510.16844>.
- Boyan Jovanovic and Peter L Rousseau. General purpose technologies. Working Paper 11093, National Bureau of Economic Research, January 2005. URL <http://www.nber.org/papers/w11093>.
- Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rodrigo Guerraoui D'Oliveira, Hubert Eichner, Salim El Rouayheb, David Evans, Judd Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Rachid Guerraoui, Zaid Harchaoui, Chaoyang He, Martin Jaggi, Swanand Kadhe, Ahmed Khaled, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Oluwasanmi Koyejo, Ananda Theertha Suresh Kumar, et al. Advances

- and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2): 1–210, 2021. doi: 10.1561/22000000083. URL <https://doi.org/10.1561/22000000083>.
- Eirini Kalliamvakou. Research: Quantifying github copilot’s impact on developer productivity and happiness, September 2022. URL <https://github.blog/news-insights/research/research-quantifying-github-copilots-impact-on-developer-productivity-and-happiness>.
- Naoaki Kanazawa, Kento Kawaharazuka, Yoshiki Obinata, Kei Okada, and Masayuki Inaba. Real-world cooking robot system from recipes based on food state recognition using foundation models and PDDL. *Advanced Robotics*, 38(18):1318–1334, 2024. doi: 10.1080/01691864.2024.2407136. URL <https://doi.org/10.1080/01691864.2024.2407136>.
- Deus F. Kandamali, Wesley M. Porter, Erin Porter, Alex McLemore, and Glen C. Rains. Cotton-Bot: An AI-driven cotton farming assistant and irrigation advisor using LLM-RAG and agentic AI tools. *Smart Agricultural Technology*, 12:101640, 2025. doi: 10.1016/j.atech.2025.101640. URL <https://doi.org/10.1016/j.atech.2025.101640>.
- Yan Ke, Xin Yu, Heming Du, Scott Chapman, and Helen Huang. Dynamic orchestration of multi-agent system for real-world multi-image agricultural vqa. *CoRR*, abs/2509.24350, 2025. doi: 10.48550/arXiv.2509.24350. URL <https://arxiv.org/abs/2509.24350>.
- Kees Hertogh. Meet microsoft dragon copilot: Your new ai assistant for clinical workflow. Microsoft Industry Blog, March 2025. URL <https://www.microsoft.com/en-us/industry/blog/healthcare/2025/03/03/meet-microsoft-dragon-copilot-your-new-ai-assistant-for-clinical-workflow/>. Announcement of Microsoft Dragon Copilot, an AI-powered assistant designed to streamline clinical workflows by automating documentation, surfacing key information, and reducing administrative burden for clinicians.
- Byungjun Kim, Minju Kim, Dayeon Seo, and Bugeun Kim. Leveraging large language models for active merchant non-player characters. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 10108–10116. International Joint Conferences on Artificial Intelligence Organization, August 2025. doi: 10.24963/ijcai.2025/1123. URL <https://doi.org/10.24963/ijcai.2025/1123>.
- Michal Kvet, Miroslav Potočár, and Slavomír Tatarka. Real estate attribute value extraction using large language models. *IEEE Access*, 13:13489–13506, 2025. doi: 10.1109/ACCESS.2025.3564511.
- Thomas Kwa, Ben West, Joel Becker, Amy Deng, Katharyn Garcia, Max Hasin, Sami Jawhar, Megan Kinniment, Nate Rush, Sydney Von Arx, et al. Measuring ai ability to complete long software tasks. *Advances in Neural Information Processing Systems*, 38:92213–92266, 2026.
- Siqi Lai, Zhao Xu, Weijia Zhang, Hao Liu, and Hui Xiong. LLMlight: Large language models as traffic signal control agents. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD ’25)*, pages 2335–2346. Association for Computing Machinery, 2025. doi: 10.1145/3690624.3709379. URL <https://doi.org/10.1145/3690624.3709379>.
- Wuyang Lan, Wenzheng Wang, Changwei Ji, Guoxing Yang, Yongbo Zhang, Xiaohong Liu, Song Wu, and Guangyu Wang. Clinicalgpt-r1: Pushing reasoning capability of generalist disease diagnosis with large language model. *arXiv (Cornell University)*, 2025. doi: <https://doi.org/10.48550/arxiv.2504.09421>.

- Sahaya Jestus Lazer, Kshitiz Aryal, Maanak Gupta, and Elisa Bertino. A survey of agentic AI and cybersecurity: Challenges, opportunities and use-case prototypes. *CoRR*, abs/2601.05293, 2026. doi: 10.48550/ARXIV.2601.05293. URL <https://doi.org/10.48550/arXiv.2601.05293>.
- Hung Le, Yue Wang, Akhilesh Deepak Gotmare, Silvio Savarese, and Steven Chu-Hong Hoi. Coderl: Mastering code generation through pretrained models and deep reinforcement learning. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/8636419dea1aa9fbd25fc4248e702da4-Abstract-Conference.html.
- LearnLM Team, Google. Learnlm: Improving gemini for learning. *arXiv preprint arXiv:2412.16429*, 2024. doi: 10.48550/arXiv.2412.16429. URL <https://arxiv.org/abs/2412.16429>.
- Dahyun Lee, Jonghyeon Choi, Jiyoung Han, and Kunwoo Park. Journalism-guided agentic in-context learning for news stance detection. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 15404–15427, 2025a.
- Jean Lee, Nicholas Stevens, et al. Large language models in finance: A survey. *Neural Computing and Applications*, 2024. doi: 10.1007/s00521-024-10495-6. URL <https://doi.org/10.1007/s00521-024-10495-6>.
- Unggi Lee, Jaeyong Lee, Jiyeong Bae, Yeil Jeong, Junbo Koh, Gyeonggeon Boaz Lee, Gunho Lee, Taekyung Ahn, and Hyeoncheol Kim. Pedagogy-r1: Pedagogical large reasoning model and well-balanced educational benchmark. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management, CIKM '25*, pages 1497–1507, Seoul, Republic of Korea, 2025b. Association for Computing Machinery. doi: 10.1145/3746252.3761133. URL <https://doi.org/10.1145/3746252.3761133>.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive NLP tasks. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>.
