

The Impact of Generative AI on Individual and Society: Prospects and Future Possibilities

Zhenyu Wang¹, Yuqing Li², Haoran Xu³, Meilin Chen⁴ *

¹ School of Economics and Management, Beijing Forestry University, No. 35 Qinghua East Road, Haidian District, Beijing 100083, China

² School of Public Administration, Hunan Normal University, 36 Lushan Road, Yuelu District, Changsha, Hunan 410081, China

³ School of Management, Hangzhou Dianzi University, No. 1158 Wenzhe Street, Xiasha Higher Education Zone, Hangzhou, Zhejiang 310018, China

⁴ School of Economics and Management, Shanxi University, No. 92 Wucheng Road, Taiyuan, Shanxi 030006, China

*Email: meilinchen@sxu.edu.cn (Corresponding Author)

Abstract

This article develops a non-review, evidence-based analytical argument about how generative artificial intelligence is reshaping personal capability, organizational practice, and the architecture of society. Rather than merely cataloguing prior studies, the paper advances a socio-technical thesis: generative AI is becoming a general-purpose layer of cognitive infrastructure that reconfigures how people access knowledge, produce content, make decisions, coordinate work, and negotiate identity. The paper first clarifies the technical foundations that make contemporary generative systems scalable, multimodal, and interactive. It then examines how those foundations support concrete applications across education, healthcare, business, creative industries, software development, and public services. Building on this, the article identifies the distinctive characteristics of generative AI in use, including probabilistic generation, conversational interfaces, personalization, co-creation, tool integration, and continuous adaptation. The central contribution is a structured discussion of impacts at two levels: individual and societal. At the individual level, generative AI affects learning, productivity, self-expression, well-being, and employability while also introducing risks of overreliance, cognitive offloading, privacy loss, and unequal access. At the societal level, the technology alters media systems, labor markets, institutional trust, governance capacity, and the distribution of opportunity. The article concludes that the future significance of generative AI will depend less on raw model scale alone than on governance design, human oversight, sector-specific implementation, and the ability of societies to convert generative capacity into trustworthy public value.

Keywords: Generative AI; individual impact; societal transformation; human-AI collaboration; future possibilities

Article History

Received November 20, 2025

Revised January 10, 2026

Accepted March 22, 2026

Available Online March 30, 2026

The Impact of Generative AI on Individual and Society: Prospects and Future Possibilities

1. Introduction

Generative artificial intelligence has moved from a specialized research frontier to a visible social infrastructure within only a few years. Systems based on large language models, diffusion models, and multimodal architectures are now used to draft text, write software, summarize legal material, generate educational resources, create images, support clinical documentation, and simulate strategic alternatives. What makes this technological wave historically significant is not only the quality of output, but also the way it changes the interface between human intention and digital execution. Earlier automation systems often required formal commands, explicit workflows, or fixed rules. Generative AI reduces those barriers by accepting natural language, partial intentions, and ambiguous goals, and then returning output that appears context-sensitive, stylistically adaptable, and increasingly action-oriented [Floridi & Chiriatti, 2020; Bommasani et al., 2021; Brown et al., 2020; Achiam et al., 2023; Dwivedi et al., 2023].

This article argues that generative AI should be understood as a socio-cognitive technology rather than merely as a new software tool. It does not simply automate narrow tasks; it redistributes access to expression, analysis, and coordination. A student can now obtain personalized explanations at any time. A professional can convert rough notes into presentations, code, proposals, or reports. A physician can use a system to draft patient-facing explanations or administrative summaries. A designer can iterate visually at a speed that previously required teams and longer cycles. A public institution can test how citizen-facing language, forms, and service responses might be adapted to different needs. In each case, generative AI changes the marginal cost of producing communicative and symbolic work, which is why its implications extend from personal productivity to institutional legitimacy and social trust [Raj et al., 2023; Noy & Zhang, 2023; Dell'Acqua et al., 2026; Ray, 2023; Stokel-Walker & Van Noorden, 2023].

The novelty of the current moment also lies in convergence. Recent systems combine scale, multimodality, instruction following, retrieval, tool use, and conversational memory into a unified experience. This makes generative AI more than a sequence model or image engine. It is becoming an interface layer through which people reach information, organize knowledge, and act in digital environments. Such convergence introduces both empowerment and fragility. It empowers because more users can now perform tasks once reserved for specialists. It is fragile because output quality remains probabilistic, context-sensitive, and vulnerable to hallucination, bias, manipulation, or overconfidence. The same properties that make these systems flexible also make them governance-intensive [Vaswani et al., 2017; Ouyang et al., 2022; Bender et al., 2021; Ji et al., 2023; Weidinger et al., 2022].

Unlike a conventional review article, the purpose of this paper is analytical and argumentative. The paper develops a structured thesis about the deep impact of generative AI on individuals and society, drawing on 100 published studies and technical reports with DOI-identified sources. The objective is not to summarize every strand of literature exhaustively, but to synthesize technical evidence, sector observations, and governance debates into a coherent account of why generative AI matters, how it operates in practice, and what future possibilities it opens. This orientation is especially important because public discourse often oscillates between utopian promises and catastrophic warnings. A more useful academic position is to explain the mechanisms through which value and risk are jointly produced. Figure 1 presents a concise technology timeline of Generative AI.




Figure 1. Technological Milestones in the Rise of Generative AI

2. Generative AI: Technical Background

The technical background of generative AI can be understood as the evolution from representation learning to controllable content generation (Figure 2). Earlier machine learning systems primarily mapped inputs to labels or predictions. Generative systems instead learn distributions from which new outputs may be sampled, reconstructed, transformed, or completed. Variational autoencoders and related probabilistic models provided one early route to latent representation learning, while generative adversarial networks demonstrated that competitive training could produce visually convincing synthetic outputs. These developments mattered not only because they improved data generation, but because they introduced a new computational logic: the machine could internalize structure and then externalize plausible novelty [Kingma & Welling, 2014; Rezende et al., 2014; Goodfellow et al., 2014; Gui et al., 2021].

