Volume 53 Number 3
Spring 2026

Published by the Association for Computing Machinery
Special Interest Group on Computers and Society

ACM SIGCAS Computers and Society

TABLE OF CONTENTS

SIGCAS

Introduction to SIGCAS

SIGCAS Computers and Society is the ACM Special Interest Group that addresses the social aspects and ethical consequences of widespread computer usage. The main goals of SIGCAS are to raise awareness about the impact that technology has on society and to support and advance the efforts of those who are involved in this important work. Our members are computer professionals from both industry and academia, as well as ethicists, psychologists, sociologists, and others. We welcome students from a variety of disciplines. Our areas of involvement include computer ethics, universal access to computer technology, security, privacy, and reliability. We collaborate with other ACM bodies engaged in related work, such as COPE, USACM, SIGITE, and SIGCSE. ACM *Computers & Society* is an online publication accessible via the ACM Digital Library and the SIGCAS website. The newsletter aims to be an effective communication vehicle between the members of the group and the outside world.

SIGCAS Computers and Society Readers and writers are invited to join and participate actively in this Special Interest Group. Membership is open to all for $25 USD per year and to students for $10 USD per year. The link to join can be found on our website at www.sigcas.org.

Contribute. The Editor-in-Chief invites contributions of all types of written material (i.e., articles, working papers, news, interviews, reports, book reviews, bibliographies of relevant literature, and letters) on all aspects of computing that have a bearing on society and culture.

Please note that it is *NOT* a peer-reviewed publication. Submissions are checked for relevance, accessibility, and basic suitability by the editors but are not fully peer reviewed.

For the latest Call(s) for Submissions, or instructions regarding formatting guidelines and copyright policy, please see the website: http://www.sigcas.org/. Submissions may be sent to editors_sigcas@acm.org.

However, as a contributing author, you retain copyright to your article or other material, and ACM will refer requests for republication directly to you.

Membership Benefits

- Subscription to the online publication ACM SIGCAS Computers and Society, which is published three to four times a year.
- Access to the full archive of the online publication and its printed predecessor in the ACM DL.
- Discounted registration fees for SIGCAS-sponsored conferences and workshops. “In cooperation” sponsor of several ACM and non-ACM conferences related to SIGCAS’ interests, including LIMITS.
- SIGCAS presents two annual awards: The Making a Difference Award and the SIGCAS Outstanding Service Award.
- SIGCAS-ANNOUNCE mailing list: includes regular announcements of upcoming conferences and calls for participation. SIGCAS-Talk mailing list: enables member-to-member interactions, and the committee will seek to stimulate discussion on this list amongst members. Subscription to the list is restricted to SIGCAS members and is optional.

From the Editor

by Chinasa T. Okolo

Welcome to the latest edition of ACM SIGCAS *Computers and Society*! This is the third edition published in my capacity as Editor-in-Chief, and I'm glad to keep the momentum going with another great set of articles.

This edition covers a wide range of topics, including short pieces on the future of the SIGCAS community, the risks of AI-induced psychosis, and the role of emotions in AI creativity; articles covering cybersecurity and compliance in fintech, open-source data infrastructure for embodied AI, and AI governance and democratic resilience in Africa; and finally, a zine exploring responsible computing patterns.

Over the next year, the SIGCAS Executive Committee plans to further expand our impact by engaging more with our members, finding new opportunities to strengthen membership, and collaborating with organizations sharing our mission of responsible computing. We are currently distributing a community survey, and request all members to complete it, as this will provide the Executive Committee with valuable insights to understand how we can improve your experience with SIGCAS.

As always, please get in touch if you'd like to volunteer or have any questions or concerns! We would love to have your feedback.

Chinasa T. Okolo, Ph.D.
Editor-in-Chief, *ACM SIGCAS Computers and Society*

Call for Short Pieces

The SIGCAS Publications Committee aims to encourage more voices and varied perspectives on topics relevant to computers and society. In addition to longer pieces, we seek short pieces that are relevant, provocative, diverse, and unexpected for our issues of *Computers and Society*. Your short piece could raise arguments, issues, critical questions, resource needs, current work, research, reviews, discussions, etc. We currently plan to run several short pieces per issue.

To be considered for the next newsletter, please submit your short piece to the SIGCAS *Computers and Society* Editor-in-Chief, Chinasa T. Okolo, through email at editors_sigcas@acm.org. Submissions are due on March 15th, 2026 (the subsequent issue will be published in April 2026). Please include "Short Piece" in the subject line. Submissions should be a maximum of 1,200 words, drafted in Microsoft Word or Google Docs format, and references formatted in IEEE style.

SIGCAS Editors may be contacted at editors_sigcas@acm.org with questions about the *Computers and Society* newsletter submission process.

SIGCAS *Computers and Society* Areas of Interest

Advanced Computing
Art, Media, Computing, and Society
Climate, Biodiversity, & the Environment
Computing Ethics Education
Conflict & Militarization
Civic Technology
Digital Privacy
Economic Development
Ethics & Responsible Computing
For or From The Archives
Gender & Women's Studies
Global Development Studies
Health & Medicine
History of Computing
Human Computer Interaction
Human Rights & Immigration
Information Security & Cybersecurity
Labor & the Workforce
Peace & Conflict Studies
Policy & Computing
Practitioner Perspectives
Public Opinion
Recent News & Reviews (Articles, Books, Movies, etc.)
Religion & Spirituality
Social Justice
Student Voices
Systemic Racism & Marginalization
Urban Computing & Smart Cities
User Experience & Design

SIGCAS Community Survey

ACM SIGCAS is asking for your input on what SIGCAS should become to support you, society, and ACM. Please take a few minutes to complete the survey from ACM SIGCAS! Your input is important to inform the future of responsible computing.

SIGCAS strives to create an enriching and supportive environment that enables discussion and understanding of the direct and indirect effects of computing, both positive and negative, on all. This survey will help inform SIGCAS on critical issues and directions, identify areas where we're doing well, and suggest things that we can improve or start up. We look forward to your responses and to sharing the results!

The survey will remain open until September 15th, 2026.
https://utdallas.qualtrics.com/jfe/form/SV_ePtHNG4zVKZHpjM

The Future is Digital: So What's the Problem? ACM SIGCAS Competition on Big Ideas!

SIGCAS is pleased to announce an exploration/competition to encourage critical thinking of all kinds and across diverse communities about the opportunities and challenges of computing in today's world, and to help cultivate future discussions and actions that build on this thinking.

Three awards will be given in each category - professional submissions and student submissions. The awards in each category will be a Best Paper Award of $500 USD and two Honorable mentions of $250 USD each.

Submissions will be accepted until October 15th, 2026.

More details on this announcement, format and length requirements, and judging information can be found here: https://www.sigcas.org/the-future-is-digital-so-whats-the-problem/

SHORT PIECE

Reflecting on the Future of SIGCAS

by Brian Krupp
Case Western Reserve University

I first joined SIGCAS in 2020, in some ways, by accident. I was up for my ACM renewal, and some of the SIGs that I was a member of didn't seem to focus on the issues that I had started to become more interested in. The lab that I was leading at the time started to work on research that was focusing more on issues in communities where computing could provide solutions to some of these issues. As I scrolled through the list of available SIGs, I came across SIGCAS and read the description, and it seemed like a perfect fit: *"The ACM Special Interest Group on Computers and Society brings together computer professionals, specialists in other fields, and the public at large to address concerns and raise awareness about the ethical and societal impact of computers."* The work our lab was doing was focused on societal impact, and I was interested to join a community that had similar interests. This seemed like a perfect fit, so I joined.