- Junkai Li, Siyu Wang, Meng Zhang, Weitao Li, Yunghwei Lai, Xinhui Kang, Weizhi Ma, and Yang Liu. Agent hospital: A simulacrum of hospital with evolvable medical agents. *arXiv (Cornell University)*, abs/2405.02957, 2024. doi: 10.48550/arxiv.2405.02957.
- Pei-Yun Lin and Yen-lung Tsai. Scorerag: A retrieval-augmented generation framework with consistency-relevance scoring and structured summarization for news generation. *CoRR*, abs/2506.03704, 2025. doi: 10.48550/ARXIV.2506.03704. URL <https://doi.org/10.48550/arXiv.2506.03704>.
- Xinyuan Lin, Chi Xu, Zhaoping Xiong, Xinfeng Zhang, Ningxi Ni, Bolin Ni, Jianlong Chang, Ruiqing Pan, Zidong Wang, Fan Yu, Qi Tian, Hualiang Jiang, Mingyue Zheng, and Nan Qiao. Pangu drug model: learn a molecule like a human. *Science China Life Sciences*, 66(4):879–882, 2022. doi: 10.1007/s11427-022-2239-y.

- Jiayu Liu, Zhenya Huang, Tong Xiao, Jing Sha, Jinze Wu, Qi Liu, Shijin Wang, and Enhong Chen. Socraticlm: exploring socratic personalized teaching with large language models. *Advances in Neural Information Processing Systems*, 37:85693–85721, 2024.
- Zhaowei Liu, Xin Guo, Zhi Yang, Fangqi Lou, Lingfeng Zeng, Jinyi Niu, Mengping Li, Qi Qi, Zhiqiang Liu, Yiyang Han, Dongpo Cheng, Ronghao Chen, Huacan Wang, Xingdong Feng, Huixia Judy Wang, Chengchun Shi, and Liwen Zhang. Fin-r1: A large language model for financial reasoning through reinforcement learning. *arXiv preprint arXiv:2503.16252*, 2025. doi: 10.48550/arXiv.2503.16252. URL <https://arxiv.org/abs/2503.16252>.
- Michael Luo, Naman Jain, Jaskirat Singh, Sijun Tan, Ameen Patel, Qingyang Wu, Alpay Ariyak, Colin Cai, Shang Zhu Tarun Venkat, Ben Athiwaratkun, Manan Roongta, Ce Zhang, Li Erran Li, Raluca Ada Popa, Koushik Sen, and Ion Stoica. DeepSWE: Training a state-of-the-art coding agent from scratch by scaling rl. <https://pretty-radio-b75.notion.site/DeepSWE-Training-a-Fully-Open-sourced-State-of-the-Art-Coding-Agent-by-Scaling-RL-2025>. Notion Blog.
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. Wizardcoder: Empowering code large language models with evol-instruct. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=UnUwSIgK5W>.
- Wanqin Ma, Chenyang Yang, and Christian Kästner. (why) is my prompt getting worse? rethinking regression testing for evolving llm apis. *Proceedings of the IEEE/ACM 3rd International Conference on AI Engineering - Software Engineering for AI*, pages 166–171, 2024. doi: 10.1145/3644815.3644950.
- Jakub Macina, Nico Daheim, Ido Hakimi, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. Math-tutorbench: A benchmark for measuring open-ended pedagogical capabilities of LLM tutors. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 204–221, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.11. URL <https://aclanthology.org/2025.emnlp-main.11/>.
- MarketsandMarkets. Artificial intelligence in manufacturing market worth \$155.04 billion by 2030 – exclusive report by marketsandmarkets™. Press Release, August 2025. URL <https://www.prnewswire.com/news-releases/artificial-intelligence-in-manufacturing-market-worth-155-04-billion-by-2030---exclusive-report-by-marketsandmarkets-2025-08-05.html>. The global AI in manufacturing market is projected to grow from USD 34.18 billion in 2025 to USD 155.04 billion by 2030, registering a CAGR of 35.3 % according to a new report by MarketsandMarkets™.
- Karl Marx. *Capital: Critique of Political Economy*, volume 1. 1867.
- Nestor Maslej, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Njenga Kariuki, Emily Capstick, Anka Reuel, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, Juan Carlos Niebles, Yoav Shoham, Russell Wald, Tobi Walsh, Armin Hamrah, Lapo Santarasci, Julia Betts Lotufo, Alexandra Rome, Andrew Shi, and Sukrut Oak. The ai index 2025 annual report, 2025. URL <https://hai.stanford.edu/ai-index/2025-ai-index-report>.

- Mantas Mazeika, Alice Gatti, Cristina Menghini, Udari Madhushani Sehwag, Shivam Singhal, Yury Orlovskiy, Steven Basart, Manasi Sharma, Denis Peskoff, Elaine Lau, Jaehyuk Lim, Lachlan Carroll, Alice Blair, Vinaya Sivakumar, Sumana Basu, Brad Kenstler, Yuntao Ma, Julian Michael, Xiaoke Li, Oliver Ingebrechtsen, Aditya Mehta, Jean Mottola, John Teichmann, Kevin Yu, Zaina Shaik, Adam Khoja, Richard Ren, Jason Hausenloy, Long Phan, Ye Htet, Ankit Aich, Tahseen Rabbani, Vivswan Shah, Andriy Novykov, Felix Binder, Kirill Chugunov, Luis Ramirez, Matias Geralnik, Hernán Mesura, Dean Lee, Ed-Yeremai Hernandez-Cardona, Annette Diamond, Summer Yue, Alexandr Wang, Bing Liu, Ernesto Hernandez, and Dan Hendrycks. Remote labor index: Measuring AI automation of remote work. *CoRR*, abs/2510.26787, 2025. doi: 10.48550/ARXIV.2510.26787. URL <https://doi.org/10.48550/arXiv.2510.26787>.
