

Unpacking Generative AI: A Literature Review on Its Transformative Potential and Ethical Challenges

Pratiksha Manjrekar¹, Aasiya Kazmi², Hasan Phudinawala³

P.G. Students, Department of Data Science^{1,2}

Coordinator, Department of Data Science³

Royal College of Arts, Science and Commerce (Autonomous), Mira Road (East)

Abstract: *Generative Artificial Intelligence (GenAI) has emerged as both a catalyst of creativity and a disruptor of ethics. This paper offers a comparative, non-technical synthesis of GenAI's role in text, image, and video generation, integrating insights from fifteen peer-reviewed studies and thirty global case analyses. It finds that GenAI profoundly enhances efficiency and accessibility, automating writing, design, and visual storytelling, yet simultaneously redefines authorship, authenticity, and accountability. From automated newsrooms to AI-generated art and deepfake media, the same innovation that empowers expression also amplifies bias, erodes creative integrity, and destabilizes trust. The analysis reveals a unifying paradox: the pursuit of productivity often compromises originality and ethical coherence. GenAI's value therefore lies not merely in its generative power but in how societies govern its use. The study concludes that ethical literacy, transparency, and human oversight are essential to ensure that automation strengthens, rather than supplants, human creativity.*

Keywords: Generative Artificial Intelligence, Ethics, Automation, Text Generation, Image Generation, Deepfakes, Authorship, Accountability

I. INTRODUCTION

The rapid evolution of Generative Artificial Intelligence (GenAI) has transformed the way humans create, communicate, and interact with digital content. Across industries, from education to entertainment, AI-driven tools now produce text, images, and videos that mirror human expression with remarkable accuracy. These technologies, though primarily designed to enhance efficiency and creativity, have sparked an equally powerful debate over their ethical, cultural, and societal implications. Generative AI not only automates creative labor but also redefines authorship, authenticity, and accountability—forcing society to reconsider what it means to “create.” Generative AI’s capacity to automate writing, illustration, and video production has made it a crucial driver of productivity and innovation. Organizations increasingly rely on text generators for content creation, image models for design and visualization, and video synthesizers for marketing and education. For example, newsroom automation through Associated Press’ AI-written financial reports [17] and enterprise-level assistants such as Klarna’s AI customer service [23] demonstrate tangible gains in cost reduction and time efficiency. Similarly, educational tools like Duolingo Max and Grammarly [24][25] have expanded accessibility and personalization in learning environments. These applications exemplify how GenAI fulfills its promise to reduce manual labor, enhance creative output, and scale human effort beyond traditional boundaries.

However, these benefits coexist with significant ethical dilemmas and practical drawbacks. As Bender et al. [1] caution in their landmark paper *On the Dangers of Stochastic Parrots*, the vast scale of modern language models can perpetuate misinformation, reinforce social biases, and obscure authorship. Similarly, Al-Kfairy et al. [4] emphasize that without transparency and governance, generative tools risk amplifying ethical blind spots rather than closing them. In the visual domain, Bendel [6] and Bird et al. [7] point out how diffusion-based image models democratize creativity yet exploit the intellectual labor of artists whose copyrighted works are scraped for training. The backlash against AI-generated



“Ghibli-style” art [21] is emblematic of this conflict between accessibility and authenticity. In the case of video synthesis, or “deepfakes,” philosophers such as those in *The Distinct Wrong of Deepfakes* [11] argue that this technology fundamentally challenges social trust and the concept of consent in the digital age.

The scope of this paper is to provide a comprehensive, non-technical comparison of how generative AI for text, image, and video generation offers both transformative potential and disruptive risks. Rather than exploring algorithmic architectures, this paper synthesizes peer-reviewed ethical analyses [1]–[15] alongside thirty real-world case studies [16]–[45] to investigate recurring patterns of advantage, disadvantage, and moral tension. It draws on established work across disciplines such as media ethics, philosophy of technology, and digital policy to examine how GenAI simultaneously enhances creativity and destabilizes cultural and professional norms. The central hypothesis underpinning this review is that generative AI improves productivity and creativity by automating repetitive, time-intensive tasks but also generates new ethical concerns surrounding ownership, accountability, bias, and transparency. As Yan et al. [2] demonstrate in the educational context, while language models facilitate scalable learning, they also pose threats to academic integrity and critical thinking. Similarly, Porlezza [5] observes that the use of AI in journalism expands workflow efficiency but risks undermining credibility and public trust. Across modalities, a pattern emerges: every increase in efficiency comes with a proportional increase in ethical complexity.