During those initial years of joining SIGCAS, the organization seemed to be flourishing. The community was active and had several events, including the SIGCAS Showcase, Works in Progress, two sponsored conferences, including GoodIT and COMPASS, meetings at SIGCSE, while also maintaining a fairly active Listserv and an active Discord server. Personally, the organization had a direct benefit within our own lab, where we were invited to discuss our research with Campus Plate, a platform that helps reduce food waste and food insecurity on college campuses. Through this invited talk, I connected with other members, which then led to more connections and opportunities. Our work was highlighted in an article on Computing for Social Good, and I was invited to attend the excellent POSSE workshop that Greg Hislop and Heidi Ellis organize[1]. These connections, this collaboration, were all made possible through the SIGCAS community.

Then all of a sudden, the activity within SIGCAS came almost to a full stop. I took over the Works in Progress events, and we had some excellent talks, but attendance waned. The Listserv and Discord server became less active. Around that time, I was able to organize a meeting at SIGCSE, and some of the members of SIGCAS who attended all had the same question: What is going on with SIGCAS? I did not have an answer, and I was wondering the same. I knew there were people out there doing amazing and impactful work, so why are we not sharing it? Why did the activity wane?

I cannot explain what happened in that quieter time of SIGCAS, and I don't think it is fair or a worthwhile exercise to direct blame on anyone. But when the email came out in September 2025 from the past-chair of SIGCAS, who found out that the organization was looking to transition from a traditional SIG model, I was

[1] https://teachingopensource.org/Main_Page

caught off guard, as were other members, based on the replies. The encouraging response to this, though, was that past leaders and current members felt the urge to do something, and they stepped up to the plate to do so. The group worked quickly to intervene, and our current chair, Samuel Mann, presented the "Phoenix Plan" to the SIG governing board, which received approval and praise.

Ultimately, it is not just on these individuals to help resurrect SIGCAS; it is on all of us to do so. We can shape what SIGCAS is in the future. If there is a lesson to be learned from this, it is that we cannot rely on any one individual or a group of individuals for SIGCAS to flourish. It takes the collective effort of **all of us** for SIGCAS to flourish.

SIGCAS is important to me because it is a space, a community, where people care about similar issues that I do. It is a community where people ask the critical questions that I find myself asking, which others may ignore. It is a community that, in many ways, wants a future for the next generation that is more open, sustainable, and beneficial to society, where it sees a future where computing can fulfill its promise in being a positive force in society.

Our organization is needed now more than ever. Our professional society needs a community where practitioners, educators, students, researchers, and anyone can share ideas, have support, discuss, and debate the challenges that face us today. The future of SIGCAS is in our hands. It is a shared responsibility among our members to shape what the organization focuses our efforts on. In many ways, we have a clean slate to set the direction of what we do. We need to think big, we need to be collaborative, and we need to share our ideas, so we can seize this opportunity where SIGCAS can flourish as an organization and live up to its vision.

So, what can you do? Tell a friend about SIGCAS, bring them to a coffee hour or future event, join our Discord server, volunteer to help with running current or future events, ask questions, and propose ideas. If there is one thing I learned from this group, it is that it is an open and welcoming community. While the executive committee has been meeting fairly regularly in an effort to help resurrect SIGCAS, we need your ideas and your efforts. The lesson I learned since joining SIGCAS is that it takes the community for it to flourish. You are part of that community, and I hope you will join this effort!

SHORT PIECE

AI Psychosis Is a Measured Risk of Capitalist Technology

by Sean Cho Ayres
Kennesaw State University

AI psychosis was a measured risk of capitalist technology.

The cases of AI psychosis continue to rise and increasingly appear in news feeds. In 2024, a Florida teen committed suicide after prolonged conversations with a chatbot [1]. In 2025, a similar incident involving a California teen was reported [2]. However, it is not just a problem among youth. A Toronto man was reportedly convinced by a chatbot that he had discovered a revolutionary mathematical framework with world-altering implications [3]. These are not isolated incidents; data suggests that roughly half a million users each week display signs of psychosis or mania related to AI chatbots [4]. Sure, at some level, this is a result of the modern-day mental health crisis. However, I also want to highlight that AI psychosis is a measured risk for technology companies, encouraged by the design of large language models and their marketing.

While user responsibility plays a role, users can read the fine print stating that "AI can fabricate information," LLMs are also marketed as more human and capable than they actually are. Contemporary large language models are "next best word guessers" with access to vast amounts of web knowledge [5]. Yet, their interfaces and marketing foster the impression that they are thinking, collaborating, and even empathizing [6].

The machine is marketed as human. Consider the subtle but effective UX choices that create an illusion of humanness. When a user interacts with an LLM, it does not simply process tokens; instead, the system is personified as "thinking" [6]. LLMs are not just agentic tools—they are positioned as "coworkers" [6]. ChatGPT and Claude greet users with phrases like "How can I help you today?" and adopt a warm, welcoming tone. These UX decisions are intentional and aim to make users believe they are interacting with something much more complex and human than lines of code [6].

LLMs are now marketed for almost every aspect of life, not just for simple tasks like cleaning up spreadsheets or writing poems [7]. Instead, companies promote LLMs as replacements or augmentations for human thinking and experience. In May 2026, OpenAI introduced a personal finance option that allows users to connect their banking accounts to the AI [8]. For medical concerns, products like Claude for Healthcare offer plain-language interpretations of lab results, potentially bypassing traditional healthcare touchpoints [9]. Digital companions like Replika are promoted as "the AI companion who cares," available 24/7 in users' pockets [10].

I don't mean to sound overly snarky, nor do I blame the profit-driven companies. Market demand is clear: overwhelmed medical systems, loneliness, and the need for financial advice all drive adoption [8, 9]. Hopping onto digital technologies for advice in high-risk domains is not new—WebMD and online self-tests (e.g., for ADHD) have long existed [11]. However, unlike LLMs, these sites do not encourage users to believe in unproven home remedies or claim to be world-class experts [11].

User responsibility remains important, but the persistent suggestion—through marketing, media, and public statements—that LLMs are approaching AGI or even consciousness further blurs the line between tool and trusted advisor [12]. With figures like Richard Dawkins suggesting AI may represent the next phase of evolution [12], it is no wonder users are encouraged to place significant trust in these systems.

Look, I know AI, generative AI, and commercial products like ChatGPT and Claude can be useful. But as we rush to deploy new LLM wrappers, chase benchmark improvements, and scale toward general intelligence, we must consider the end users—the everyday people whose well-being depends on these technologies. Their safety and mental health should be at the forefront of every design and deployment decision.