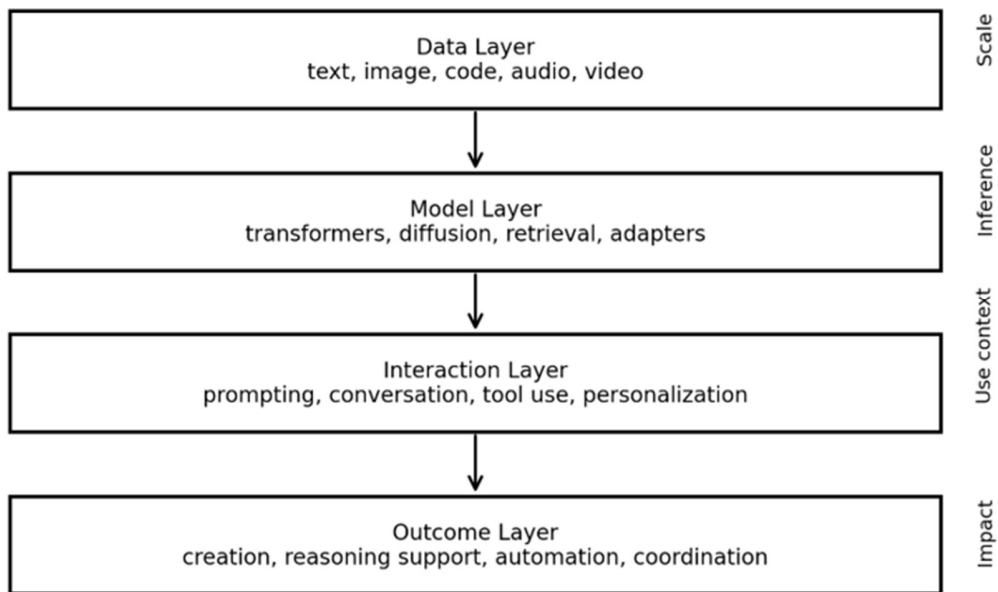


Figure 2. Socio-Technical Stack of Generative AI Systems

The second major shift came with sequence modeling and attention. Encoder-decoder systems first showed that learning could support translation and sequence transformation at scale, but the transformer architecture created a more general and scalable route to contextual representation. Attention mechanisms allowed models to process long dependencies efficiently and to generalize across tasks once trained on sufficiently large corpora. In this sense, transformers changed the economics of generality. One pretraining process could support downstream summarization, translation, question answering, coding, dialogue, and reasoning-style behavior. Models such as BERT, GPT-3, PaLM, LLaMA, and GPT-4 reflect this transition from task-specific systems to

foundation-model architectures whose behavior can be adapted through prompting, finetuning, feedback alignment, and tool augmentation [Bahdanau et al., 2014; Vaswani et al., 2017; Devlin et al., 2019; Brown et al., 2020; Chowdhery et al., 2022; Touvron et al., 2023a; Achiam et al., 2023].

The third shift involved alignment and interaction. Raw scale improved fluency but did not guarantee usefulness or safety. Instruction tuning, reinforcement learning from human feedback, constitutional design, retrieval augmentation, and prompting techniques made generative AI more responsive to user goals. These methods reduced the gap between a model that can statistically continue text and a system that can perform social or organizationally meaningful work. The importance of this change cannot be overstated. When systems became instruction-following, people without technical expertise could meaningfully engage with them. A large part of generative AI's social diffusion therefore comes from interface design and training alignment, not from parameter count alone [Ouyang et al., 2022; Wei et al., 2021; Chung et al., 2022; Bai et al., 2022; Lewis et al., 2020; White et al., 2023].

The fourth shift is multimodality. Contemporary systems increasingly handle combinations of text, image, audio, code, and video, while also calling external tools or APIs. Diffusion models drastically improved image quality and controllability; CLIP-like alignment connected language and vision; prompt conditioning and control networks increased editability; and agentic frameworks linked models to browsers, calculators, databases, and software environments. In practice, multimodality means that generative AI is moving from content generation toward environment orchestration. Users do not simply ask for a paragraph or picture; they ask for a plan, a simulation, a draft, a workflow, or a coordinated action [Ho et al., 2020; Nichol & Dhariwal, 2021; Dhariwal & Nichol, 2021; Rombach et al., 2022; Ramesh et al., 2022; Schick et al., 2023; Wu et al., 2023].

A final technical point concerns limits. These systems remain probabilistic approximators whose apparent competence can mask hidden brittleness (Table 1). Hallucinations, prompt sensitivity, data contamination, hidden bias, environmental cost, and safety vulnerabilities remain fundamental rather than accidental problems. Moreover, emergent capability is not equivalent to robust understanding. It often reflects scale, pattern interpolation, and interface scaffolding. The technological power of generative AI is real, but its reliability remains uneven across domains. That is why deployment context matters so much: the same model can be benign in brainstorming and harmful in diagnosis, harmless in note drafting and risky in legal advice [Floridi & Chiriatti, 2020; Bender et al., 2021; Perez et al., 2022; Ji et al., 2023; Weidinger et al., 2022; Ganguli et al., 2023].

Table 1. Major technological milestones and their societal meaning

Period	Technical milestone	Representative models	Implication for society
2014-2016	Latent and adversarial generation	VAE; GAN; context encoders	Machine generation becomes visibly creative
2017-2019	Attention and foundation pretraining	Transformer; BERT	Language becomes a general representation space
2020	Large-scale in-context generation	GPT-3	Natural language turns into an action interface
2021-2022	Alignment and retrieval	RLHF; RAG; tool use	Systems become more useful to non-experts
2022	Diffusion breakthrough	DDPM; latent diffusion; DALL-E 2	High-quality image generation scales

2023	Open foundation model diffusion	LLaMA; instruction chat models	Access broadens beyond a few firms
2023-2024	Agentic and multimodal integration	GPT-4; Gemini; AutoGen	Generative AI moves toward workflow orchestration

3. Generative AI in Concrete Applications

Generative AI matters because it is no longer confined to a single vertical domain. In education, it functions as a tutoring partner, assessment assistant, language practice partner, accessibility enhancer, and content generator (Figure 3). Educators use it to build lecture notes, examples, rubrics, and adaptive materials, while students use it for explanation, brainstorming, summarization, practice, and revision (Table 2). The educational significance lies in personalization at scale. A learner who previously depended on scarce office hours now has persistent interactive support. Yet the same affordance also changes assessment, authorship, and the meaning of effort. If systems can instantly draft essays, answer conceptual questions, and solve programming exercises, institutions must redesign evaluation around process transparency, critical reflection, and human judgment rather than output alone [Kasneji et al., 2023; Lim et al., 2023; Chiu, 2024; Tlili et al., 2023; Baidoo-Anu & Ansah, 2023; Pang et al., 2025; Qu et al., 2025].

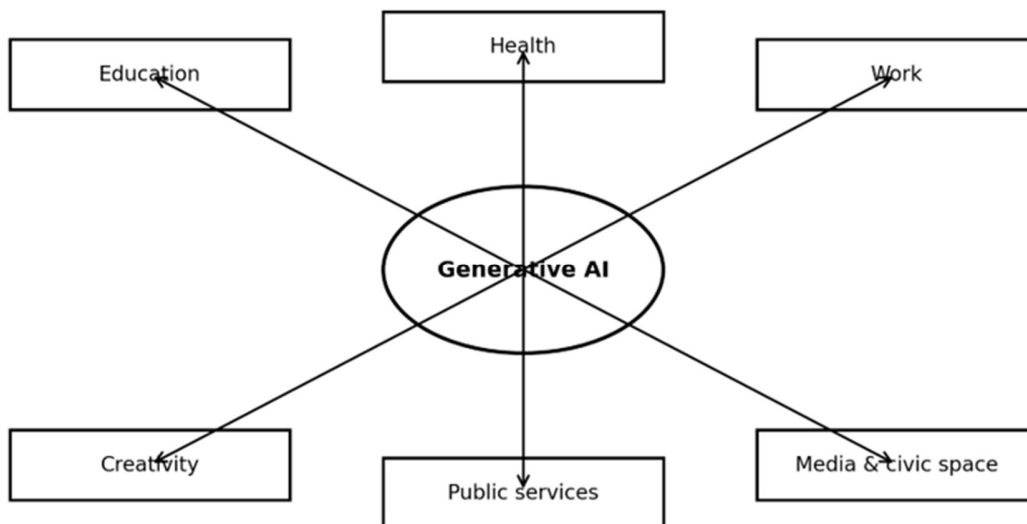


Figure 3. Application Landscape of Generative AI across Major Domains

In healthcare, generative AI has become visible first in administrative and communicative tasks rather than in autonomous clinical decision making. Typical uses include drafting discharge notes, translating patient instructions into plain language, triaging knowledge requests, supporting literature review, and assisting clinical documentation. Studies suggest that language models can perform surprisingly well on standardized medical examinations or patient-facing communication tasks, but this does not imply that they should replace clinical judgment. The more defensible path is augmentation under oversight: reduce clerical burden, improve explanation quality, expand patient access, and strengthen evidence retrieval while preserving professional accountability. As implementation studies make clear, translation into care settings depends on governance, evaluation, and workflow fit, not just model performance [Gilson et al., 2023; Kung et al., 2023; Ayers et al., 2023; Lee et al., 2023; Sallam, 2023; Reddy et al., 2024; Templin et al., 2024; Yim et al., 2024; Abbasian et al., 2024].