A significant gap identified across the reviewed scholarship is the fragmentation of discourse. Studies tend to examine text, image, or video generation in isolation, leaving limited cross-modal understanding of shared ethical principles. For instance, Bender et al. [1] and Bendel [6] both discuss bias and creative erosion, yet few comparative frameworks connect their arguments across creative forms. Moreover, real-world case studies—such as the Getty Images vs. Stability AI lawsuit [20] or the Tom Cruise deepfake phenomenon [34]—are rarely integrated with academic literature despite their value in demonstrating how theoretical concerns manifest in practice. This paper bridges that divide by combining empirical and philosophical evidence to articulate a unified view of automation, creativity, and ethics in generative media. In doing so, this study aims to contribute a comparative ethical framework that helps understand how GenAI’s advantages and disadvantages manifest differently across modalities. Text generation often challenges authorship and intellectual integrity, image generation reshapes artistic and aesthetic conventions, while video synthesis questions the very foundations of visual truth and human consent. Yet at their core, all three share the same moral tension: the pursuit of efficiency at the cost of authenticity. By drawing together findings from interdisciplinary sources and lived case examples, this paper argues that responsible governance, ethical transparency, and digital literacy are essential for ensuring that automation serves humanity rather than replacing its creative essence.

II. LITERATURE REVIEW

2.1 Text Generation

Text-based generative AI has been the earliest and most widely adopted branch of creative automation. Tools such as ChatGPT, Jasper, and Grammarly now handle tasks ranging from drafting essays to writing corporate press releases, enabling organizations to scale communication and cut down on manual labor. Scholars emphasize that these tools transform writing from an exclusive intellectual practice into a collaborative human–machine process. Yan et al. [2] observe in their systematic scoping review that large language models in education enhance **accessibility, personalization, and learning scalability**, helping students receive individualized feedback that human instructors could not easily provide. Similarly, Porlezza [5] reports that in journalism, automated writing systems expand newsroom productivity by generating routine reports at unprecedented speed, echoing real-world implementations such as the **Associated Press automated earnings stories** [17], which reduced manual reporting time and allowed human journalists to focus on analysis and interpretation. The **advantages** highlighted across the literature converge around three axes: productivity, creativity support, and inclusivity. Hosseini et al. [3] point out that GenAI tools democratize authorship by enabling non-experts to engage in professional writing and multilingual translation. In corporate settings, systems such as **Klarna’s AI customer-service assistant** [23] showcase measurable improvements in efficiency and profitability through task automation. Even creative sectors benefit—copywriting assistants like **Jasper AI** [25] accelerate content pipelines and empower small businesses lacking professional writing staff.



Nevertheless, most authors stress equally serious **ethical concerns**. Bender et al. [1] warn that large language models risk functioning as “stochastic parrots,” producing fluent yet **unverifiable or biased text** without true comprehension. This echoes scandals such as **Microsoft Tay** [16], which rapidly learned and replicated offensive speech, exposing weaknesses in model supervision. Educational deployments have also prompted **academic-integrity crises**, with students submitting AI-written essays that evade plagiarism detection [20]. Hosseini et al. [3] further raise questions about **transparency and disclosure**—should AI-assisted authorship be explicitly credited, and does failure to do so constitute academic or journalistic misconduct? Ethical implications extend beyond bias to issues of **deskilling and dependency**. Yan et al. [2] and Al-Kfairy et al. [4] note that overreliance on automated text generation may erode critical thinking, language mastery, and human creativity. The **CNET AI-written article controversy** [18], where factual errors and plagiarism were discovered, exemplifies how automation without oversight can damage institutional credibility. Healthcare experiments like **Med-PaLM** [22] demonstrate both the potential and peril of language models: they can ease documentation burdens yet pose safety risks if hallucinated text enters clinical records.