References

[1] Angela Yang. 2025. “Mom who sued Character.AI over son’s suicide says the platform’s new teen policy comes ‘too late’,” NBC News. https://www.nbcnews.com/tech/tech-news/characterai-bans-minors-response-megan-garcia-parent-suing-company-rcna240985

[2] Clare Duffy. 2025. “Parents of 16-year-old Adam Raine sue OpenAI, claiming ChatGPT advised on his suicide,” CNN Business. https://www.cnn.com/2025/08/26/tech/openai-chatgpt-teen-suicide-lawsuit

[3] Kashmir Hill and Dylan Freedman. 2025. “Chatbots Can Go Into a Delusional Spiral. Here’s How It Happens,” The New York Times. https://www.nytimes.com/2025/08/08/technology/ai-chatbots-delusions-chatgpt.html

[4] OpenAI. 2025. “Strengthening ChatGPT’s responses in sensitive conversations,” OpenAI. https://openai.com/index/strengthening-chatgpt-responses-in-sensitive-conversations/

[5] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?,” in Proc. 2021 ACM Conf. on Fairness, Accountability, and Transparency (FAccT ’21), New York, NY, USA, pp. 610–623.

[6] Anthropic, “Claude Cowork,” Anthropic. https://claude.com/product/cowork

[7] OpenAI. 2022. “Introducing ChatGPT,” OpenAI. https://openai.com/blog/chatgpt

[8] OpenAI. 2026. “A new personal finance experience in ChatGPT,” OpenAI. https://openai.com/index/personal-finance-chatgpt/

[9] Anthropic. 2026. “Advancing Claude in healthcare and the life sciences,” Anthropic. https://www.anthropic.com/news/healthcare-life-sciences

[10] Luka, Inc., “Replika: The AI companion who cares,” Replika. https://replika.com

[11] Psychology Today, “ADHD Test,” Psychology Today. https://www.psychologytoday.com/us/tests/health/adhd-attention-deficit-disorder-test

[12] Richard Dawkins. 2026. “Is AI the next phase of evolution? Claude appears to be conscious,” UnHerd. https://unherd.com/2026/05/is-ai-the-next-phase-of-evolution/

SHORT PIECE

Creativity without Emotions? Irrational Intelligence: Why AI Research Needs Emotions

by Maritza Mills

University of South Carolina

Introduction

Since inception, AI research has been inspired by the idea of building machines that think and behave like humans. Or put another way, machines that think and act rationally [14]. This paradigm has its historical foundations in the assumption that what makes humans different from other forms of life is our capacity for self-reflective, rational thought. This assumption has caused research into artificial intelligence and human intelligence more generally to undervalue the role of emotions in generalized problem-solving. The very definition of rationality has, for much of Western history, discounted emotions from its conception. Where emotions are investigated, they are often treated as irrational, immaterial, or illusory aspects of human experience. This has been an unfortunate mistake.

My goal is to bring awareness to a different way of viewing emotions, with a focus on how emotion research may better advance subsequent research into human and machine intelligence. I take the position that emotions are essential to human reasoning, creativity, and decision-making. They do not interfere with or suppress human rationality but are instead constitutive of it. I'll begin by briefly referencing a bit of history regarding how we got to think of emotions as being somehow different from rationality. I then show how those historical origins continue to have an observable impact on AI research. This particular article will examine how such attitudes influence model design and evaluation with regard to computational creativity. I offer my own view as to how this limits the potential for creative AI. The importance of emotions in human cognition provides support for the necessity of humanistic approaches to artificial intelligence.

A Very Brief History of Emotion-Cognition Interaction

The prevailing paradigm is that emotions are somehow separate from other aspects of our cognition: attention, memory, perception, language, and most of all, reasoning [13]. Therefore, when emotions are studied, they are assumed to be separate phenomena that interact with human cognition to magnify, dampen, or filter it, and by implication, are not essential to it. These views can be traced back to early Western philosophy. Plato and Aristotle spoke of human desire as something that must be subordinated to reason and habituated through education. The modern popularity of Stoicism can be partially attributed to the widespread demand for an emotional education aligned with a cultural imperative to subdue, restrain, or otherwise regulate emotional expression. With a few notable exceptions [3, 8], these views have not changed significantly in the past 2000 years. I would encourage anyone who wants to dive deeper into this history to consider Andrea Scarantino's comprehensive guide to emotion theory [15].

How This Influences AI Research: A Case Study from AI Creativity

Can you have artistic creativity without emotions? Research into AI creativity would lead you to believe this is the case. Our most cutting-edge approaches to generative artificial intelligence rely on throwing increasing amounts of data into neural network architecture, implicitly treating creativity as a form of pattern matching and combinatorial exploration. This has led to impressive results. It is not my intention to discount the work of the thousands of bright researchers who have led us to an unprecedented diversity of skill automation. But is generative AI creative? That is a hotly debated question.

I will sidestep this philosophical argument for now and lean on the consensus view so as to illustrate where emotions play a role. The current consensus on creativity asserts that creative intelligence comprises the intentional agency to produce something novel that is valued by others [1, 2, 4, 6, 17]. This consensus view defines creativity as more than mere novelty. It involves intentional novelty and novelty that is congruent with our needs. Novelty itself may be defined historically (e.g., such as when a creative work is copyrighted) or psychologically (i.e., 'new to me') [2]. Higher levels of creativity surprise us.

Each of these has its origin in emotional experience. Novelty seeking is a personality trait that is frequently associated with higher creative output and is often motivated by physiological need and emotional reward. Surprise itself is considered one of our most basic emotional experiences [5], and the value we place on art often has as much to do with how we relate to it emotionally as its functional purpose or skilled craftsmanship. A lot of these considerations are largely missing from current AI research. The exclusion of emotions from AI creativity can most readily be observed in its performance, its evaluation criteria, and in current approaches to model design.

Many have noted that the dispassionate production of AI art, while frequently persuasive, has also led to a noticeable uncanniness, an abundance of 'AI slop' [9], and technical challenges such as model decay, semantic drift, or mode collapse. These are the symptoms of imposing a rationalist worldview on what is widely considered a sentimental and predominantly social form of expression.

Measures of AI creativity tend toward the exclusion of emotional experience in favor of self-contained, quantitative measures such as randomness, diversity, statistical irregularity, or novelty relative to the training dataset [12]. While human evaluation is sometimes pursued, humans are often brought in for a binary Turing-style evaluation of model output, or to evaluate second-order perceptions of AI creativity: originality, coherence, etc. More difficult to find is research incorporating the human perception of AI art as it relates to emotional experiences like surprise, catharsis, or emotional self-awareness.

This evaluation criteria, and its underlying assumptions, filters back into model design, where leading paradigms focus on feeding models more sources [7], compositionality [11], and post-training alignment [10]. All of these things certainly influence creativity, yet most of them, including alignment, tend to overlook the role of emotions as sources of information, motivation for recomposition and as foundational to the ethical values and intentions sought from alignment. So long as we continue to overlook the role of emotions in creativity, we may find ourselves trapped in a local optimum sooner rather than later.