In business and knowledge work, the productivity effects of generative AI are already measurable. Experimental research suggests that systems such as ChatGPT can reduce completion time and raise average quality for many writing-intensive or analysis-oriented tasks, especially for less experienced workers. Firms are therefore deploying AI copilots for drafting, coding, customer support, report generation, internal search, marketing variation, and data interpretation. However, the productivity narrative is incomplete unless paired with organizational nuance. Generative AI does not improve all work equally. It often raises average performance while compressing differences, helps workers within the model's competence frontier more than outside it, and can degrade outcomes when users overtrust confident but incorrect outputs. Thus, firms face a redesign question rather than a plug-and-play adoption question: which tasks should be accelerated, which should remain human-led, and how should responsibility be assigned when human and machine jointly produce results? [Noy & Zhang, 2023; Raj et al., 2023; Dell'Acqua et al., 2026; Eloundou et al., 2023; Yu et al., 2024].

Creative industries represent another central field of application. Text-to-image and multimodal generation have changed prototyping, ideation, concept art, advertising, product visualization, and social media production. The principal advantage is not that machines suddenly replaced creativity, but that the number of iterations possible within a fixed time dramatically increased. Generative AI lowers the cost of exploring styles, narratives, and formats. It turns the blank page into a dialogue and accelerates transition from concept to artifact. Yet creative labor also experiences identity pressure. If value once depended partly on exclusive command of technique, then systems that democratize output alter the scarcity structure of cultural work. This may widen access to creative participation while simultaneously intensifying competition for attention and authenticity [Ramesh et al., 2021; Rombach et al., 2022; Saharia et al., 2022; Brooks et al., 2023; Xu et al., 2024].

Public services and civic communication form a fifth application cluster. Governments can use generative systems to simplify administrative forms, offer multilingual assistance, draft policy options, summarize consultations, and personalize citizen-facing communication. These functions are attractive because state capacity is often constrained by paperwork, fragmentation, and linguistic barriers. At the same time, democratic legitimacy imposes a higher threshold than ordinary consumer applications. Citizens must know when they are dealing with an AI system, public records must remain auditable, and decisions that affect rights cannot be reduced to opaque probabilistic generation. The future of generative AI in government therefore depends on a balance between service efficiency and procedural justice [Li et al., 2024; Sharma & Yadav, 2024; van Dis et al., 2023].

Table 2. Representative Application Domains of Generative AI

Domain	Representative uses	Primary value logic	Key caution
Education	Tutoring, feedback, content drafting	Personalization and access	Academic integrity and dependence
Healthcare	Documentation, explanation, search	Clerical relief and communication	Safety and accountability
Business	Writing, coding, support, analysis	Productivity and speed	Overreliance and process opacity
Creative industries	Image/video generation, ideation	Iteration and democratization	Authorship and market saturation
Science	Hypothesis support, literature search, simulation assistance	Cognitive acceleration	Verification burden

Public services	Citizen assistance, translation, summarization	Service inclusion and capacity	Transparency and due process
-----------------	--	--------------------------------	------------------------------

4. Characteristics and Functions of Generative AI in Use

The power of generative AI in practice comes from a set of interacting characteristics rather than from a single technical feature. The first is probabilistic generation. These systems do not retrieve fixed answers in the way a database does; they produce context-conditioned responses by predicting plausible continuations or outputs. This gives them flexibility and stylistic range, but also means output quality must always be interpreted as probabilistic rather than guaranteed. In practical terms, generative AI is valuable when plausibility, variation, and rapid draft generation are useful, and less appropriate when determinism and exactness are mandatory [Kingma & Welling, 2014; Goodfellow et al., 2014; Ho et al., 2020; Song et al., 2021].

The second characteristic is natural-language mediation. Users can engage with complex systems through ordinary language instead of specialized code or rigid interfaces. This lowers access barriers and redistributes technical leverage. A person who cannot program may still create a workflow, generate a data explanation, or draft a formal document. From a social perspective, natural-language interaction is one reason generative AI spreads so rapidly: the learning curve is shallow enough to enable mass experimentation. Yet natural language is also ambiguous, meaning that prompting quality, hidden assumptions, and conversational framing strongly shape outcomes [Brown et al., 2020; Ouyang et al., 2022; White et al., 2023; Qin et al., 2023].

The third characteristic is iterative co-creation. Unlike a static search result, a generative system can revise, extend, simplify, translate, or stylistically transform output across multiple conversational turns. This creates a dynamic relationship in which the user steers direction while the model expands possibility space. In education and design, this iterative quality is especially significant because it supports experimentation and reflection. However, it may also encourage surface-level satisfaction with persuasive outputs before deeper verification occurs [Kasneci et al., 2023; Mollick & Mollick, 2023; Brooks et al., 2023; Xu et al., 2024].

The fourth characteristic is personalization. Systems can adapt tone, examples, reading level, task structure, or feedback style to different user needs. This makes generative AI unusually relevant for inclusion: non-native speakers, novice learners, and users with communication barriers can receive more accessible support. But personalization also raises questions about profiling, privacy, and dependence. A technology that seems empowering at the individual level may quietly intensify surveillance or behavioral nudging at scale [Baidoo-Anu & Ansah, 2023; Qu et al., 2025; Pang et al., 2025; Wang et al., 2024].

The fifth characteristic is composability. With retrieval, tool use, agent frameworks, and APIs, generative models increasingly operate as components within larger systems. They can browse, calculate, retrieve records, call external tools, or coordinate multi-step processes. This shifts their function from generation alone toward orchestration. As a result, generative AI now influences not only what content is produced, but how workflows are sequenced, delegated, and monitored. Figure 4 captures these pathways at the level of the individual, while Table 3 summarizes the core functions that recur across sectors [Lewis et al., 2020; Schick et al., 2023; Yao et al., 2022; Yao et al., 2023; Wu et al., 2023].