Across these cases, a recurring theme emerges—**efficiency versus authenticity**. GenAI undeniably saves time and expands access, but it also redefines what counts as “authorship.” Hosseini et al. [3] argue that new disclosure norms and governance mechanisms are required to preserve integrity without stifling innovation. The literature therefore positions text-generation AI as both a **productivity catalyst and an ethical testing ground** for future creative technologies.

2.2 Image Generation

The explosion of **text-to-image models** such as DALL·E 2, Midjourney, and Stable Diffusion has revolutionized visual creation, making high-quality imagery available to anyone with a short text prompt. Bendel [6] describes this as the “democratization of digital imagination,” wherein non-artists gain unprecedented creative agency. Bird et al. [7] extend this argument, identifying tangible **advantages** including cost efficiency, speed, and the ability to explore multiple stylistic variations. Real-world adoptions—such as **IKEA’s AI visualization tools** [28] or **Adobe Firefly’s licensed image generation** [26]—illustrate practical applications that merge creative experimentation with commercial productivity. Ethically, however, image generation raises complex questions about **authenticity, ownership, and bias**. Multiple studies [9][10] document how diffusion models reproduce and amplify **gender and cultural stereotypes**, revealing structural biases embedded in training datasets. Gierbach et al. [9] show that prompts containing occupational titles often yield images reflecting outdated gender norms, aligning with similar findings in Wu et al. [10]. These biases parallel controversies such as **ThisPersonDoesNotExist** [24] and **RashmiKandanna’s AI-manipulated photos** [30], where realistic yet fabricated imagery challenged public perceptions of identity and consent.

Copyright and data provenance dominate the ethical discourse. Bendel [6] and Bird et al. [7] note that millions of copyrighted artworks were scraped without consent to train image models, provoking legal battles such as **Getty Images v. Stability AI** [20] and the **artists’ class-action lawsuit (Andersen / McKernan / Ortiz)** [21]. These cases underscore how GenAI complicates traditional notions of artistic labor and ownership. At the same time, licensed approaches like **Shutterstock’s indemnified AI features** [27] and **IBM–Adobe collaborations** [29] demonstrate emerging industry standards attempting to reconcile innovation with creator rights. Cultural backlash also figures prominently. The **“Ghibli-style” AI-art trend** [19] provoked criticism from Studio Ghibli enthusiasts and artists who viewed algorithmic imitation as disrespectful to the craftsmanship of Hayao Miyazaki. Bendel [6] interprets such incidents as signs of “aesthetic homogenization,” where models flatten diverse artistic voices into marketable sameness. Furthermore, the malicious misuse of generative technologies—from DeepNude’s non-consensual imagery [22], which highlights the prevalence and societal impact of synthetic intimate content and underscores the ethical challenges and need for regulatory frameworks [8], to political photo manipulation—reveals the darker side of democratized AI generation. Collectively, these works depict image-generation AI as a **paradox of empowerment and exploitation**. It liberates visual creativity from technical barriers but jeopardizes authenticity, representation, and artistic integrity. The literature repeatedly calls for **transparent data licensing, algorithmic watermarking, and fair-compensation frameworks** as potential mitigations. As Bird et al. [7] suggest, sustainable progress requires embedding ethical reflection directly into design and deployment processes.



2.3 Video Generation

Among all generative modalities, **video synthesis and deepfakes** provoke the sharpest ethical alarm because they directly engage with **truth, identity, and consent**. Philosophical analyses such as *The Distinct Wrong of Deepfakes* [11] and *Deepfakes and Dishonesty* [12] argue that synthetic videos are not merely deceptive but fundamentally corrosive to social trust. By fabricating realistic motion and voice, these systems can manipulate evidence, undermine journalism, and violate personal autonomy. High-profile demonstrations like the **Jordan Peele / Obama PSA** [31] and viral **Tom Cruise TikTokdeepfakes** [33] have illustrated both the creative allure and the manipulative potential of this technology. The **advantages** of generative video tools, though less discussed, include **educational simulation, entertainment, and cost-efficient production**. Platforms like **Synthesia** [37] and **Khan Academy's AI tutors** [40] use synthetic avatars to deliver multilingual, accessible instruction, showing how ethical use can expand inclusivity. The Wiley study on ecological communication [15] further highlights how AI-generated visualizations aid environmental education by simplifying complex phenomena into digestible narratives. Thus, video generation, when responsibly governed, offers significant value in pedagogy, advertising, and cross-cultural communication.