- McKinsey & Company. The end of inertia: Agentic ai’s disruption of retail and sme banking. McKinsey article / industry report, August 2025. URL <https://www.mckinsey.com/industries/financial-services/our-insights/the-end-of-inertia-agentic-ais-disruption-of-retail-and-sme-banking>.
- Mely.ai. Ai-powered customs automation solution, 2025. URL <https://mely.ai/solutions/customs/>. Accessed: 2026-03-09.
- Menlo Security. Menlo security’s 2025 report uncovers 68% surge in “shadow” generative ai usage in the modern enterprise. Technical report, Menlo Security, 2025. URL <https://www.menlosecurity.com/press-releases/menlo-securitys-2025-report-uncovers-68-surge-in-shadow-generative-ai-usage-in-the>. Accessed: 2026-04-12.
- Jiayi Miao, Zhanqing Su, and Xiaojun Xiao. Mmlm-ca: A multimodal large language model framework for automated insurance claims adjudication integrating policy document understanding and visual damage assessment. In *Proceedings of the 2025 3rd International Conference on Artificial Intelligence, Systems and Network Security (AISNS 2025)*, March 2026. doi: 10.1145/3797161.3797170. URL <https://doi.org/10.1145/3797161.3797170>.
- Microsoft. Dragon copilot: Ai assistant for clinical workflow. <https://www.microsoft.com/en-us/health-solutions/clinical-workflow/dragon-copilot>, 2025. Microsoft Cloud for Healthcare solution page. Accessed: 2026-03-09.
- Microsoft Corporation. Document intelligence invoice model. Microsoft Learn Documentation, 2025. URL <https://learn.microsoft.com/en-us/azure/ai-services/document-intelligence/prebuilt/invoice?view=doc-intel-4.0.0>. Documentation for Azure Document Intelligence’s prebuilt invoice model, which uses OCR to analyze and extract structured invoice fields such as customer name, billing address, due date, amount due, and line items across various formats and languages.
- Emanuele Musumeci, Michele Brienza, Vincenzo Suriani, Daniele Nardi, and Domenico Daniele Bloisi. Llm based multi-agent generation of semi-structured documents from semantic templates in the public administration domain. In *International Conference on Human-Computer Interaction*, pages 98–117. Springer, 2024.
- Prema Nedungadi, Kai-Yu Tang, and Raghu Raman. The transformative power of generative artificial intelligence for achieving the sustainable development goal of quality education. *Sustainability*, 16(22):9779, November 2024. doi: 10.3390/su16229779. URL <https://www.mdpi.com/2071-1050/16/22/9779>.

- Minheng Ni, Chenfei Wu, Huaying Yuan, Zhengyuan Yang, Ming Gong, Lijuan Wang, Zicheng Liu, Wangmeng Zuo, and Nan Duan. Autodirector: Online auto-scheduling agents for multi-sensory composition. *arXiv preprint arXiv:2408.11564*, 2024.
- Donghun Noh, Hyunwoo Nam, Kyle Gillespie, Yeting Liu, and Dennis Hong. Yummy operations robot initiative: An autonomous cooking system utilizing a modular robotic kitchen and dual-arm proprioceptive manipulator. *IEEE Robotics & Automation Magazine*, pages 2–14, 2025. doi: 10.1109/MRA.2025.3615353. URL <https://doi.org/10.1109/MRA.2025.3615353>.
- OECD. *Measuring Productivity – OECD Manual: Measurement of Aggregate and Industry-level Productivity Growth*. OECD Publishing, Paris, 2001. doi: 10.1787/9789264194519-en. Defines labor productivity as the ratio of volume measure of output to a measure of input use.
- OpenAI. Codex | AI Coding Partner from OpenAI, 2025. URL <https://openai.com/codex/>.
- OpenAI. Ai as a healthcare ally: How americans are navigating the system with chatgpt. Technical report, OpenAI, January 2026. URL <https://cdn.openai.com/pdf/2cb29276-68cd-4ec6-a5f4-c01c5e7a36e9/OpenAI-AI-as-a-Healthcare-Ally-Jan-2026.pdf>. Accessed: 2026-03-04.
- openclaw. Openclaw: Your own personal ai assistant. <https://github.com/openclaw/openclaw>, 2026. GitHub repository, starred 185,000+.
- Leyi Ouyang. Can memory-augmented llm agents aid journalism in interpreting and framing news for diverse audiences?, 2025. URL <https://arxiv.org/abs/2507.21055>.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/b1efde53be364a73914f58805a001731-Abstract-Conference.html.
- Charles Packer, Vivian Fang, Shishir G. Patil, Kevin Lin, Sarah Wooders, and Joseph E. Gonzalez. Memgpt: Towards llms as operating systems. 2023.
- Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, and Yizhe Zhang. Training software engineering agents and verifiers with swe-gym. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu, editors, *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*, volume 267 of *Proceedings of Machine Learning Research*. PMLR / OpenReview.net, 2025. URL <https://proceedings.mlr.press/v267/pan25g.html>.
- Joon Sung Park, Joseph C. O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. In Sean Follmer, Jeff Han, Jürgen Steimle, and Nathalie Henry Riche, editors, *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, UIST 2023, San Francisco, CA, USA, 29 October 2023- 1 November 2023*, pages 2:1–2:22. ACM, 2023. doi: 10.1145/3586183.3606763. URL <https://doi.org/10.1145/3586183.3606763>.

- Joon Sung Park, Carolyn Q. Zou, Aaron Shaw, Benjamin Mako Hill, Carrie J. Cai, Meredith Ringel Morris, Robb Willer, Percy Liang, and Michael S. Bernstein. Generative agent simulations of 1,000 people. *CoRR*, abs/2411.10109, 2024. doi: 10.48550/ARXIV.2411.10109. URL <https://doi.org/10.48550/arXiv.2411.10109>.
- Tejal Patwardhan, Rachel Dias, Elizabeth Proehl, Grace Kim, Michele Wang, Olivia Watkins, Simón Posada Fishman, Marwan Aljubei, Phoebe Thacker, Laurance Fauconnet, Natalie S. Kim, Patrick Chao, Samuel Miserendino, Gildas Chabot, David Li, Michael Sharman, Alexandra Barr, Amelia Glaese, and Jerry Tworek. Gdpval: Evaluating AI model performance on real-world economically valuable tasks. *CoRR*, abs/2510.04374, 2025. doi: 10.48550/ARXIV.2510.04374. URL <https://doi.org/10.48550/arXiv.2510.04374>.