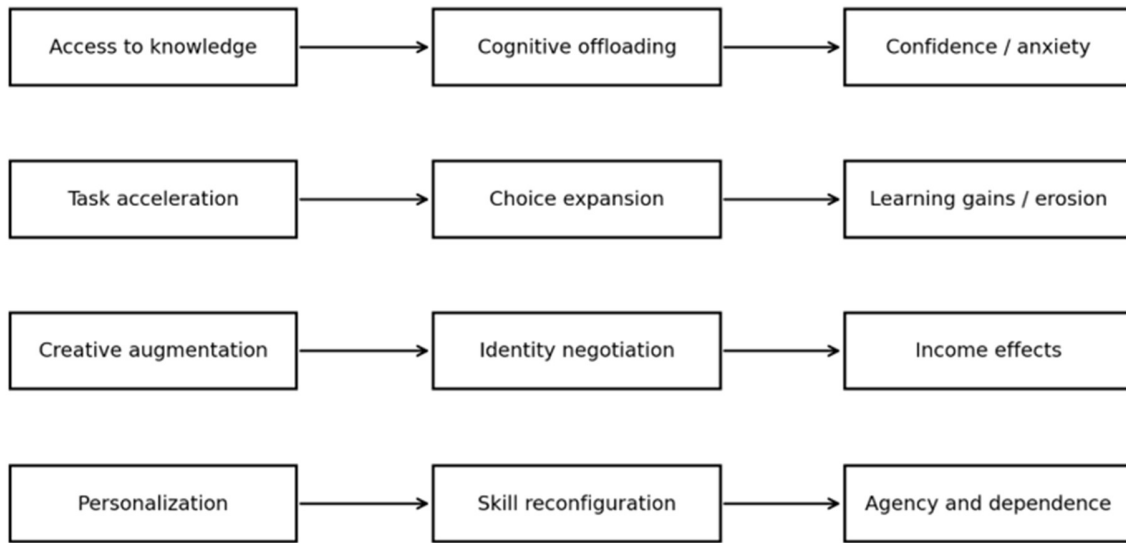


Figure 4. Individual-Level Pathways through Which Generative AI Affects Human Activity

Table 3. Core Characteristics and Functions of Generative AI in Application

Characteristic	Functional expression	Practical advantage	Typical risk
Probabilistic generation	Drafting, completion, variation	Speed and flexibility	Hallucination
Natural-language interface	Prompting and dialogue	Low access barrier	Ambiguity
Iterative co-creation	Revision, editing, transformation	Exploration and refinement	Superficial acceptance
Personalization	Tone, level, context adaptation	Inclusion and tailored support	Dependence and profiling
Composability	Tools, retrieval, agents	Workflow integration	Opaque automation

5. Deep Impact on the Individual

At the individual level, generative AI changes the relationship between effort and capability (Table 4). It functions as a cognitive amplifier in tasks involving drafting, explanation, summarization, search, coding, and ideation. For many users, this means lower entry barriers to complex work. A novice can produce a passable memo, a beginner programmer can generate a working template, and a learner can obtain multiple explanations of the same concept until comprehension improves. This has important egalitarian potential because it reduces dependence on scarce human guidance. People with fewer prior resources may gain access to forms of assistance once limited by cost, time, geography, or institutional gatekeeping [Noy & Zhang, 2023; Raj et al., 2023; Kasneci et al., 2023; Baidoo-Anu & Ansah, 2023; Pang et al., 2025].

Yet capability amplification is inseparable from cognitive offloading. When a system drafts, remembers, structures, and reformulates on the user's behalf, some forms of mental labor shift from internal processing to external delegation. In moderation this can be beneficial: it frees attention for strategy, judgment, and synthesis. In excess it may erode learning, memory consolidation, writing fluency, and problem-solving persistence. The issue is not whether cognitive offloading exists—humans have always offloaded onto paper, calculators, or search engines—but whether generative AI's fluidity makes offloading so frictionless that users stop engaging

deeply with the underlying task. This concern is especially salient for education and early-career skill formation, where competence is built partly through struggle and repetition [Bastani et al., 2025; Gill et al., 2024; Lim et al., 2023; Zhai, 2022].

Generative AI also affects identity and self-expression. Because the system can mirror tone, imitate genres, and suggest phrasing, it becomes intertwined with how people present themselves in academic, professional, and social settings. For some users this is empowering. Non-native speakers may communicate more confidently; introverted users may structure thoughts more clearly; creators may explore styles that feel newly reachable. For others, the presence of algorithmic mediation raises authenticity concerns. If personal expression is continually co-produced by a machine, where should authorship be located? The answer is unlikely to be binary. In practice, many outputs will be hybrid, and personal agency will depend on how visibly and critically the person shapes the final artifact [Taecharunroj, 2023; Xu et al., 2024; Mollick & Mollick, 2023].

The technology further changes employability and skill composition. Routine writing, translation, coding, and formatting tasks are increasingly augmented, which means the value of purely procedural knowledge may decline in some occupations. At the same time, demand may rise for problem framing, domain judgment, verification, cross-modal integration, emotional intelligence, and governance literacy. Individuals who learn to supervise, refine, and contextualize generative output may gain an advantage over those who either reject the tools entirely or trust them uncritically. In other words, generative AI is likely to reward meta-skill development: asking good questions, diagnosing weak output, integrating evidence, and understanding when not to automate [Eloundou et al., 2023; Dell'Acqua et al., 2026; Yu et al., 2024].

Well-being is another emerging dimension. Some users report reduced anxiety when the technology helps them organize information, overcome blank-page paralysis, or access support. Others experience new pressure: fear of being replaced, difficulty proving authenticity, or stress from keeping up with rapidly shifting expectations. When generative AI enters daily life as a near-constant assistant, it can create subtle dependence. People may become less willing to begin unaided, less patient with manual work, or more uncertain about their own unaided capability. This is why the impact of generative AI on the individual cannot be measured only through task efficiency. It must also be evaluated through confidence, autonomy, trust, and psychological resilience [Wang et al., 2024; De Angelis et al., 2023; Rudolph et al., 2023].

Table 4. Main Ways in Which Generative AI Affects the Individual

Impact channel	Positive possibility	Countervailing risk	Strategic implication
Learning	On-demand explanation and practice	Shallow understanding	Assess process, not only output
Productivity	Faster drafting and coding	Overtrust in incorrect outputs	Keep verification visible
Expression	Confidence and stylistic support	Authenticity concerns	Disclose meaningful AI mediation
Employment	Skill augmentation and role redesign	Task displacement	Invest in meta-skills
Well-being	Reduced friction and support	Dependence and anxiety	Promote healthy boundaries

6. Deep Impact on Society

The societal impact of generative AI begins with information abundance (Table 5). When text, images, video, code, and synthetic voices can be produced at low marginal cost, the informational environment changes qualitatively. Content scarcity gives way to filtering scarcity. This has profound implications for media systems, political communication, and public trust. Societies must increasingly distinguish between what can be generated and what deserves credibility. Search, recommendation, and platform governance therefore become more important, not less, in the age of generation. A future saturated with synthetic content is not simply a future of more communication; it is a future in which verification, provenance, and institutional trust are strategic resources [Dwivedi et al., 2023; De Angelis et al., 2023; Zhang et al., 2024; van Dis et al., 2023].