Final Thoughts

I have spoken rather generally about current norms in AI research, rather than addressing every existing approach to this topic. There are many possible reasons as to why emotions have been kept peripheral to the study of human and machine intelligence. It may be that AI researchers have avoided emotions as a consequence of disciplinary silos, preventing them from incorporating a clearer understanding of what emotions are (especially those outside of philosophy, psychology, and neuroscience circles). It may also be that many researchers implicitly accept that emotions are unique to humans, and there is therefore no point trying to study something that doesn't exist in machines. These are possibilities worth considering. My hope is that we can begin identifying, discussing and clearing away the source of these confusions to support a more robust, humanistic approach to AI research. There is burgeoning scientific literature on the vital role of emotions in human cognition, as well as much older views that, while not the prevailing paradigm at the moment, may nonetheless be familiar to many readers. It is my intention to give this more attention in future writings.

References

[1] Anna Abraham. 2018. The Neuroscience of Creativity. Cambridge University Press.

[2] Margaret A. Boden. "The creative mind: Myths and mechanisms." (2004).

[3] Seth Duncan and Lisa Feldman Barrett. 2007. Affect is a form of cognition: A neurobiological analysis. Cognition and emotion, 21(6), 1184-1211.

[4] Mihaly Csikszentmihalyi. 1997. Creativity: the psychology of discovery and invention. HarperPerennial, New York, NY.

[5] Paul Ekman, Tim Dalgleish, and M. Power. 1999. Basic emotions. San Francisco, USA, 1.

[6] Berys Gaut. "The philosophy of creativity." Philosophy Compass 5, no. 12 (2010): 1034-1046.

[7] Alon Halevy, Peter Norvig, and Fernando Pereira. 2009. The unreasonable effectiveness of data. IEEE Intelligent Systems, 24(2), 8-12.

[8] David Hume. 2000. A treatise of human nature. Oxford University Press.

[9] B. Hoffman. 2024. First came 'spam.' Now, with A.I., we've got 'slop'. The New York Times. https://www.nytimes.com/2024/06/11/style/ai-search-slop.html

[10] Jiaming Ji, Tianyi Qiu, Boyuan Chen, Borong Zhang, Hantao Lou, Kaile Wang, Yawen Duan, Zhonghao He, Lukas Vierling, Donghai Hong, Jiayi Zhou, Zhaowei Zhang, Fanzhi Zeng, Juntao Dai, Xuehai Pan, Kwan Yee Ng, Aidan O'Gara, Hua Xu, Brian Tse, Jie Fu, Stephen McAleer, Yaodong Yang, Yizhou Wang, Song-Chun Zhu, Yike Guo, and Wen Gao. 2023. AI alignment: A comprehensive survey. arXiv preprint arXiv:2310.19852.

[11] Brenden M. Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J. Gershman. 2017.

Building machines that learn and think like people. Behavioral and brain sciences, 40, e253.
[12] Li-Chun Lu, Miri Liu, Pin Chun Lu, Yufei Tian, Shao-Hua Sun, and Nanyun Peng. 2026. Rethinking creativity evaluation: A critical analysis of existing creativity evaluations. In
Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers) (pp. 6329-6352).
[13] Jeffrey R. Huntsinger and Simone Schnall. 2013. Emotion–cognition interactions. In D. Reisberg (Ed.), The Oxford Handbook of Cognitive Psychology (pp. 571–584). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780195376746.013.0036
[14] Stuart J. Russell and Peter Norvig. 2020. Artificial intelligence: A modern approach (4th ed.). Pearson. ISBN-13: 978-0-13-461099-3
[15] Andrea Scarantino. (Ed.). 2025. Emotion theory: the Routledge comprehensive guide. Routledge.
[16] Hua Shen, Tiffany Knearem, Reshmi Ghosh, Michael Xieyang Liu, Andrés Monroy-Hernández, Tongshuang Wu, Diyi Yang, Yun Huang, Tanushree Mitra, Yang Li, and Marti Hearst. 2025. Bidirectional Human-AI Alignment: Emerging Challenges and Opportunities. In Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (pp. 1-6).
[17] Dustin Stokes. 2011. "Minimally creative thought." Metaphilosophy 42, no. 5 (2011): 658-681.
[18] Kathryn Wojtkiewicz. 2023. "How Do You Solve a Problem like DALL-E 2?." The Journal of Aesthetics and Art Criticism 81.4: 454-467.

ARTICLE

The Human AI Balance to Maintain Effective Cyber Security and Compliance for Fintech Rollouts

by Stephan Kudyba
New Jersey Institute of Technology

Keywords: financial technology, cybersecurity, compliance, risk management, digital payments

Introduction

New technologies have changed operational processes across industries, from supply chain and CRM applications to the rollout of simple Apps that expedite customer touchpoints with organizations. The financial industry is no exception. Technologies have expedited payment processes, trading and investment activities, commercial banking and lending operations, to name a few. Unfortunately, digitization can also introduce a dark side to business activities in the form of cybercrimes. A growing problem in fintech is money laundering. According to the US government, money laundering involves disguising financial assets so they can be used without detection of the illegal activity that produced them and transformed into funds with a legal source. The UN has estimated that up to 5% of global GDP is laundered annually. As a result, regulators have adopted a sound risk management approach, including policies and regulations to mitigate these activities. Difficulties arise, however, as financial organizations seek to maintain their operational efficiencies while adhering to these regulations.

Limitations of Existing Cybercrime Defenses in a Dynamic Digital World

Banking institutions and compliance agencies have adopted policies and systematic approaches to combat the growth of money laundering activities. More traditional systems have been used to identify fraudulent activities in money sources and flows, along with account-based issues, through the use of rules-based techniques. These approaches can produce “red flag” activities that indicate potential money laundering attempts. However, the growing sophistication of fraudulent activities continues to evolve, given more opportunities presented by the digital world. Evolving technologies in fintech include rapid communication capabilities, increased mobile and digitized platforms, and the expansion of new e-products and services. These elements, along with the rapid growth of mobile bank accounts and new currency alternatives, have enhanced opportunities for cybercrimes, rendering existing defenses less effective and more restrictive. Rules-based systems can overcompensate to safeguard against evolving malicious activities, and this can lead to adverse effects such as false positives or the red-flagging of an account or activity as a laundering candidate when it is not.

More simplistic, rules-based systems that produce high “red flag” rates can hinder the operations of fintech operations as they must investigate the flagged activities, which utilize increased labor and time resources. They also can restrict access to financial resources from “higher risk” developing nations that are more readily flagged.

A major setback of false positives is that legitimate organizations that require access to financial resources (e.g., loans, transactional processes) have difficulty obtaining them. This limits their ability to expand and promote economic growth, including investment in regional infrastructure, business expansion, and increased employment. This, in turn, impacts corresponding societal opportunities. This issue can be particularly detrimental to developing economies, which require access to capital and financial processing to expand.

Financial Systems, Fintech, and Economic Growth

Research has consistently supported the notion that greater access to finance by small to medium-sized enterprises (SMEs) positively impacts economic development. The impact is more prevalent in lower-income, economically developing nations. Early studies [1] and more recent works [2], [3] find that greater access to finance enables small- to medium-sized organizations to pursue projects and become more productive. This is particularly beneficial to local economies. (SMEs) can experience more difficulty remaining in business and expanding, due to their higher risk profile (e.g., new and unproven revenue), which causes greater restrictions in access to capital. Innovative startups and entrepreneurial entities, however, provide economic growth through employment opportunities and the availability of increased products and services for local economies.