- Joe Pazak and Manuel Baeuml. Transform supply chain logistics with agentic ai. AWS Blog (Industries Category), October 2025. URL <https://aws.amazon.com/cn/blogs/industries/transform-supply-chain-logistics-with-agentic-ai/>. AWS discusses how agentic AI and AI agents can be used to transform supply chain logistics operations by improving autonomy, responsiveness, and decision-making across inventory, warehousing, sourcing and other logistics workflows.
- Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. The impact of AI on developer productivity: Evidence from github copilot. *CoRR*, abs/2302.06590, 2023. doi: 10.48550/ARXIV.2302.06590. URL <https://doi.org/10.48550/arXiv.2302.06590>.
- Gaurab Pokharel, Shafkat Farabi, Patrick J Fowler, and Sanmay Das. Street-level ai: Are large language models ready for real-world judgments? In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 8, pages 2043–2054, 2025.
- Practical DevSecOps. Ai security statistics 2026: Latest data, trends & research, 2026. URL <https://www.practical-devsecops.com/ai-security-statistics-2026-research-report/>. Accessed: 2026-04-12.
- PricewaterhouseCoopers (PwC). Revolutionising customs with ai. PDF Report, PricewaterhouseCoopers, November 2024. URL <https://www.pwc.com/m1/en/publications/documents/2024/revolutionising-Customs-with-AI.pdf>. Thought leadership report discussing how artificial intelligence can transform customs operations by enhancing operational efficiency, strengthening security, and enabling digital transformation across the customs value chain.
- Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. Chatdev: Communicative agents for software development. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 15174–15186. Association for Computational Linguistics, 2024. doi: 10.18653/V1/2024.ACL-LONG.810. URL <https://doi.org/10.18653/v1/2024.acl-long.810>.
- Sebastian Raisch and Sebastian Krakowski. Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1):192–210, 2021. doi: 10.5465/amr.2018.0072.
- Petros Raptopoulos, Giorgos Filandrianos, Maria Lymperaioi, and Giorgos Stamou. PAKTON: A multi-agent framework for question answering in long legal agreements. In Christos Christodouloupoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng, editors, *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, EMNLP 2025, Suzhou,*

- China, November 4-9, 2025, pages 7948–7984. Association for Computational Linguistics, 2025. doi: 10.18653/V1/2025.EMNLP-MAIN.403. URL <https://doi.org/10.18653/v1/2025.emnlp-main.403>.
- Vipula Rawte, Amit P. Sheth, and Amitava Das. A survey of hallucination in large foundation models. *CoRR*, abs/2309.05922, 2023. doi: 10.48550/ARXIV.2309.05922. URL <https://doi.org/10.48550/arXiv.2309.05922>.
- Pavan Reddy and Aditya Gujral. Echoleak: The first real-world zero-click prompt injection exploit in a production llm system. *arXiv (Cornell University)*, abs/2509.10540, 2025. doi: 10.48550/arxiv.2509.10540.
- Saed Rezayi, Zhengliang Liu, Zihao Wu, Chandra Dhakal, Bao Ge, Chen Zhen, Tianming Liu, and Sheng Li. Agribert: Knowledge-infused agricultural language models for matching food and nutrition. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, pages 5150–5156, 2022. doi: 10.24963/ijcai.2022/715. URL <https://doi.org/10.24963/ijcai.2022/715>.
- Paul Romer. Endogenous technological change. Working Paper 3210, National Bureau of Economic Research, December 1989. URL <http://www.nber.org/papers/w3210>.
- Nathan Rosenberg. *Inside the Black Box: Technology and Economics*. Cambridge University Press, 1983.
- Karsten Roth. Towards total recall in industrial anomaly detection. *TIB Data Manager*, 2021. doi: 10.57702/mm9ulkog.
- Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach (4th Edition)*. Pearson, 2020. ISBN 9780134610993. URL <http://aima.cs.berkeley.edu/>.
- Md. Sadman Sakib and Yu Sun. From cooking recipes to robot task trees – improving planning correctness and task efficiency by leveraging LLMs with a knowledge network. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12704–12711. IEEE, 2024. doi: 10.1109/ICRA57147.2024.10611369. URL <https://doi.org/10.1109/ICRA57147.2024.10611369>.
- Dinesh Jackson Samuel, Inna Skarga-Bandurova, David Sikolia, and Muhammad Awais. Agrollm: Connecting farmers and agricultural practices through large language models for enhanced knowledge transfer and practical application. *CoRR*, abs/2503.04788, 2025. doi: 10.48550/arXiv.2503.04788. URL <https://arxiv.org/abs/2503.04788>.
- Marcelo Sandoval-Castaneda, Bryan Russell, Josef Sivic, Gregory Shakhnarovich, and Fabian Caba Heilbron. Editduet: A multi-agent system for video non-linear editing. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, pages 1–11, 2025.
- Mario Sanz-Guerrero and Javier Arroyo. Credit risk meets large language models: Building a risk indicator from loan descriptions in p2p lending. *arXiv preprint arXiv:2401.16458*, 2024. doi: 10.48550/arXiv.2401.16458. URL <https://arxiv.org/abs/2401.16458>.
- Juanming Shi, Qinglang Guo, Yong Liao, Yuxing Wang, Shijia Chen, and Shenglin Liang. Legal-lm: Knowledge graph enhanced large language models for law consulting. In De-Shuang Huang, Zhanjun Si, and Chuanlei Zhang, editors, *Advanced Intelligent Computing Technology and Applications - 20th International Conference, ICIC 2024, Tianjin, China, August 5-8, 2024, Proceedings, Part IV*

- (LNAI), volume 14878 of *Lecture Notes in Computer Science*, pages 175–186. Springer, 2024. doi: 10.1007/978-981-97-5672-8_15. URL https://doi.org/10.1007/978-981-97-5672-8_15.