Table 5. Societal Opportunities and Risks Associated with Generative AI

Societal arena	Transformative promise	Systemic risk	Governance need
Media sphere	More voices and formats	Synthetic misinformation	Provenance and moderation
Labor market	Productivity and new roles	Displacement and inequality	Reskilling and fair diffusion
Institutions	Better services and lower friction	Opacity and legitimacy loss	Auditability and disclosure
Science and innovation	Faster search and hypothesis support	Quality dilution and reproducibility burdens	Verification standards
Civic life	Accessibility and multilingual inclusion	Trust erosion and manipulation	Public-interest governance

Labor markets represent a second societal transformation pathway. Generative AI does not automate only manual repetition; it increasingly affects symbolic and knowledge-intensive work. This does not mean mass replacement is inevitable, but it does mean occupational boundaries will be redefined. Some tasks will be automated, many will be reorganized, and new supervisory or integrative roles will appear. The broader social question is distributional: who captures the gains in productivity, and who bears the transitional costs? If generative AI raises output while concentrating value in a small number of firms, platforms, or high-skill professionals, inequality may widen. If, however, its benefits are diffused through education, public tools, and organizational redesign, it could broaden participation in high-value work [Eloundou et al., 2023; Noy & Zhang, 2023; Dell'Acqua et al., 2026].

A third societal effect concerns institutions. Schools, hospitals, firms, courts, media organizations, and governments increasingly confront a shared dilemma: whether to treat generative AI as a prohibited shortcut, a managed assistant, or a structural redesign catalyst. Institutions that focus only on prohibition may lose relevance, while institutions that embrace generative AI without procedural safeguards may lose legitimacy. The likely long-run equilibrium is institutional adaptation. Assessment systems will shift toward reasoning traces and oral defense; professional practice will emphasize accountable review; public organizations will add disclosure and audit mechanisms; and firms will formalize rules about acceptable use, responsibility, and data protection [Kasneci et al., 2023; Reddy et al., 2024; Li et al., 2024; Sharma & Yadav, 2024].

Culture and social stratification constitute a fourth pathway. Generative AI may democratize creation, but it can also deepen attention inequality. When content production is cheap, discoverability becomes more valuable. Platform dynamics may favor already visible actors, while smaller creators compete within a more crowded symbolic marketplace. Furthermore, access to high-quality models, premium subscriptions, enterprise deployments, and AI literacy training is unevenly distributed across countries, organizations, and social groups.

The result may be a paradoxical society: more people can create than ever before, yet the capacity to convert creation into recognition, income, or influence may remain highly unequal. Figure 5 illustrates these linked societal channels.

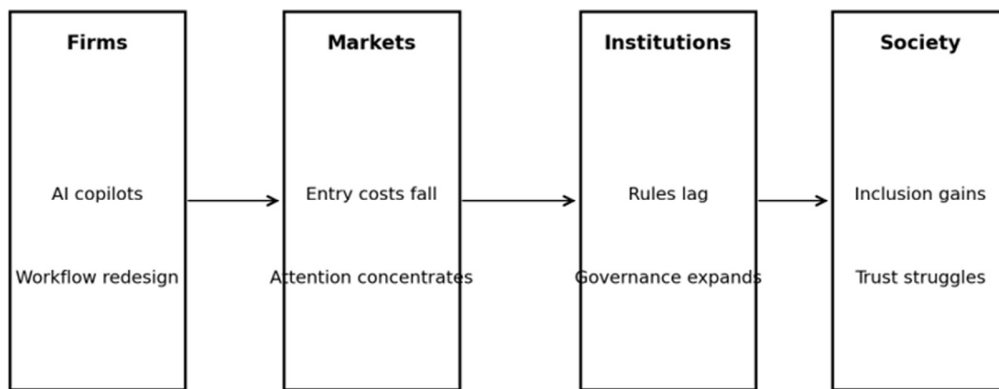


Figure 5. Pathways from firm-level adoption to broader social transformation

7. Future Possibilities and Governance Directions

The future of generative AI will likely be shaped by three interacting trends: deeper integration into everyday environments, greater model autonomy through agentic design, and stronger governance demands. Integration means that generative functions will disappear into software infrastructure rather than remain separate destination products. Writing tools, office suites, search systems, educational platforms, scientific instruments, and public-service portals will all absorb generative capabilities. As that occurs, the meaningful question will shift from whether one 'uses AI' to how much of the environment is AI-mediated by default [Gemini Team, 2023; OpenAI, 2023; Wu et al., 2023].

Agentic development introduces new possibilities and new risks. When a model can plan, retrieve, call tools, coordinate with other agents, and revise its own output, it begins to approximate delegated workflow execution. This raises efficiency potential in customer service, software engineering, scientific assistance, and administrative processing. Yet it also magnifies concerns over traceability, alignment drift, and error propagation. The more steps an AI system takes on a user's behalf, the more important it becomes to design visible checkpoints, override mechanisms, and bounded authority. High-autonomy systems should not be evaluated by the same criteria as simple drafting assistants [Yao et al., 2022; Yao et al., 2023; Shinn et al., 2023; Madaan et al., 2023].

A productive future therefore depends on governance by design. Governance should begin with application context rather than abstract fear or generic enthusiasm. Low-risk contexts such as brainstorming, translation support, or first-draft generation can tolerate higher experimentation. High-stakes contexts such as healthcare, legal advice, credit, policing, and public entitlements require more conservative deployment, domain evaluation, logging, and human accountability. This differentiated view avoids the mistake of treating all generative AI use as equivalent. It also allows innovation to continue in settings where value is clear and harms are manageable [Reddy et al., 2024; Templin et al., 2024; Thomas et al., 2024].

The broader possibility is that generative AI may become a foundational complement to human intelligence rather than a substitute for it. In science, it may accelerate hypothesis generation, simulation, and literature navigation; in education, it may scale personalized support; in work, it may reduce drudgery and strengthen high-level synthesis; in government, it may improve accessibility and responsiveness. Whether these possibilities materialize depends on institutional choices about openness, training, competition, infrastructure,

and standards. A society that treats generative AI only as a source of private productivity may underinvest in public-interest uses. A society that treats it only as a source of danger may miss significant gains in inclusion and capability. Figure 6 positions future scenarios along the axes of autonomy and oversight.