Unfortunately, lower-income, economically developing nations are sometimes plagued with a “high risk profile” for other reasons that impede their access to financial resources. These lower-income, less economically developed zones may lack structured regulatory controls over existing banking institutions, increasing the likelihood of cybercrimes (e.g., money laundering) and making investors less likely to engage with them [4].

The recent expansion of technology characteristics of fintech systems shows the same positive effects on economic development. The progression of fintech (e.g., digital, e-transactions, mobile) positively impacts economic growth across various provinces in China [5], while fintech e-payments positively affect economic growth in APEC countries [6]. Additional research has shown that fintech attributes help achieve more sustainable, green economic growth [7]. These more technologically savvy systems also help alleviate poverty, reduce inequality, and improve the quality of a country’s environment. These systems enable smaller, often excluded consumers, to access financial resources (e.g., through credit facilities and alternative lending options), thus closing the exclusion gap. The diverse and broad functionality of fintech facilitates consumers’ daily consumption activities, which promotes individual development within their society. [8]

Despite consistent evidence supporting the expansion of fintech to achieve more prosperous economies and opportunities for corresponding societies, the evolution of negative characteristics (e.g., sophisticated money laundering activities) is a reality that must be managed to preserve effective financial systems. A major technological tool that can achieve this is the application of sophisticated AI-based methodologies that are able to keep pace and safeguard against the ever-changing tactics of money launderers, while keeping financial organizations compliant with updates on policies and regulations. However, the key to maintaining an edge on evolving cybercrime tactics is the human element.

The Rise of AI in Finance

Risk management remains an essential attribute in the realm of banking and finance. The digitization of the world economy has enhanced this sector through innovative technologies, and affiliated services and products. However, digitization has also introduced dynamic opportunities for cybercrime. The evolution of Artificial Intelligence methodologies has enabled financial institutions to combat this dark element. Through the combination of increased data resources and the computational power of its algorithms, this mechanism has provided an improvement over more static rules-based systems. Machine Learning and Natural Language Processing (NLP), processes both structured and unstructured data, which are relevant to financial operations [9].

However, despite the power of these methodologies, they include some limitations when considering the ever-changing character of fintech, consumers accessing them, and regulations that arise to maintain compliance. One of these entails a collaboration gap between AI implementers and business counterparts who require the systems. A near-seamless communication between business practitioners who know the intricacies of operations can help guide and provide insights to the deployment of AI models [10].

The following case study illustrates an effective approach to combating money laundering in fintech through the use of an ensemble AI approach, along with knowledge resources or operational experts in the field to guide it.

A More Effective Approach

Identifying money laundering activities involves analyzing an array of data and content. Both structured and unstructured data sources are prevalent in the network of activities involved in these fraudulent schemes. Text sources describing the origins of individuals seeking to move currency and establish accounts, along with more structured data (e.g., behavioral and operational activities), are involved as well. Effectively analyzing this content requires both predictive AI and NLP-based AI, along with guidance by experts in the field.

To fully address a complex network of activities involved in money laundering, the proper application of both machine learning and NLP, along with the guidance and input of skilled investigators, is required. Constant optimization of these approaches yields more consistent accuracy in money laundering identification, serving as an ongoing, dynamic defense against the continuous evolution of cyber tactics.

Possible results of this more advanced approach are significant. The greater accuracy of identifying authentic money laundering reduces the number of false positive instances. In other words, increased accuracy of the ensemble AI approach enables fintech organizations to function more effectively.

This more accurate AI approach addresses a major concern in the realm of fintech, which refers to risk management. By more effectively identifying bad actors in opening accounts and moving currency, financial organizations can enter a wider marketplace (e.g., regions of higher risk) where legitimate customers need financial services. More traditional ROI is a gain in market share due to more effective risk management for financial institutions. ROI is extended to the positive societal impacts from an expanding economy in developing nations due to increased access to capital.

Case Study

Onafriq, a leading Pan-African payment network, connects over 500 million mobile wallets and bank accounts across more than 40 African markets, with nearly 1,500 payment corridors linking banks and telecom partners to its infrastructure. As is the case with its fintech digital operations, Onafriq is confronted with effectively managing its operations to safeguard against malicious cyber activities.

It sought to enhance its effectiveness in combating money laundering activities, adhere to compliance standards, and implement ThetaRay, a provider of AI-based analytics to apply a "rules-free" ensemble-based AI approach. They adopted the platform, which provides data resources and full forensic investigative capabilities built on an AI foundation. This platform enables Onagfriq investigators to produce a customized solution for their situation by analyzing both SWIFT (a standardized messaging system for financial institutions to communicate transactions) and KYC (Client identification and authorization data) using an ensemble approach, with guidance from investigative experts. Figure 1 illustrates the analytic method.

This AI approach, guided by skilled operational personnel, identifies patterns of adverse behavioral and operational activities that indicate money laundering. The combination of AI technology, effective data resources, and subject matter experts enables the platform to quickly adjust to the dynamic nature of cyber threats.

Examples indicating possible fraudulent activities:

- Small repetitive payments to multiple recipients
- Financial activity inconsistent with a particular profile
- Multiple deposits below reporting thresholds

Onafrig was able to more accurately identify cyber incidents and reduce false positives. In the first cycle of alert investigations, 85% of sender alerts and 94% of receiver alerts were considered detection and investigation-worthy. In other words, the platform's red flags identified potential fraudulent activities, thereby mitigating unnecessary investigations of legitimate customers. These time savings from reduced inspection activity enable the company to allocate more resources (skilled labor) to expand its services to new customers.

This ensemble approach has produced ROI beyond simple reduction in false positives and corresponding “reduction in investigative labor hours” for the company. More accurate identification of legitimate threats enables Onafrig to expand and more effectively provide access to financial resources to developing nations. This more effective AI approach has enabled the organization to approach a wider marketplace to gain legitimate customers, producing significant market share (internal ROI) along with benefits to society (e.g., developing nations in need of capital to expand their economies).

Evolving ROI Considerations

As the deployment of digitally based initiatives increases throughout industries, organizations, and researchers need to consider not just standardized systems and traditional ROI metrics. They need to adapt to dynamic, AI-based systems that involve human expertise and consider the wider, more far-reaching value of their applications. Organizational leaders must not only consider integrating AI to enhance their operations' KPIs but also increasingly adopt strategies that balance human skills with advanced systems, and consider how their companies can gain market share and benefit consumers and society while maintaining compliance standards. The expansion of effectively functioning fintech provides access to capital to developing nations to promote economic expansion. This, in turn, benefits societies of these lower-income nations.

References

1] Matthew O. Odedokun. 1998. Financial intermediation and economic growth in developing countries, Journal of Economic Studies 25 (3): 203–224.

[2] Karthik Krishnan, Debarshi Nandy, and Manju Puri. 2015. Does Financing Spur Small Business Productivity? Evidence from a Natural Experiment, The Review of Financial Studies, Volume 28, Issue 6, Pages 1768–1809.

[3] Sanjaya Kumar Lenka and Ruchi Sharma. 2017. "Does Financial Inclusion Spur Economic Growth in India?," Journal of Developing Areas, vol. 51(3), pages 215-228.