- Parshin Shojaee, Aneesh Jain, Sindhu Tipirneni, and Chandan K. Reddy. Execution-based code generation using deep reinforcement learning. *Trans. Mach. Learn. Res.*, 2023, 2023. URL <https://openreview.net/forum?id=0XBuaxqEcG>.
- Siemens. Siemens xcelerator: Scaling roll-out of generative ai with siemens industrial copilot. Press Release, April 2024. URL <https://press.siemens.com/global/en/pressrelease/siemens-xcelerator-scaling-roll-out-generative-ai-siemens-industrial-copilot>. Press release on Siemens’ expansion of generative AI via the Siemens Industrial Copilot connected to the Siemens Xcelerator platform at Hannover Messe 2024.
- Siemens. Siemens introduces ai agents for industrial automation. Press Release, May 2025. URL <https://press.siemens.com/global/en/pressrelease/siemens-introduces-ai-agents-industrial-automation>. At Automate 2025 in Detroit, Siemens announced an expansion of its industrial AI offerings with advanced AI agents designed to operate autonomously within its Industrial Copilot ecosystem.
- Mandeep Singh, Anurag Rana, Nishant Chintala, Charles Shum, and Steven Tseng. Generative ai races toward \$1.3 trillion in revenue by 2032, March 2024. URL <https://www.bloomberg.com/professional/insights/data/generative-ai-races-toward-1-3-trillion-in-revenue-by-2032/>.
- Alex Singla, Alexander Sukharevsky, Lareina Yee, Michael Chui, Bryce Hall, and Tara Balakrishnan. The state of AI in 2025: Agents, innovation, and transformation, November 2025. URL <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>. McKinsey Global Survey on AI.
- Adam Smith. *The wealth of nations [1776]*, volume 11937. 1937.
- Robert M. Solow. Technical change and the aggregate production function. *The Review of Economics and Statistics*, 39(3):312–320, 1957. ISSN 00346535, 15309142. URL <http://www.jstor.org/stable/1926047>.
- Leon Staufer, Kevin Feng, Kevin Wei, Luke Bailey, Yawen Duan, Mick Yang, A. Pinar Ozisik, Stephen Casper, and Noam Kolt. The 2025 AI agent index: Documenting technical and safety features of deployed agentic AI systems. *CoRR*, abs/2602.17753, 2026. doi: 10.48550/ARXIV.2602.17753. URL <https://doi.org/10.48550/arXiv.2602.17753>.
- Kameshbhai Pravinbhai Suryavanshi and Sendhil kumar ks. Ai-driven financial risk management with llm agents and adaptive learning. ResearchGate Preprint, April 2025. URL https://www.researchgate.net/publication/390483692_AI-Driven_Financial_Risk_Management_with_LLM_Agents_and_Adaptive_Learning.
- Jakub Swacha and Michał Gracel. Retrieval-augmented generation (rag) chatbots for education: A survey of applications. *Applied Sciences*, 15(8):4234, 2025. doi: 10.3390/app15084234.
- Pilar Talón-Ballesteró, Marta Nieto-García, and Lydia González-Serrano. The wheel of dynamic pricing: towards open pricing and one to one pricing in hotel revenue management. *International Journal of Hospitality Management*, 102:103184, 2022. doi: 10.1016/j.ijhm.2022.103184. URL <https://doi.org/10.1016/j.ijhm.2022.103184>.

- Alex Tamkin and Peter McCrory. Estimating ai productivity gains from claude conversations, 2025. URL <https://www.anthropic.com/research/estimating-productivity-gains>.
- Technavio. Generative ai in manufacturing market: On-premises is expected to lead the deployment segment during 2025-2029. Press Release / Market Report Summary, August 2025. URL <https://newsroom.technavio.org/generative-ai-in-manufacturing-market>. The global generative AI in manufacturing market is expected to grow at a CAGR of 37.7 % from 2024 to 2029, driven by the need for accelerated product innovation, operational efficiency and human-AI collaboration in manufacturing environments.
- Vasilis Tsiavos and Fotis C. Kitsios. The digital transformation of the film industry: How artificial intelligence is changing the seventh art. *Telecommunications Policy*, 49(8):103021, 2025. doi: 10.1016/j.telpol.2025.103021. URL <https://doi.org/10.1016/j.telpol.2025.103021>.
- Oliver van Erven, Konstantinos Zafeirakis, Jacobus Smit, Julio Smidi, and Luc Buijs. [re] cooperate or collapse: Emergence of sustainable cooperation in a society of LLM agents. *Trans. Mach. Learn. Res.*, 2025, 2025. URL <https://openreview.net/forum?id=EWwXSkUch0>.
- Ivana Vichentijevikj, Kostadin Mishev, and Monika Misheva. Prompt-to-pill: Multi-agent drug discovery and clinical simulation pipeline. *Bioinformatics Advances*, 6(1):vbaf323–vbaf323, 2025. doi: 10.1093/bioadv/vbaf323.
- Bertie Vidgen, Abby Fennelly, Evan Pinnix, Chirag Mahapatra, Zach Richards, Austin Bridges, Calix Huang, Ben Hunsberger, Fez Zafar, Brendan Foody, Dominic Barton, Cass R. Sunstein, Eric Topol, and Osvald Nitski. The AI productivity index (APEX). *CoRR*, abs/2509.25721, 2025. doi: 10.48550/ARXIV.2509.25721. URL <https://doi.org/10.48550/arXiv.2509.25721>.
- Haochun Wang, Sendong Zhao, Zewen Qiang, Zijian Li, Chi Liu, Nuwa Xi, Yanrui Du, Bing Qin, and Ting Liu. Knowledge-tuning large language models with structured medical knowledge bases for trustworthy response generation in chinese. *ACM Trans. Knowl. Discov. Data*, August 2025a. ISSN 1556-4681. doi: 10.1145/3686807. URL <https://doi.org/10.1145/3686807>. Just Accepted.