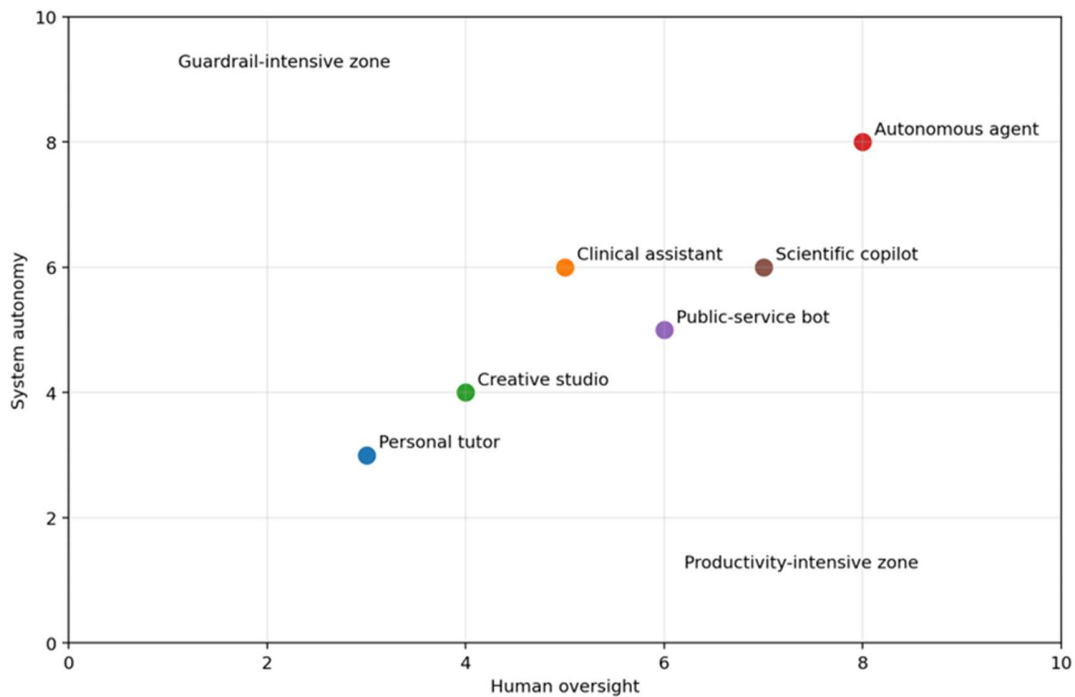


Figure 6. Future Deployment Scenarios: Autonomy versus Human Oversight

8. Conclusion

Generative AI is already reshaping the conditions under which people learn, work, communicate, create, and govern. Its significance lies not simply in better text generation or more realistic images, but in the emergence of a new interface between human intention and machine execution. As natural-language mediation, multimodality, retrieval, and tool use converge, generative AI becomes a general-purpose layer of cognitive infrastructure. It lowers barriers to expression and analysis, expands access to assistance, and accelerates many forms of symbolic labor. At the same time, it introduces new risks related to hallucination, bias, manipulation, overreliance, opacity, and unequal access.

The article has argued that the impact of generative AI should be interpreted across two intertwined levels. At the individual level, the technology changes capability, confidence, learning patterns, self-expression, and employability. At the societal level, it alters information systems, labor markets, institutions, and cultural competition. The central challenge is therefore not whether generative AI is inherently good or bad. The challenge is how to build conditions under which its generative capacity creates trustworthy public value. This requires application-specific governance, visible human oversight, verification cultures, inclusive access, and institutional redesign rather than passive adoption.

Future possibilities remain substantial. Generative AI could widen access to expertise, improve public communication, augment scientific discovery, and free human attention for more relational, strategic, and creative work. But these outcomes will not occur automatically. They depend on who controls models and infrastructure, how gains are distributed, what standards shape deployment, and whether societies cultivate the human capabilities needed to work critically with generative systems. The long-run question is not whether

generative AI will influence individuals and society. It already does. The real question is whether that influence will deepen inequality and distrust, or whether it will be governed in ways that expand human capability, institutional quality, and social resilience.

ACKNOWLEDGEMENT

The authors acknowledge the intellectual contributions of the broader interdisciplinary research community whose work on generative models, human-AI interaction, education, healthcare, governance, and digital transformation informed the analytical development of this article.

Reference

- Abbasian, M., Azamfirei, R., Hauser, P., Shen, Y., Sönksen, P., & Althoff, T. (2024). Foundation metrics for evaluating effectiveness of healthcare conversational agents. *npj Digital Medicine*, 7, 258. DOI: 10.1038/s41746-024-01074-z.
- Achiam, J., Adler, S., Agarwal, S., et al. (2023). GPT-4 technical report. arXiv. DOI: 10.48550/arXiv.2303.08774.
- Ayers, J. W., Poliak, A., Dredze, M., et al. (2023). Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA Internal Medicine*, 183(6), 589-596. DOI: 10.1001/jamainternmed.2023.1838.
- Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. arXiv. DOI: 10.48550/arXiv.1409.0473.
- Bai, Y., Kadavath, S., Kundu, S., et al. (2022). Constitutional AI: Harmlessness from AI feedback. arXiv. DOI: 10.48550/arXiv.2212.08073.
- Baidoo-Anu, D., & Ansah, L. O. (2023). Education in the era of generative artificial intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning. *Journal of AI*, 7, 52-62. DOI: 10.61969/jai.1337500.
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakçı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning. *Proceedings of the National Academy of Sciences*, 122(13), e2422633122. DOI: 10.1073/pnas.2422633122.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *FAccT '21* (pp. 610-623). DOI: 10.1145/3442188.3445922.
- Bommasani, R., Hudson, D. A., Adeli, E., et al. (2021). On the opportunities and risks of foundation models. arXiv. DOI: 10.48550/arXiv.2108.07258.
- Brooks, T., Holynski, A., & Efros, A. A. (2023). InstructPix2Pix: Learning to follow image editing instructions. In *CVPR* (pp. 18392-18402). DOI: 10.1109/CVPR52729.2023.01764.
- Brown, T., Mann, B., Ryder, N., et al. (2020). Language models are few-shot learners. arXiv. DOI: 10.48550/arXiv.2005.14165.
- Bubeck, S., Chandrasekaran, V., Eldan, R., et al. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4. arXiv. DOI: 10.48550/arXiv.2303.12712.

- Chiu, T. K. F. (2024). The impact of Generative AI (GenAI) on practices, policies and research direction in education: A case of ChatGPT and Midjourney. *Interactive Learning Environments*, 32(6), 2375-2391. DOI: 10.1080/10494820.2023.2253861.
- Chowdhery, A., Narang, S., Devlin, J., et al. (2022). PaLM: Scaling language modeling with pathways. arXiv. DOI: 10.48550/arXiv.2204.02311.
- Chung, H. W., Hou, L., Longpre, S., et al. (2022). Scaling instruction-finetuned language models. arXiv. DOI: 10.48550/arXiv.2210.11416.
- De Angelis, L., Baglivo, F., Arzilli, G., et al. (2023). ChatGPT and the rise of large language models: The new AI-driven infodemic threat in public health. *Frontiers in Public Health*, 11, 1166120. DOI: 10.3389/fpubh.2023.1166120.
- Dell'Acqua, F., McFowland, E., Mollick, E., et al. (2026). Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organization Science*, 37(2), 403-423. DOI: 10.1287/orsc.2025.21838.
- Dettmers, T., Pagnoni, A., Holtzman, A., & Zettlemoyer, L. (2023). QLoRA: Efficient finetuning of quantized LLMs. arXiv. DOI: 10.48550/arXiv.2305.14314.
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. arXiv. DOI: 10.48550/arXiv.1810.04805.
- Dhariwal, P., & Nichol, A. (2021). Diffusion models beat GANs on image synthesis. arXiv. DOI: 10.48550/arXiv.2105.05233.
- Dwivedi, Y. K., Kshetri, N., Hughes, L., et al. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, 102642. DOI: 10.1016/j.ijinfomgt.2023.102642.
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. arXiv. DOI: 10.48550/arXiv.2303.10130.
- Esser, P., Rombach, R., Blattmann, A., & Ommer, B. (2021). Taming transformers for high-resolution image synthesis. In *CVPR* (pp. 12873-12883). DOI: 10.1109/CVPR46437.2021.01269.
- Floridi, L., & Chiriatti, M. (2020). GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30, 681-694. DOI: 10.1007/s11023-020-09548-1.
- Ganguli, D., Madaan, A., Anthropic, et al. (2023). Predictability and surprise in large generative models. arXiv. DOI: 10.48550/arXiv.2307.03507.
- Gao, L., Biderman, S., Black, S., et al. (2020). The pile: An 800GB dataset of diverse text for language modeling. arXiv. DOI: 10.48550/arXiv.2101.00027.
- Gemini Team. (2023). Gemini: A family of highly capable multimodal models. arXiv. DOI: 10.48550/arXiv.2312.11805.
- Gill, S. S., Xu, M., Patros, P., et al. (2024). Transformative effects of ChatGPT on modern education: Emerging era of AI chatbots. *Internet of Things and Cyber-Physical Systems*, 4, 19-23. DOI: 10.1016/j.iotcps.2023.06.002.
- Gilson, A., Safranek, C. W., Huang, T., et al. (2023). How does ChatGPT perform on the United States Medical Licensing Examination? The implications of large language models for medical education and knowledge assessment. *npj Digital Medicine*, 6, 85. DOI: 10.1038/s41746-023-00881-0.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al. (2014). Generative adversarial networks. *Communications of the ACM*, 63(11), 139-144. DOI: 10.1145/3422622.
- Gui, J., Sun, Z., Wen, Y., Tao, D., & Ye, J. (2021). A review on generative adversarial networks: Algorithms, theory, and applications. *IEEE Transactions on Knowledge and Data Engineering*, 35(4), 3313-3332. DOI: 10.1109/TKDE.2021.3130191.