[4] Badar un Nisa, Waqas Shair, Ameera Muzammal, and Rizwan ul Hassan. 2024. Economic Consequences of Anti-money Laundering Regulations: Investigating the Impact on GDP Per Capita across Global Economies, Bulletin of Business and Economics, Vol 13 No. 3.

[5] Na Song and Isaac Appiah-Otoo. 2022. The Impact of Fintech on Economic Growth: Evidence from China, Sustainability, 14(10), 6211.

[6] Chatchai Khiewngamdee and Ho-don Yan. 2019. The role of Fintech e-payment on APEC economic development, Journal of Physics, Volume 1324.

[7] Minahil Awais, Ayesha Afzal, Saba Firdousi, Amir Hasnaoui. 2023. Is fintech the new path to sustainable resource utilisation and economic development?, Resource Policy, Vol 81.

[8] Attila Kovari. 2024. AI for Decision Support: Balancing Accuracy, Transparency, and Trust Across Sectors, Information, Vol 15 (11).

[9] Jasmin Praful Bharadiya, Reji Kurien Thomas, and Farhan Ahmed. 2023. Rise of Artificial Intelligence in Business and Industry, Journal of Engineering Research and Reports, 25(3).

[10] Al-Alawi, Adel Ismail, and Sara Abdulrahman Al-Bassam. 2020. The Significance of Cybersecurity Systems in Helping Manage Risk in Banking and the Financial Sector, Journal of Xidan University, 14(7).

ARTICLE

Technology Transfer of Open-Source Data Infrastructure for Embodied AI: The AIRSPEED Case

by Tiantian Zhang and Shaoshan Liu
Shenzhen Institute of Artificial Intelligence and Robotics for Society

Keywords: Technology transfer, open-source data infrastructure, embodied AI, computing and society, robotics data governance, AI standardization, research commercialization

1. Introduction: Why Embodied AI Needs Data Infrastructure

Embodied AI changes the data problem in artificial intelligence. Text, images, and videos can be collected from the digital world at large scale, but robot interaction data must be produced through physical execution [1], [2]. A useful dataset may include camera streams, robot trajectories, force and tactile signals, teleoperation commands, task outcomes, and environmental context. These records are costly because they depend on hardware, operators, calibration, safety checks, and repeated trials. As a result, many embodied AI systems still rely on small demonstrations, isolated laboratory datasets, or task-specific data pipelines.

The importance of this bottleneck is social as well as technical. Data infrastructure influences who can participate in embodied AI development, whether different robot platforms can interoperate, how sensitive recordings of human activity are governed, and whether the emerging robotics ecosystem becomes a set of closed vendor silos or an extensible public-private infrastructure.

AIRSPEED was motivated by a recurring engineering problem in embodied AI data production [3, 4]. When a team changed from one robot arm to another, or from a mobile robot to a humanoid platform, much of the data pipeline had to be rebuilt. Sensor streams, control commands, coordinate frames, timestamps, task labels, and teleoperation records were often stored and processed differently. This slowed data collection, but the larger problem arose after collection: teams still had to clean, align, label, convert, version, and inspect the data before it could be used for training or evaluation. AIRSPEED was created to make this process repeatable across robot embodiments.

The project, therefore, focused on the layer between robot hardware and model training. Instead of building another robot or another task-specific dataset, AIRSPEED connects teleoperation devices, robot adapters, simulation tools, and dataset construction workflows. The goal is practical: once a team has collected real interaction data, the same platform should help turn it into training-ready datasets, expand it through controlled simulation, and reuse it across related tasks or robot platforms. This made AIRSPEED suitable for technology transfer because the value was not limited to one laboratory prototype; it lay in a workflow that other teams could adopt.

Open source played a specific role in this process. It allowed outside researchers and developers to inspect the interfaces, test the data format, and build adapters for their own devices or robots [5]. This matters in embodied AI because no single institution controls all robot embodiments, sensors, simulators, and teleoperation tools. A closed platform may work within one vendor's system, but it is less likely to become a shared workflow. AIRSPEED's open-source strategy was therefore not only a dissemination choice; it was a way to test whether a common data-production interface could attract users beyond the original development team.

This article examines AIRSPEED as a technology-transfer and computing-and-society case. It asks how a research institute can move from a technical prototype to an industrially usable infrastructure project when the core asset is not a single algorithm, but a reusable engineering workflow that affects access, governance, and standardization [3, 4]. The following sections discuss the problem AIRSPEED addressed, the platform architecture, the role of AIRS in incubation, the transition from technical validation to venture formation, and the lessons that may apply to other deep-tech transfer projects.

2. AIRS as an Incubation Platform for Technology Transfer

AIRSPEED began inside AIRS because the problem was still too early and too engineering-heavy to be solved as a conventional software startup. The team first had to understand how different robots, teleoperation devices, simulators, and data formats could be connected in one workflow. This required access to hardware, engineers, industrial partners, and testing scenarios. AIRS provided that environment before the project was ready to become a market-facing venture. This role is consistent with the broader view that universities and research institutes increasingly support a "third mission" beyond research and education: translating research outputs into economic and social impact through technology transfer and commercialization [7].

AIRS is one of Shenzhen's major public research institutions, established to connect frontier research with industrial deployment in artificial intelligence and robotics. Affiliated with The Chinese University of Hong Kong, Shenzhen, AIRS operates as a new type of research institute that combines academic research, engineering development, industry collaboration, venture incubation, and technology transfer. This institutional model is particularly relevant for embodied AI, where successful commercialization depends not only on algorithms but also on hardware access, engineering integration, industrial partners, pilot scenarios, and standardization pathways.

From a computing-and-society perspective, this incubation role matters because infrastructure decisions made in research settings can shape future industrial norms. If data formats, adapters, and governance practices are defined only by dominant vendors, smaller laboratories, local integrators, and public-interest applications may face higher entry barriers. Public research institutes can partially counter this tendency by producing shared technical assets and open interfaces before markets fully consolidate.

Within this environment, AIRSPEED was developed as an open-source data production platform for embodied AI. Its technical objective is to reduce the cost and fragmentation of embodied AI data production through a modular architecture that connects teleoperation devices, heterogeneous robots, simulation environments, and dataset construction pipelines. Rather than building another vertically integrated robot system, AIRSPEED focuses on the missing horizontal layer: the data infrastructure that allows different robots, sensors, operators, and simulation tools to contribute to reusable embodied AI datasets.

AIRS helped AIRSPEED in three concrete ways. First, it gave the team an engineering environment in which data alignment, simulation interfaces, and robot integration could be tested before the platform had to satisfy paying customers. Second, it exposed the team to industrial users who cared less about the elegance of the software architecture and more about whether the system could reduce repeated integration work. Third, it helped the team translate the prototype into a commercialization plan: which parts should remain open source, which services could become enterprise offerings, and what kind of investors or partners would understand an infrastructure-style project.

AIRSPEED followed an internal incubation pathway. Under this model, the technical team develops the core technology within the institute, validates the system through engineering and pilot projects, and then establishes an independent company once the technology and market opportunity become sufficiently clear. This model reduces the early risk of commercialization: the project can rely on institutional resources during the uncertain prototype stage, while later gaining the autonomy needed for venture financing, customer delivery, and product scaling.