- Liqiong Wang, Teng Jin, Jinyu Yang, Ales Leonardis, Fangyi Wang, and Feng Zheng. AgriLLaVA: Knowledge-infused large multimodal assistant on agricultural pests and diseases. *CoRR*, abs/2412.02158, 2024a. doi: 10.48550/arXiv.2412.02158. URL <https://arxiv.org/abs/2412.02158>.
- Qian Wang, Ziqi Huang, Ruoxi Jia, Paul Debevec, and Ning Yu. MAViS: A multi-agent framework for long-sequence video storytelling. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2273–2295, Rabat, Morocco, March 2026. Association for Computational Linguistics. doi: 10.18653/v1/2026.eacl-long.101. URL <https://aclanthology.org/2026.eacl-long.101/>.
- Saizhuo Wang, Hang Yuan, Leon Zhou, Lionel M. Ni, Heung-Yeung Shum, and Jian Guo. Alpha-gpt: Human-ai interactive alpha mining for quantitative investment. *arXiv preprint arXiv:2308.00016*, 2023. doi: 10.48550/arXiv.2308.00016. URL <https://arxiv.org/abs/2308.00016>.
- Shen Wang, Tianlong Xu, Hang Li, Chaoli Zhang, Joleen Liang, Jiliang Tang, Philip S. Yu, and Qingsong Wen. Large language models for education: A survey and outlook. *arXiv preprint arXiv:2403.18105*, 2024b. doi: 10.48550/arXiv.2403.18105. URL <https://arxiv.org/abs/2403.18105>.

- Weixuan Wang, Dongge Han, Daniel Diaz, Jin Xu, Victor Rühle, and Saravan Rajmohan. Odyssey-bench: Evaluating llm agents on long-horizon complex office application workflows. *arXiv (Cornell University)*, abs/2508.09124, 2025b. doi: 10.48550/arxiv.2508.09124.
- Xidong Wang, Jianquan Li, Shunian Chen, Yuxuan Zhu, Xiangbo Wu, Zhiyi Zhang, Xiaolong Xu, Junying Chen, Jie Fu, Xiang Wan, et al. Huatuo-26m, a large-scale chinese medical qa dataset. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 3828–3848, 2025c.
- Yiding Wang, Yuxuan Chen, Fanxu Meng, Xifan Chen, Xiaolei Yang, and Muhan Zhang. Law in silico: Simulating legal society with llm-based agents. *CoRR*, abs/2510.24442, 2025d. doi: 10.48550/ARXIV.2510.24442. URL <https://doi.org/10.48550/arXiv.2510.24442>.
- Zhefan Wang, Yuanqing Yu, Wendi Zheng, Weizhi Ma, and Min Zhang. Macrec: A multi-agent collaboration framework for recommendation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2760–2764, 2024c.
- Ziqi Wang and Boqin Yuan. L-MARS: legal multi-agent workflow with orchestrated reasoning and agentic search. *CoRR*, abs/2509.00761, 2025. doi: 10.48550/ARXIV.2509.00761. URL <https://doi.org/10.48550/arXiv.2509.00761>.
- Ziyue Wang, Junde Wu, Chang Low, Linghan Cai, Xihong Yang, Qiaxuan Li, and Yueming Jin. Medagent-pro: Towards multi-modal evidence-based medical diagnosis via reasoning agentic workflow. *arXiv (Cornell University)*, abs/2503.18968, 2025e. doi: 10.48550/arxiv.2503.18968.
- William Watson, Nicole Cho, Nishan Srishankar, Zhen Zeng, Lucas Cecchi, Daniel Scott, Suchetha Siddagangappa, Rachneet Kaur, Tucker Balch, and Manuela Veloso. LAW: legal agentic workflows for custody and fund services contracts. In Owen Rambow, Leo Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, Steven Schockaert, Kareem Darwish, and Apoorv Agarwal, editors, *Proceedings of the 31st International Conference on Computational Linguistics, COLING 2025 - Industry Track, Abu Dhabi, UAE, January 19-24, 2025*, pages 583–594. Association for Computational Linguistics, 2025. URL <https://aclanthology.org/2025.coling-industry.50/>.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html.
- Jiaxin Wen, Jian Guan, Hongning Wang, Wei Wu, and Minlie Huang. Codeplan: Unlocking reasoning potential in large language models by scaling code-form planning. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL <https://openreview.net/forum?id=dCPF1wlqj8>.
- Jingbo Wen and Baoxia Ma. Enhancing museum experience through deep learning and multimedia technology. *Heliyon*, 10(12):e32706, 2024. doi: 10.1016/j.heliyon.2024.e32706. URL <https://doi.org/10.1016/j.heliyon.2024.e32706>.
- Yihan (Logon) Wen and Xin Chen. X-gridagent: An LLM-powered agentic AI system for assisting power grid analysis. *arXiv preprint arXiv:2512.20789*, 2025. doi: 10.48550/arXiv.2512.20789. URL <https://arxiv.org/abs/2512.20789>.

- Hjalmar Wijk, Tao Roa Lin, Joel Becker, Sami Jawhar, Neev Parikh, Thomas Broadley, Lawrence Chan, Michael Chen, Joshua Clymer, Jai Dhyani, Elena Elicheva, Katharyn Garcia, Brian Goodrich, Nikola Jurkovic, Megan Kinniment, Aron Lajko, Seraphina Nix, Lucas Jun Koba Sato, William Saunders, Maksym Taran, Ben West, and Elizabeth Barnes. Re-bench: Evaluating frontier AI r&d capabilities of language model agents against human experts. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu, editors, *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*, volume 267 of *Proceedings of Machine Learning Research*. PMLR / OpenReview.net, 2025. URL <https://proceedings.mlr.press/v267/wijk25a.html>.
- World Economic Forum. Chatwto: An analysis of generative artificial intelligence and international trade. White Paper, September 2024. URL <https://www.weforum.org/publications/chatwto-an-analysis-of-generative-artificial-intelligence-and-international-trade/>. White paper exploring how generative artificial intelligence is transforming global trade, its opportunities, challenges and the need for international governance frameworks.