- Hertz, A., Mokady, R., Tenenbaum, J., et al. (2022). Prompt-to-prompt image editing with cross attention control. arXiv. DOI: 10.48550/arXiv.2208.01626.
- Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. arXiv. DOI: 10.48550/arXiv.2006.11239.
- Hoffmann, J., Borgeaud, S., Mensch, A., et al. (2022). Training compute-optimal large language models. arXiv. DOI: 10.48550/arXiv.2203.15556.
- Hu, E. J., Shen, Y., Wallis, P., et al. (2021). LoRA: Low-rank adaptation of large language models. arXiv. DOI: 10.48550/arXiv.2106.09685.
- Ji, Z., Lee, N., Frieske, R., et al. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), 248. DOI: 10.1145/3571730.
- Jiao, W., Wang, W., Huang, J.-T., Wang, X., & Tu, Z. (2023). Is ChatGPT a good translator? A preliminary study. arXiv. DOI: 10.48550/arXiv.2301.08745.
- Kaddour, J., Harris, J., Mozes, M., et al. (2023). Challenges and applications of large language models. arXiv. DOI: 10.48550/arXiv.2307.10169.
- Kasneci, E., Sessler, K., Küchemann, S., et al. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. DOI: 10.1016/j.lindif.2023.102274.
- Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. arXiv. DOI: 10.48550/arXiv.1312.6114.
- Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., & Iwasawa, Y. (2022). Large language models are zero-shot reasoners. arXiv. DOI: 10.48550/arXiv.2205.11916.
- Kung, T. H., Cheatham, M., Medinilla, A., et al. (2023). Performance of ChatGPT on USMLE: Potential for AI-assisted medical education using large language models. *PLOS Digital Health*, 2(2), e0000198. DOI: 10.1371/journal.pdig.0000198.
- Lee, P., Bubeck, S., & Petro, J. (2023). Benefits, limits, and risks of GPT-4 as an AI chatbot for medicine. *New England Journal of Medicine*, 388, 1233-1239. DOI: 10.1056/NEJMs2214184.
- Lewis, P., Perez, E., Piktus, A., et al. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. arXiv. DOI: 10.48550/arXiv.2005.11401.
- Li, Y., Li, S., & Cipit, A. (2024). Large language models in public administration: Opportunities and governance dilemmas. *Government Information Quarterly*, 41(4), 101973. DOI: 10.1016/j.giq.2024.101973.
- Lim, W. M., Gunasekara, A. N., Pallant, J. L., Pallant, J. I., & Pechenkina, E. (2023). Generative AI and the future of education: Ragnarök or reformation? A paradoxical perspective from management educators. *International Journal of Management Education*, 21(2), 100790. DOI: 10.1016/j.ijme.2023.100790.
- Liu, Y., Han, T., Ma, S., et al. (2023). Summary of ChatGPT/GPT-4 research and perspective towards the future of large language models. arXiv. DOI: 10.48550/arXiv.2304.01852.
- Madaan, A., Tandon, N., Clark, P., & Yang, Y. (2023). Self-refine: Iterative refinement with self-feedback. arXiv. DOI: 10.48550/arXiv.2303.17651.
- Mialon, G., Dessì, R., Lombardi, D., et al. (2023). Augmented language models: A survey. arXiv. DOI: 10.48550/arXiv.2302.07842.
- Mollick, E., & Mollick, L. (2023). Assigning AI: Seven approaches for students, with prompts. *SSRN Electronic Journal*. DOI: 10.2139/ssrn.4475995.
- Nakano, R., Hilton, J., Balaji, S., et al. (2021). WebGPT: Browser-assisted question-answering with human feedback. arXiv. DOI: 10.48550/arXiv.2112.09332.
- Nazir, A., Rao, N., Wu, L., & Sun, L. (2023). A comprehensive survey of ChatGPT: Advancements, applications, prospects, and challenges. *Meta-Radiology*, 1(3), 100022. DOI: 10.1016/j.metrad.2023.100022.

- Nichol, A. Q., & Dhariwal, P. (2021). Improved denoising diffusion probabilistic models. arXiv. DOI: 10.48550/arXiv.2102.09672.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187-192. DOI: 10.1126/science.adh2586.
- OpenAI. (2023). GPT-4 technical report. arXiv. DOI: 10.48550/arXiv.2303.08774.
- Ouyang, L., Wu, J., Jiang, X., et al. (2022). Training language models to follow instructions with human feedback. arXiv. DOI: 10.48550/arXiv.2203.02155.
- Pang, W., Wang, Y., Wang, S., & Liu, H. (2025). A technology usage study on generative AI innovations in higher education. *Information*, 16(2), 95. DOI: 10.3390/info16020095.
- Park, J. S., O'Brien, J., Cai, C. J., et al. (2023). Generative agents: Interactive simulacra of human behavior. In CHI '23. DOI: 10.1145/3544548.3581218.
- Pathak, D., Krähenbühl, P., Donahue, J., Darrell, T., & Efros, A. A. (2016). Context encoders: Feature learning by inpainting. In CVPR (pp. 2536-2544). DOI: 10.1109/CVPR.2016.278.
- Perez, E., Kiela, D., Cho, K., Khashabi, D., & Weston, J. (2022). Red teaming language models with language models. arXiv. DOI: 10.48550/arXiv.2202.03286.
- Qu, Y., Wu, H., & Huang, R. (2025). Generative artificial intelligence in higher education: Disclosure, dependence and academic integrity. *British Journal of Educational Technology*, 56(1), 1-18. DOI: 10.1111/bjjet.70029.
- Raj, R., Shrivastava, G., Dhar, P., & Singh, A. (2023). Analyzing the potential benefits and use cases of ChatGPT as a tool for improving productivity and efficiency in business operations. *Technology in Society*, 75, 102320. DOI: 10.1016/j.techsoc.2023.102320.
- Ramesh, A., Dhariwal, P., Nichol, A., et al. (2022). Hierarchical text-conditional image generation with CLIP latents. arXiv. DOI: 10.48550/arXiv.2204.06125.
- Ramesh, A., Pavlov, M., Goh, G., et al. (2021). Zero-shot text-to-image generation. arXiv. DOI: 10.48550/arXiv.2102.12092.
- Ray, P. P. (2023). ChatGPT: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope. *Internet of Things and Cyber-Physical Systems*, 3, 121-154. DOI: 10.1016/j.ioteps.2023.04.003.
- Reddy, S., Allan, S., Coghlan, S., & Cooper, P. (2024). Generative AI in healthcare: An implementation science informed translational path on application, integration and governance. *Implementation Science*, 19, 32. DOI: 10.1186/s13012-024-01357-9.
- Reed, S., Akata, Z., Yan, X., et al. (2016). Generative adversarial text to image synthesis. In ICML (pp. 1060-1069). DOI: 10.48550/arXiv.1605.05396.
- Rezende, D. J., Mohamed, S., & Wierstra, D. (2014). Stochastic backpropagation and approximate inference in deep generative models. arXiv. DOI: 10.48550/arXiv.1401.4082.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In CVPR (pp. 10684-10695). DOI: 10.1109/CVPR52688.2022.01042.
- Saharia, C., Chan, W., Saxena, S., et al. (2022). Photorealistic text-to-image diffusion models with deep language understanding. arXiv. DOI: 10.48550/arXiv.2205.11487.
- Sallam, M. (2023). ChatGPT utility in healthcare education, research, and practice: Systematic review on the promising perspectives and valid concerns. *Healthcare*, 11(6), 887. DOI: 10.3390/healthcare11060887.
- Schick, T., Dwivedi-Yu, J., Dessì, R., et al. (2023). Toolformer: Language models can teach themselves to use tools. arXiv. DOI: 10.48550/arXiv.2302.04761.
- Shinn, N., Labash, B., Gopinath, A., et al. (2023). Reflexion: Language agents with verbal reinforcement learning. arXiv. DOI: 10.48550/arXiv.2303.11366.