This incubation pathway is important because embodied AI infrastructure is difficult to commercialize through a simple licensing model. Much of the value lies in tacit engineering knowledge, interface adaptation, customer feedback, dataset workflows, and ecosystem trust. AIRSPEED's transfer, therefore, required more than moving intellectual property from a laboratory to a company. It required moving a team, a technical architecture, partner relationships, and a market thesis together.

The AIRSPEED case suggests that infrastructure-oriented AI projects often need an institutional bridge before they can become companies. The early asset is rarely just source code. It also includes interface knowledge, hardware adaptation experience, user feedback, testing environments, and trust from partners. AIRSPEED's transfer, therefore, involved moving a working engineering practice, not simply transferring a piece of intellectual property.

3. The AIRSPEED Platform: Architecture for Scalable Data Production

AIRSPEED treats embodied AI data production as an engineering workflow rather than a one-time collection task [6]. A demonstration collected on one robot becomes more valuable if it can be synchronized, checked, converted, replayed, expanded in simulation, and reused in later training cycles. The platform was therefore designed around the steps that make data reusable: device connection, robot adaptation, simulation linkage, data cleaning, dataset construction, and version control.

Although these elements are technical, they also encode institutional choices. A common adapter layer can make the ecosystem less dependent on a single robot vendor. A shared dataset construction workflow can make evaluation more transparent. Private deployment options can reduce pressure to centralize sensitive human-machine interaction data. These design choices determine not only platform performance but also the social shape of embodied AI adoption.

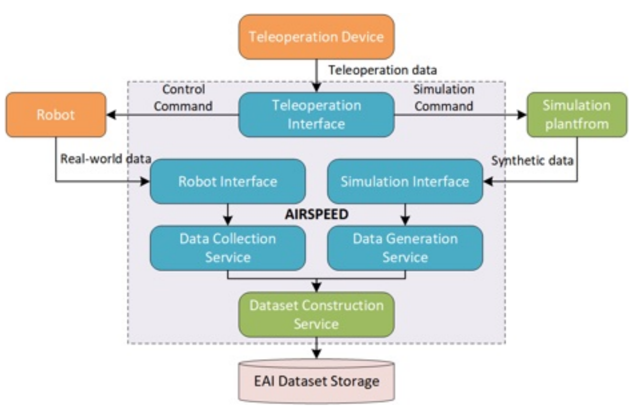

Figure 1. AIRSPEED architecture design.

As illustrated in Figure 1, the platform is organized around three interfaces and three services. The three interfaces connect AIRSPEED to the key sources of embodied AI data: teleoperation devices, robot embodiments, and simulation environments. The teleoperation interface supports devices such as VR headsets, motion-capture systems, grippers, and other control terminals. The robot interface allows heterogeneous platforms, including robotic arms, mobile robots, quadrupeds, and humanoids, to be integrated through adapters. The simulation interface connects the platform to environments such as NVIDIA Isaac Sim and MuJoCo, allowing real-world demonstrations to be expanded through synthetic variation and simulation-based data generation.

As illustrated in Figure 2, on top of these interfaces, AIRSPEED provides three core services: data collection, data generation, and dataset construction. The data collection service receives real-world interaction data and performs synchronization, timestamp alignment, and initial cleaning. The data generation service uses simulation and data-derivation methods to expand limited real-world demonstrations into more diverse samples. The dataset construction service converts raw multimodal signals into structured datasets through cleaning, annotation, format conversion, quality control, and version management.

The key architectural contribution of AIRSPEED is the integration of these components into a closed-loop data production process. Real-world demonstrations serve as high-value seed data. These data are standardized and converted into training-ready datasets. The datasets are then used for model training, simulation rollout, and further data generation.

Newly generated data are merged with real-world data, creating larger and more diverse datasets for the next training cycle. This mechanism forms a data flywheel in which each round improves dataset scale, diversity, and reusability while reducing the marginal cost of additional data production.

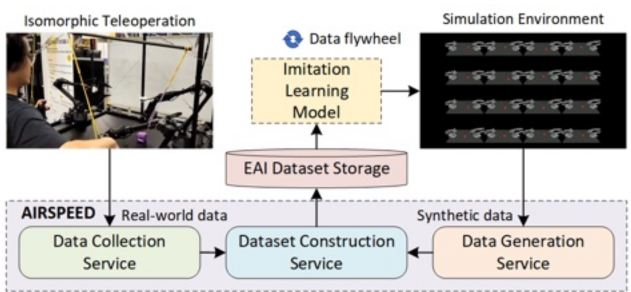

Figure 2. Services provided by AIRSPEED.

4. AIRSPEED's Incubation-Stage Development

During incubation, AIRSPEED's main task was to turn a research idea into a repeatable engineering capability. The team did not begin with a finished product. It began with a practical problem: embodied AI data pipelines were repeatedly rebuilt whenever the robot, teleoperation device, simulator, or task changed. The incubation stage therefore focused on accumulating reusable technical assets rather than delivering one-off demonstrations. These assets included device interfaces, robot adapters, simulation connectors, data-alignment methods, dataset construction workflows, and early reference examples.

The first capability of AIRSPEED was hardware and device compatibility. AIRSPEED had to connect different sources of embodied AI data, including teleoperation devices, robot embodiments, and simulation environments. This required work that is often invisible in a research prototype: driver integration, control-signal conversion, timestamp alignment, coordinate-frame mapping, sensor synchronization, and interface testing. The value of this work was cumulative. Each new adapter made the next integration easier and helped turn scattered engineering fixes into a reusable compatibility layer.

The second capability was a modular data-production pipeline. AIRSPEED separated teleoperation interfaces, robot adapters, simulation connectors, data collection, data generation, and dataset construction into different parts of the platform. This separation mattered because customer scenarios rarely use the same hardware stack. A modular architecture allowed the team to change one part of the workflow without rebuilding the whole system.

The third capability was simulation-assisted data expansion. Real robot demonstrations remained the anchor, but they could be extended through simulation and data-derivation methods. This was not intended to replace real data. Its purpose was to increase variation, test related scenarios, and reduce the marginal cost of producing additional training samples.

The fourth capability was dataset construction. AIRSPEED treated raw robot logs as unfinished material. Before such data can support model training, it must be cleaned, synchronized, labeled, formatted, checked, versioned, and made searchable. By building workflows around these steps, AIRSPEED moved beyond data capture and became a data-production platform.

The fifth capability was interface discipline. Because AIRSPEED was designed as an open-source platform, the team had to make interfaces understandable beyond the original developers. This required cleaner SDK boundaries, reference workflows, documentation, and dataset templates. In this sense, open source was not only a release strategy; it forced the project to make its engineering assumptions explicit.

By the end of incubation, AIRSPEED was no longer only a software prototype. The team had produced concrete components that could be used again: robot adapters, teleoperation interfaces, simulation connectors, timestamp-alignment tools, data-cleaning scripts, dataset templates, and reference workflows. This changed the nature of the project. Its value was no longer tied to one demonstration, one robot, or one dataset. It came from the fact that the same components could shorten the next integration cycle. The key output of incubation was therefore practical: AIRSPEED gave the team a way to connect a new device, collect interaction data, clean and synchronize the records, and convert them into a dataset without rebuilding the whole pipeline. Whether this repeatable workflow could support customers and commercialization is the question addressed in the next section.