- Chaoyi Wu, Weixiong Lin, Xiaoman Zhang, Ya Zhang, Weidi Xie, and Yanfeng Wang. Pmc-llama: toward building open-source language models for medicine. *Journal of the American Medical Informatics Association*, 31(9):1833–1843, 2024. doi: 10.1093/jamia/ocae045.
- Jibang Wu, Chenghao Yang, Yi Wu, Simon Mahns, Chaoqi Wang, Hao Zhu, Fei Fang, and Haifeng Xu. AI realtor: Towards grounded persuasive language generation for automated copywriting, 2025a. URL <https://arxiv.org/abs/2502.16810>.
- Weijia Wu, Zeyu Zhu, and Mike Zheng Shou. Automated movie generation via multi-agent cot planning. *arXiv preprint arXiv:2503.07314*, 2025b.
- Zhenhua Wu, Zhuohao Chen, Di Zhu, Christos Mousas, and Dominic Kao. A systematic review of generative ai on game character creation: Applications, challenges, and future trends. *IEEE Transactions on Games*, pages 1–15, 2025c. doi: 10.1109/TG.2025.3564869. URL <https://doi.org/10.1109/TG.2025.3564869>. Early Access / in press.
- Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. The rise and potential of large language model based agents: A survey. *Science China Information Sciences*, 68(2):121101, 2025.
- Yijia Xiao, Edward Sun, Di Luo, and Wei Wang. Tradingagents: Multi-agents llm financial trading framework. *arXiv preprint arXiv:2412.20138*, 2024. doi: 10.48550/arXiv.2412.20138. URL <https://arxiv.org/abs/2412.20138>.
- Fei Xiong, Xiang Zhang, Aosong Feng, Siqi Sun, and Chenyu You. Quantagent: Price-driven multi-agent llms for high-frequency trading. *arXiv preprint arXiv:2509.09995*, 2025a. doi: 10.48550/arXiv.2509.09995. URL <https://arxiv.org/abs/2509.09995>.
- Weimin Xiong, Yifan Song, Qingxiu Dong, Bingchan Zhao, Feifan Song, Xun Wang, and Sujian Li. Mpo: Boosting llm agents with meta plan optimization. *arXiv preprint arXiv:2503.02682*, 2025b.
- Songlin Xu, Hao-Ning Wen, Hongyi Pan, Dallas Dominguez, Dongyin Hu, and Xinyu Zhang. Classroom simulacra: Building contextual student generative agents in online education for learning behavioral simulation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–26, 2025a.

- Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Wanxiang Che, Min Zhang, and Lidong Zhou. Layoutlmv2: Multi-modal pre-training for visually-rich document understanding. *arXiv (Cornell University)*, 2020. doi: 10.48550/arxiv.2012.14740.
- Zhenran Xu, Longyue Wang, Jifang Wang, Zhouyi Li, Senbao Shi, Xue Yang, Yiyu Wang, Baotian Hu, Jun Yu, and Min Zhang. Filmagent: A multi-agent framework for end-to-end film automation in virtual 3d spaces. *arXiv preprint arXiv:2501.12909*, 2025b.
- Sikuan Yan, Xiufeng Yang, Zuchao Huang, Ercong Nie, Zifeng Ding, Zonggen Li, Xiaowen Ma, Hinrich Schütze, Volker Tresp, and Yunpu Ma. Memory-r1: Enhancing large language model agents to manage and utilize memories via reinforcement learning. *CoRR*, abs/2508.19828, 2025. doi: 10.48550/ARXIV.2508.19828. URL <https://doi.org/10.48550/arXiv.2508.19828>.
- Bo Yang, Yunkui Chen, Lanfei Feng, Yu Zhang, Xiao Xu, Jianyu Zhang, Nueraili Aierken, Runhe Huang, Hongjian Lin, Yibin Ying, and Shijian Li. AgriGPT-VL: Agricultural vision-language understanding suite. OpenReview, 2025a. URL <https://openreview.net/forum?id=RZrLdTQSrg>. ICLR 2026 withdrawn submission.
- Bo Yang, Lanfei Feng, Yunkui Chen, Yu Zhang, Jianyu Zhang, Xiao Xu, Nueraili Aierken, and Shijian Li. AgriGPT-Omni: A unified speech-vision-text framework for multilingual agricultural intelligence. *CoRR*, abs/2512.10624, 2025b. doi: 10.48550/arXiv.2512.10624. URL <https://arxiv.org/abs/2512.10624>.
- Bo Yang, Yu Zhang, Lanfei Feng, Yunkui Chen, Jianyu Zhang, Xiao Xu, Nueraili Aierken, Yurui Li, Yuxuan Chen, Guijun Yang, Yong He, Runhe Huang, and Shijian Li. AgriGPT: a large language model ecosystem for agriculture. *CoRR*, abs/2508.08632, 2025c. doi: 10.48550/arXiv.2508.08632. URL <https://arxiv.org/abs/2508.08632>.
- Bo Yang, Yu Zhang, Yunkui Chen, Lanfei Feng, Xiao Xu, Nueraili Aierken, and Shijian Li. AgriAgent: Contract-driven planning and capability-aware tool orchestration in real-world agriculture. *CoRR*, abs/2601.08308, 2026a. doi: 10.48550/arXiv.2601.08308. URL <https://arxiv.org/abs/2601.08308>.
- Hongyang Yang, Boyu Zhang, Neng Wang, Cheng Guo, Xiaoli Zhang, Likun Lin, Junlin Wang, Tianyu Zhou, Mao Guan, Runjia Zhang, and Christina Dan Wang. Finrobot: An open-source ai agent platform for financial applications using large language models. *arXiv preprint arXiv:2405.14767*, 2024. doi: 10.48550/arXiv.2405.14767. URL <https://arxiv.org/abs/2405.14767>.
- Qianyu Yang, Yang Liu, Jiaqi Li, Jun Bai, Hao Chen, Kaiyuan Chen, Tiliang Duan, Jiayun Dong, Xiaobo Hu, Zixia Jia, et al. Onemillion-bench: How far are language agents from human experts? *arXiv preprint arXiv:2603.07980*, 2026b.