- Solaiman, I., Brundage, M., Clark, J., et al. (2019). Release strategies and the social impacts of language models. arXiv. DOI: 10.48550/arXiv.1908.09203.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., & Poole, B. (2021). Score-based generative modeling through stochastic differential equations. arXiv. DOI: 10.48550/arXiv.2011.13456.
- Stokel-Walker, C., & Van Noorden, R. (2023). What ChatGPT and generative AI mean for science. *Nature*, 614, 214-216. DOI: 10.1038/d41586-023-00340-6.
- Taecharungroj, V. (2023). What can ChatGPT do? Analyzing early reactions to the innovative AI chatbot on Twitter. *Big Data and Cognitive Computing*, 7(1), 35. DOI: 10.3390/bdcc7010035.
- Templin, T., Ackland, M. J., Evans, J., et al. (2024). Addressing 6 challenges in generative AI for digital health. *PLOS Digital Health*, 3(9), e0000503. DOI: 10.1371/journal.pdig.0000503.
- Tlili, A., Shehata, B., Adarkwah, M. A., et al. (2023). What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education. *Smart Learning Environments*, 10, 15. DOI: 10.1186/s40561-023-00237-x.
- Touvron, H., Lavril, T., Izacard, G., et al. (2023). LLaMA: Open and efficient foundation language models. arXiv. DOI: 10.48550/arXiv.2302.13971.
- Touvron, H., Martin, L., Stone, K., et al. (2023). Llama 2: Open foundation and fine-tuned chat models. arXiv. DOI: 10.48550/arXiv.2307.09288.
- van den Oord, A., Vinyals, O., & Kavukcuoglu, K. (2017). Neural discrete representation learning. arXiv. DOI: 10.48550/arXiv.1711.00937.
- Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need. arXiv. DOI: 10.48550/arXiv.1706.03762.
- Wang, X., Wei, J., Schuurmans, D., et al. (2022). Self-consistency improves chain of thought reasoning in language models. arXiv. DOI: 10.48550/arXiv.2203.11171.
- Wei, J., Bosma, M., Zhao, V. Y., et al. (2021). Finetuned language models are zero-shot learners. arXiv. DOI: 10.48550/arXiv.2109.01652.
- Wei, J., Tay, Y., Bommasani, R., et al. (2022). Emergent abilities of large language models. arXiv. DOI: 10.48550/arXiv.2206.07682.
- Wei, J., Wang, X., Schuurmans, D., et al. (2022). Chain-of-thought prompting elicits reasoning in large language models. arXiv. DOI: 10.48550/arXiv.2201.11903.
- Weidinger, L., Uesato, J., Rauh, M., et al. (2022). Taxonomy of risks posed by language models. In *FACCT '22* (pp. 214-229). DOI: 10.1145/3531146.3533088.
- White, J., Fu, Q., Hays, S., et al. (2023). A prompt pattern catalog to enhance prompt engineering with ChatGPT. arXiv. DOI: 10.48550/arXiv.2302.11382.
- Wu, Q., Bansal, G., Zhang, J., et al. (2023). AutoGen: Enabling next-gen LLM applications via multi-agent conversation. arXiv. DOI: 10.48550/arXiv.2308.08155.
- Xu, S., Ruan, Y., & Liang, H. (2024). Generative AI and digital creativity: Toward human-AI co-creation in the cultural industries. *Journal of Business Research*, 179, 114753. DOI: 10.1016/j.jbusres.2024.114753.
- Yao, S., Yu, D., Zhao, J., et al. (2023). Tree of thoughts: Deliberate problem solving with large language models. arXiv. DOI: 10.48550/arXiv.2305.10601.
- Yao, S., Zhao, J., Yu, D., et al. (2022). ReAct: Synergizing reasoning and acting in language models. arXiv. DOI: 10.48550/arXiv.2210.03629.
- Yim, M., Yetisgen, M., Harris, W. P., & Kwan, S. W. (2024). Preliminary evidence of the use of generative AI in health care clinical services and patient encounter workflows. *Journal of Medical Internet Research*, 26, e52073. DOI: 10.2196/52073.
- Yu, J., Wang, X., Koh, J. Y., et al. (2022). Scaling autoregressive models for content-rich text-to-image generation. arXiv. DOI: 10.48550/arXiv.2206.10789.

Yu, L., Simianer, P., Eger, S., & Kobayashi, S. (2024). LLMs and the changing skill architecture of work: Evidence from task decomposition experiments. *Information Systems Frontiers*, 26, 1-18. DOI: 10.1007/s10796-024-10559-4.

Zhai, X. (2022). ChatGPT user experience: Implications for education. *SSRN Electronic Journal*. DOI: 10.2139/ssrn.4312418.

Zhang, L., Rao, A., & Agrawala, M. (2023). Adding conditional control to text-to-image diffusion models. In *ICCV* (pp. 3836-3847). DOI: 10.1109/ICCV51070.2023.00358.

Zhu, Z., Xu, Y., Ouyang, W., et al. (2022). Prompt learning for vision-language models: A survey. *arXiv*. DOI: 10.48550/arXiv.2201.08623.

Zhuo, T. Y., Huang, Y., Chen, C., & Xing, Z. (2023). Exploring AI-assisted software development with ChatGPT and GPT-4. *arXiv*. DOI: 10.48550/arXiv.2304.05357.