The darker side revolves around **misinformation, consent, and financial fraud**. The *Journal of Ethics and Social Philosophy* paper *Deepfakes, Deep Harms* [13] details how fabricated videos erode collective epistemic trust—viewers begin to doubt even authentic footage. Real-world incidents reinforce these worries: **corporate deepfake fraud in Hong Kong (Arup case)** [35] caused multimillion-dollar losses, and **celebrity impersonation scams such as Steve Burton's** [36] exploited AI-generated likenesses to deceive victims. In India, **political deepfakes of actors like Ranveer Singh** [39] have led to lawsuits and calls for stronger regulation [38]. Academic discussions emphasize that unlike textual or pictorial outputs, video synthesis **collapses temporal and auditory realism**, making detection and accountability harder. The *Journal of Interactive Advertising* study [14] empirically shows that even when audiences are informed a video is synthetic, **trust and purchase intention still decline sharply**, confirming tangible reputational risks for brands. Environmental perspectives [15] add another dimension, noting the high energy demands of generative video training, linking digital ethics to ecological sustainability.

Overall, the literature portrays video-generation AI as the **most ethically volatile** of the three domains—its realism magnifies harm potential. Scholars call for **mandatory disclosure policies, consent-based data curation, and public-awareness campaigns** to mitigate misuse. Yet, like its textual and visual counterparts, the same system that threatens authenticity also offers tools for education, creativity, and inclusion when applied transparently and responsibly.

Table 1: Comparative Summary of Key Literature on Ethical, Societal, and Technical Aspects of Generative AI

Author(s) & Paper Title	Focus Objective	Key Findings	Advantages / Benefits	Ethical Concerns / Risks
Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). <i>On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?</i>	Risks of large language models	LLMs may propagate bias and misinformation due to dataset limitations	Automation, improved efficiency in language tasks	Bias, misinformation, overreliance on models
Yan, L., Sha, L., Zhao, L., Li, Y., Martinez-Maldonado, R., Chen, G., & Gašević, D. (2023). <i>Practical and Ethical Challenges of Large Language Models in Education: A Systematic Scoping Review</i>	LLMs in education	Identified practical and ethical challenges in educational applications	Personalized learning, reduced workload	Academic integrity, fairness, bias



Hosseini, M., Resnik, D. B., & Holmes, K. (2023). <i>The Ethics of Disclosing the Use of Artificial Intelligence Tools in Writing Scholarly Manuscripts</i>	Disclosure of AI in academic writing	Examines ethical requirements for transparency	Encourages responsible AI use	Misrepresentation, lack of transparency
Al-kfairy, M., Mustafa, D., Kshetri, N., Insiew, M., & Alfandi, O. (2024). <i>Ethical Challenges and Solutions of Generative AI: An Interdisciplinary Perspective</i>	Ethical challenges in Generative AI	Offers interdisciplinary perspective on ethics and solutions	Increased productivity and creative output	Bias, misuse, intellectual property issues
Porlezza, C. (2024). <i>AI Ethics in Journalism (Studies): An Evolving Field Between Research and Practice</i>	AI ethics in journalism	Reviews evolving AI ethics in media	Speed and efficiency in news production	Misinformation, accountability
Bendel, O. (2023). <i>Image Synthesis from an Ethical Perspective</i>	Ethics of image synthesis	Highlights ethical concerns in AI-generated images	Creativity and automation	Copyright, bias, deepfakes
Bird, C., Ungless, E. L., & Kasirzadeh, A. (2023). <i>Typology of Risks of Generative Text-to-Image Models</i>	Risks of text-to-image models	Typology of ethical risks in text-to-image generation	Visual content generation	Stereotypes, harmful content
McGlynn, L., Downing, J., & Datt, A. (2024). <i>Non-Consensual Synthetic Intimate Imagery: Prevalence, Attitudes, and Behaviors</i>	Non-consensual synthetic imagery	Prevalence and societal impact of non-consensual images	Raises awareness and promotes regulation	Privacy violation, non-consent
Girrbach, L., Alaniz, S., Smith, G., & Akata, Z. (2025). <i>A Large-Scale Analysis of Gender Biases in Text-to-Image Generative Models</i>	Gender bias in text-to-image models	Large-scale analysis shows models reinforce gender stereotypes	Automates content creation	Gender bias, stereotype amplification
Wu, Y., Nakashima, Y., & Garcia, N. (2024). <i>Gender Bias Evaluation in Text-to-Image Generation: A Survey</i>	Survey on gender bias evaluation	Summarizes evaluation methods for bias in text-to-image generation	Insight into fairness improvement	Gender bias, cultural bias
De Ruiter, A. (2021). <i>The Distinct Wrong of Deepfakes</i>	Ethics of deepfakes	Philosophical analysis of deepfake harms	Critical awareness	Misinformation, reputational harm
Flattery, T., & Miller, C. B.	Deepfakes and	Explores impact on	Understanding	Deception,