- Shuting Yang, Zehui Liu, Wolfgang Mayer, Ningpei Ding, Ying Wang, Yu Huang, Pengfei Wu, Wanli Li, Lin Li, Hong-Yu Zhang, and Zaiwen Feng. Shizishangpt: An agricultural large language model integrating tools and resources. In *Web Information Systems Engineering – WISE 2024*, volume 15439 of *Lecture Notes in Computer Science*, pages 284–298, Singapore, 2025d. Springer. doi: 10.1007/978-981-96-0573-6_21. URL https://doi.org/10.1007/978-981-96-0573-6_21.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R. Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL https://openreview.net/forum?id=WE_vluYUL-X.

- Yossi Matias, Aashima Gupta. Introducing medlm for the healthcare industry. Google Cloud Blog (Healthcare & Life Sciences), December 2023. URL <https://cloud.google.com/blog/topics/healthcare-life-sciences/introducing-medlm-for-the-healthcare-industry>. Blog announcing MedLM, a family of foundation models fine-tuned for healthcare use cases, now available to Google Cloud customers through Vertex AI platform.
- Shengbin Yue, Wei Chen, Siyuan Wang, Bingxuan Li, Chenchen Shen, Shujun Liu, Yuxuan Zhou, Yao Xiao, Song Yun, Xuanjing Huang, and Zhongyu Wei. Disc-lawllm: Fine-tuning large language models for intelligent legal services. *CoRR*, abs/2309.11325, 2023. doi: 10.48550/ARXIV.2309.11325. URL <https://doi.org/10.48550/arXiv.2309.11325>.
- An Zhang, Yuxin Chen, Leheng Sheng, Xiang Wang, and Tat-Seng Chua. On generative agents in recommendation. In *Proceedings of the 47th international ACM SIGIR conference on research and development in Information Retrieval*, pages 1807–1817, 2024a.
- H. G. Zhang, Peter Eckmann, J. Miao, A. B. Mahon, and James Zou. The virtual biotech: A multi-agent ai framework for therapeutic discovery and development. *bioRxiv*, February 2026a. bioRxiv preprint.
- Mingqing Zhang, Zhuoning Xu, Peijie Wang, Rongji Li, Liang Wang, Qiang Liu, Jian Xu, Xuyao Zhang, Shu Wu, and Liang Wang. AgriDoctor: A multimodal intelligent assistant for agriculture. *CoRR*, abs/2509.17044, 2025a. doi: 10.48550/arXiv.2509.17044. URL <https://arxiv.org/abs/2509.17044>.
- Ni Zhang, Zhiguang Cao, Jianan Zhou, Cong Zhang, and Yew-Soon Ong. An agentic framework with LLMs for solving complex vehicle routing problems. In *The Fourteenth International Conference on Learning Representations*, 2026b. URL <https://openreview.net/forum?id=BM0gYw4EhQ>. Poster.
- Wentao Zhang, Lingxuan Zhao, Haochong Xia, Shuo Sun, Jiaze Sun, Molei Qin, Xinyi Li, Yuqing Zhao, Yilei Zhao, Xinyu Cai, Longtao Zheng, Xinrun Wang, and Bo An. A multimodal foundation agent for financial trading: Tool-augmented, diversified, and generalist. *arXiv preprint arXiv:2402.18485*, 2024b. doi: 10.48550/arXiv.2402.18485. URL <https://arxiv.org/abs/2402.18485>.
- Yan Zhang, Ahmad Mohammad Saber, Amr Youssef, and Deepa Kundur. Grid-agent: An LLM-powered multi-agent system for power grid control. *arXiv preprint arXiv:2508.05702*, 2025b. doi: 10.48550/arXiv.2508.05702. URL <https://arxiv.org/abs/2508.05702>.
- Yuqin Zhang, Qijie Fan, Xuan Chen, Min Li, Zeying Zhao, Fuzhong Li, and Leifeng Guo. IPM-AgriGPT: A large language model for pest and disease management with a G-EA framework and agricultural contextual reasoning. *Mathematics*, 13(4):566, 2025c. doi: 10.3390/math13040566. URL <https://doi.org/10.3390/math13040566>.
- Li Zhao, Rui Sun, Zuoyou Jiang, Bo Yang, Yuxiao Bai, Mengting Chen, Xinyang Wang, Jing Li, and Zuo Bai. Contesttrade: A multi-agent trading system based on internal contest mechanism. *arXiv preprint arXiv:2508.00554*, 2025. doi: 10.48550/arXiv.2508.00554. URL <https://arxiv.org/abs/2508.00554>.
- Yu Zhao and Haoxiang Gao. Utilizing large language models for information extraction from real estate transactions. *CoRR*, abs/2404.18043, 2024. doi: 10.48550/ARXIV.2404.18043. URL <https://doi.org/10.48550/arXiv.2404.18043>.

- Enyu Zhou, Zhiheng Xi, Long Ma, Zhihao Zhang, Shihan Dou, Zhikai Lei, Guoteng Wang, Rui Zheng, Hang Yan, Tao Gui, et al. Steering llms via scalable interactive oversight. *arXiv preprint arXiv:2602.04210*, 2026.
- Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. Webarena: A realistic web environment for building autonomous agents. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024a. URL <https://openreview.net/forum?id=oKn9c6ytLx>.
- Zhi Zhou, Jiang-Xin Shi, Peng-Xiao Song, Xiaowen Yang, Yi-Xuan Jin, Lan-Zhe Guo, and Yufeng Li. Lawgpt: A chinese legal knowledge-enhanced large language model. *CoRR*, abs/2406.04614, 2024b. doi: 10.48550/ARXIV.2406.04614. URL <https://doi.org/10.48550/arXiv.2406.04614>.
- Jie Zhu, Qian Chen, Huaixia Dou, Junhui Li, Lifan Guo, Feng Chen, and Chi Zhang. Dianjin-r1: Evaluating and enhancing financial reasoning in large language models. *arXiv preprint arXiv:2504.15716*, 2025. doi: 10.48550/arXiv.2504.15716. URL <https://arxiv.org/abs/2504.15716>.