(2024). <i>Deepfakes and Dishonesty</i>	dishonesty	trust and deception	media literacy	manipulation
Rini, R., & Cohen, L. (2022). <i>Deepfakes, Deep Harms</i>	Deepfakes and societal harms	Analyzes social and ethical implications of deepfake media	Awareness of deepfake risks	Privacy, misinformation, manipulation
Powers, G., Johnson, J. P., & Killian, G. (2023). <i>To Tell or Not to Tell: The Effects of Disclosing Deepfake Video on US and Indian Consumers' Purchase Intention</i>	Effects of disclosing deepfake videos	Investigates consumer perception	Consumer protection	Trust erosion, ethical disclosure issues
Rillig, M. C., Mansour, I., Hempel, S., Bi, M., König-Ries, B., & Kasirzadeh, A. (2024). <i>How Widespread Use of Generative AI for Images and Video Can Affect the Environment and the Science of Ecology</i>	Environmental impact of generative AI	Explores ecological consequences of AI computation	Supports large-scale analysis and research	Carbon footprint, environmental sustainability

III. ANALYSIS AND CONCLUSION

A cross-comparison of generative AI's applications in text, image, and video generation reveals a unifying narrative: **automation and creativity are intertwined in both empowerment and ethical compromise**. Across all modalities, GenAI enhances productivity by reducing manual labor and enabling scalability, yet this same efficiency raises profound questions about authenticity, authorship, and accountability. The analysis shows that while the tools differ in output—words, visuals, or motion—their social consequences follow remarkably similar trajectories.

In **Text Generation**, the dominant advantages lie in **productivity and accessibility**. Language models streamline communication, facilitate learning, and democratize writing for non-experts. Case studies like the **Associated Press automated journalism initiative** [17] and **Klarna's customer-service automation** [23] demonstrate measurable economic and operational benefits. However, the same mechanisms that enable such scale also generate **ethical friction**. The “stochastic parroting” effect identified by Bender et al. [1] encapsulates the central issue—machines that simulate understanding without accountability. Incidents such as **CNET's AI-written factual errors** [18] and **students' misuse of ChatGPT** [20] illustrate the thin boundary between assistance and deception. Consequently, text generation embodies the dual nature of GenAI: an efficient partner in creativity that simultaneously threatens the intellectual integrity of authorship. In **Image Generation**, the ethical balance becomes even more visible due to its immediate sensory impact. Models like Stable Diffusion and Midjourney have redefined how design, marketing, and art are produced, exemplified by **Adobe Firefly** [26] and **IKEA's visualization systems** [28]. These systems amplify creativity and accessibility, yet they are also the epicenter of legal and cultural backlash. The **Getty Images lawsuit** [20] and **artists' class action against AI platforms** [21] mark the beginning of a global debate on ownership in the age of algorithmic art. The **Ghibli-style art controversy** [19] and **DeepNude's misuse** [22] further highlight moral boundaries—how far can inspiration stretch before it becomes exploitation? The reviewed literature [6][7][9] suggests that while image generators empower creative freedom, they also **commodify originality**, leading to homogenization and the erosion of artistic identity. Hence, the ethical burden in image generation lies not merely in data bias or copyright but in the preservation of cultural authenticity. In **Video Generation**, the stakes escalate dramatically. Deepfakes introduce a **crisis of trust**—where seeing is no longer believing. While enterprise tools like **Synthesisia** [37]



and educational applications such as **Khan Academy's AI tutors** [40] showcase positive, lawful uses, malicious adaptations cause disproportionate harm. The **Hong Kong corporate fraud deepfake** [35] and **celebrity impersonation scandals** [36][39] reveal that the technology's realism amplifies its potential for deception, reputational damage, and financial crime. As philosophers in *The Distinct Wrong of Deepfakes* [11] argue, the moral gravity of video synthesis arises not only from its content but from its **capacity to manipulate trust at a societal level**. In contrast to text or image, which may mislead through interpretation, synthetic video **manufactures false reality**, challenging epistemic and ethical foundations alike.

When examined comparatively, several **common ethical threads** emerge. First, **data provenance and consent** recur as unresolved dilemmas. Whether scraping copyrighted texts, artistic works, or personal likenesses, the training of GenAI consistently tests the limits of informed consent and ownership. Second, **bias and exclusion** persist across modalities. Text models reproduce linguistic stereotypes, image models reinforce visual hierarchies, and video generators inherit representational gaps—showing that algorithmic ethics cannot be siloed by medium. Third, **authorship and accountability** remain ambiguous. In all three domains, human users bear nominal responsibility, yet the opacity of model behavior diffuses ethical liability, leading to governance gaps in academic, creative, and commercial contexts. Despite these challenges, the literature does not frame GenAI solely as a threat. Rather, it presents it as a **transformative catalyst that demands moral adaptation**. Authors such as Al-Kfairy et al. [4] and Bendel [6] emphasize that governance, transparency, and participatory policymaking can align AI's benefits with human values. Case studies like **Shutterstock's licensing framework** [27] and **Adobe's indemnified AI model** [26] demonstrate that ethical design and corporate responsibility are viable counterbalances. Similarly, educational initiatives that disclose AI use, as suggested by Hosseini et al. [3] and Yan et al. [2], illustrate how ethical literacy can mitigate misuse without suppressing innovation. Ultimately, the comparative review affirms the paper's hypothesis: **Generative AI enhances creativity and efficiency but simultaneously generates ethical tensions proportionate to its power**. Each modality embodies a unique intersection of opportunity and risk—text generation challenges intellectual honesty, image generation redefines artistic authenticity, and video generation destabilizes epistemic trust. Yet all converge on one critical insight: technological advancement without ethical alignment threatens to automate not only creativity but also accountability. The future of GenAI, therefore, depends not merely on technical progress but on the cultivation of **moral frameworks, transparent governance, and collective digital responsibility** that ensure automation serves as an extension of human ingenuity rather than its replacement.

REFERENCES

- [1] Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?* Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT).
- [2] Yan, L., Sha, L., Zhao, L., Li, Y., Martinez-Maldonado, R., Chen, G., & Gašević, D. (2023). *Practical and Ethical Challenges of Large Language Models in Education: A Systematic Scoping Review*. British Journal of Educational Technology, Wiley.
- [3] Hosseini, M., Resnik, D. B., & Holmes, K. (2023). *The Ethics of Disclosing the Use of Artificial Intelligence Tools in Writing Scholarly Manuscripts*. Research Ethics, SAGE Publications.
- [4] Al-kfairy, M., Mustafa, D., Kshetri, N., Insiew, M., & Alfandi, O. (2024). *Ethical Challenges and Solutions of Generative AI: An Interdisciplinary Perspective*. Informatics, MDPI.
- [5] Porlezza, C. (2024). *AI Ethics in Journalism (Studies): An Evolving Field Between Research and Practice*. Journalism Studies, SAGE.
- [6] Bendel, O. (2023). *Image Synthesis from an Ethical Perspective*. AI & Society, Springer.
- [7] Bird, C., Ungless, E. L., & Kasirzadeh, A. (2023). *Typology of Risks of Generative Text-to-Image Models*. Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES).
- [8] L. McGlynn, J. Downing, and A. Datt, *Non-Consensual Synthetic Intimate Imagery: Prevalence, Attitudes, and Behaviors*, *Proc. ACM Hum.-Comput. Interact.*, vol. 8, no. CSCW2, pp. 1–28, 2024.



- [9] Girrback, L., Alaniz, S., Smith, G., & Akata, Z. (2025). *A Large-Scale Analysis of Gender Biases in Text-to-Image Generative Models*. arXiv.
- [10] Wu, Y., Nakashima, Y., & Garcia, N. (2024). *Gender Bias Evaluation in Text-to-Image Generation: A Survey*. arXiv.
- [11] De Ruiter, A. (2021). *The Distinct Wrong of Deepfakes*. Philosophy & Technology, Springer.
- [12] Flattery, T., & Miller, C. B. (2024). *Deepfakes and Dishonesty*. Philosophy & Technology, Springer.
- [13] Rini, R., & Cohen, L. (2022). *Deepfakes, Deep Harms*. Journal of Ethics and Social Philosophy.
- [14] Powers, G., Johnson, J. P., & Killian, G. (2023). *To Tell or Not to Tell: The Effects of Disclosing Deepfake Video on US and Indian Consumers' Purchase Intention*. Journal of Interactive Advertising, Taylor & Francis.
- [15] Rillig, M. C., Mansour, I., Hempel, S., Bi, M., König-Ries, B., & Kasirzadeh, A. (2024). *How Widespread Use of Generative AI for Images and Video Can Affect the Environment and the Science of Ecology*. Ecology Letters, Wiley.
- [16] *Microsoft Tay Chatbot Controversy*. (2016). The Guardian.
- [17] *Associated Press Automated Earnings Reports*. (2014–present). The Associated Press.
- [18] *CNET / Red Ventures AI-Written Articles Investigation*. (2022–2023). The Verge.
- [19] *Studio Ghibli-Style AI Art Backlash*. (2025). AP News.
- [20] *Fake Legal Citations in ChatGPT-Generated Briefs (Mata v. Avianca)*. (2023). Reuters.
- [21] *Getty Images v. Stability AI Copyright Litigation*. (2023–2025). Reuters.
- [22] *Healthcare LLMs / Med-PaLM Hospital Pilots*. (2023–2024). AI Magazine.
- [23] *Klarna AI Customer-Service Assistant Deployment*. (2024). Klarna Press Release.
- [24] *Duolingo Max (GPT-4-Powered Tutoring Features)*. (2023). TechCrunch.
- [25] *Grammarly and Jasper AI Writing Assistants*. (2020s–present). Company Websites.
- [26] *Adobe Firefly: Licensed Dataset & Ethical AI*. (2023). Adobe Newsroom.
- [27] *Shutterstock Generative AI and Indemnification Policy*. (2023). Shutterstock Investor Relations.
- [28] *IKEA Virtual Room Visualization Using AI Imagery*. (2023). IKEA Design Labs.
- [29] *IBM–Adobe Collaboration on Licensed Content for AI*. (2023–2024). Adobe Newsroom.
- [30] *Rashmi Mandanna AI Image Misuse Case*. (2024), NDTV India.
- [31] *Jordan Peele / Obama Deepfake PSA Demonstration*. (2018). Vox Media.
- [32] *ZAO Face-Swap App Privacy Backlash*. (2019). The Guardian.
- [33] *Tom Cruise Deepfakes on TikTok*. (2021). The Guardian.
- [34] *Mark Zuckerberg Deepfake Artwork Exhibition*. (2019). Artnet News.
- [35] *Corporate Deepfake Fraud in Hong Kong (Arup Case)*. (2024). The Guardian.
- [36] *Steve Burton Celebrity Deepfake Scam*. (2025). People.com.
- [37] *Synthesia Enterprise Synthetic Video Platform*. (2019–present). Synthesia Case Studies.
- [38] *Indian Public Figure Deepfakes and Legal Responses*. (2023–2025). NDTV India.
- [39] *Ranveer Singh Political Deepfake Incident*. (2024). NDTV India.
- [40] *Khan Academy AI Tutors and Educational Video Integration*. (2023). WPBS Public Media.
- [41] *Rashmi Mandanna / Multiple Indian Celebrities Deepfake Reactions*. (2023–2025). NDTV India.
- [42] *Corporate Use of AI in Marketing Videos: Disclosure Policy Cases*. (2024). Taylor & Francis Reports.
- [43] *Deepfake Awareness Campaigns in Journalism*. (2023). SAGE News Briefs.
- [44] *Environmental Implications of AI Video Generation*. (2024). Wiley Online Library.
- [45] *Media Policy Debates on AI-Generated Content Governance*. (2025). Reuters Technology