5. From Engineering Capability to Commercial Validation

The engineering work described in Section 4 did not automatically create a business case. AIRSPEED still had to show that these components solved problems customers faced repeatedly. The key question was not only whether the platform could connect robots, devices, simulators, and datasets, but whether it could reduce the time spent on repeated integration, data cleaning, synchronization, and format conversion. If a pilot project only produced a custom script, it remained a service project. If it produced a robot adapter, dataset template, deployment script, or quality-control workflow that could be reused in the next project, it strengthened the platform. In platform ecosystems, such reusable complements can increase the value of the core platform because external users and partners help extend its applicability across use cases [11].

The first question was whether the problem would appear again. AIRSPEED had to show that data production was not a one-time integration task. Each new robot, sensor setup, teleoperation device, simulator, or task could require new alignment, cleaning, conversion, and dataset-building work. This made the platform useful only if it could reduce the work needed for the next project, not just finish the current one.

The second question was what each pilot left behind. A pilot was valuable only when it produced something that could be reused: a robot adapter, a teleoperation interface, a dataset template, a quality-checking script, or a deployment workflow. If the output was only a customer-specific script, the project moved AIRSPEED toward services. If the output became part of the common platform, the project made AIRSPEED stronger.

The third question was what should remain open and what customers would pay for. The open-source layer should allow users to inspect the code, test the interfaces, and decide whether AIRSPEED can fit their own robots and data workflows [9]. It should not be a marketing wrapper around a closed product; it should contain interfaces, data schemas, adapters, and reference workflows that outside users can actually evaluate. The commercial layer should focus on work that customers are less likely to handle alone: private deployment, integration with their own hardware, large-scale simulation runs, dataset management, technical support, governance support, and workflow customization [10]. Because customer data often contains proprietary and personally sensitive information, the commercial model should prioritize integration and trusted deployment rather than data ownership.

The fourth question was what evidence could convince customers and partners. A vision for embodied AI data infrastructure was not enough. AIRSPEED needed working integrations, tested data workflows, pilot feedback, and examples showing that the same adapter, template, or cleaning process could be used again. These were stronger signals than a research demo because they showed that AIRSPEED could reduce repeated engineering work.

The core lesson is that a research prototype becomes commercially credible when it helps users avoid doing the same engineering work again and again. Section 4 showed what AIRSPEED built during incubation. This section explained why those components mattered outside the laboratory: they shortened repeated integration work, turned pilots into reusable modules, gave developers an open interface to test, and created service opportunities around deployment, data management, and technical support. At that point, AIRSPEED was no longer only a prototype. It became a practical workflow that other teams could use.

6. Governance and Risk Mitigation: Keeping AIRSPEED Scalable and Trustworthy

AIRSPEED's main risks are not only execution risks. They are also governance risks. The project can fail not because the idea is wrong, but because it becomes too dependent on simulation, too customized for early customers, too difficult for outside developers to adopt, or too weak in handling sensitive robot interaction data.

The first risk is over-reliance on simulation. AIRSPEED should treat real-world demonstrations as the calibration anchor and simulation as a controlled expansion tool. A simulation-generated dataset should be linked to real demonstrations, checked against real-robot execution when possible, and assigned a quality score before it is used for training. For high-precision manipulation, real data should remain the benchmark.

The second risk is becoming a customization service rather than a platform. Early customers will naturally ask for special adapters, special workflows, and special data formats. AIRSPEED should accept such work only when it strengthens the common platform. Each customer project should produce at least one reusable asset: a robot adapter, a device interface, a dataset template, a quality-control workflow, or a deployment script.

The third risk is weak adoption. Open source can reduce this risk, but only if the platform is easy to inspect, test, and extend. AIRSPEED should therefore prioritize reference examples, clean SDK interfaces, documentation, and small reproducible workflows. Official standards may take time, so the practical objective is to let developers use AIRSPEED-compatible formats before formal standardization is complete.

The fourth risk is customer trust. Embodied AI data can include sensitive videos, task records, operator behavior, and proprietary robot information. AIRSPEED should therefore support private-cloud or on-premises deployment, role-based access control, audit logs, encrypted storage, and clear data-governance rules. The business model should not depend on taking ownership of customer data.

The fifth risk is ecosystem capture. A nominally open platform can still become de facto closed if key adapters, formats, or validation tools are controlled by a small set of vendors. AIRSPEED should keep core schemas, reference implementations, and compatibility tests open enough for independent users to build against them. This would make open source a practical interoperability mechanism rather than only a release label.

In short, AIRSPEED should keep real data as the source of truth, use simulation to scale carefully, use open source to spread interfaces, use standards to build legitimacy, treat data governance as part of the architecture, and use paid pilots to test whether the platform solves repeatable customer problems.

7. Summary: Success Factors and Forward Outlook

The AIRSPEED case shows that technology transfer in embodied AI is not simply a handoff from laboratory to company. It is also a process through which choices about data formats, open interfaces, privacy-preserving deployment, and platform governance can become embedded in emerging AI infrastructure. The project had to show that the same data workflow could work across different robots, teleoperation devices, simulators, and task settings. AIRS played an important role by giving the team time to test adapters, debug data alignment, work with industrial users, and decide which parts of the system should be open source and which parts could support paid deployment.

The broader lesson is that socially meaningful deep-tech transfer should begin with a problem that users face more than once. AIRSPEED was valuable because teams repeatedly faced the same practical work: connecting new hardware, synchronizing sensor streams, converting control records, cleaning robot logs, and preparing datasets for training. Other technology-transfer projects can ask a similar question: does the research output help many users avoid rebuilding the same toolchain, and does it make the resulting ecosystem more open, governable, and reusable?

Looking forward, AIRSPEED's challenge is to prove that it can support more projects without becoming a custom-service company [8]. This will depend on practical progress: adding more robot adapters, improving multimodal synchronization, making dataset construction easier to inspect, supporting private deployment, and turning repeated customer requests into standard software modules. If AIRSPEED can do this, its importance will go beyond one platform. It will offer a concrete example of how a public research institute can turn open-source AI infrastructure into an industrial tool while preserving openness, interoperability, and governance capacity. For computing and society, that is the central significance: the infrastructure used to build embodied AI will shape not only technical performance, but also access, trust, and institutional power in the next stage of AI deployment.

References

[1] Fan Wang and Shaoshan Liu. 2025. "Putting the Smarts into Robot Bodies," Communications of the ACM, vol. 68, no. 3, pp. 6–8, doi: 10.1145/3703761.

[2] Shuang Wu, Bo Yu, Shaoshan Liu, and Yuhao Zhu. "Autonomy 2.0: The Quest for Economies of Scale," Communications of the ACM, vol. 68, no. 4, pp. 28–32, Apr. 2025, doi: 10.1145/3708012.

[3] Xuan Xia, Xianqiao Tong, Bo Yu, Jialin Jiao, Xinmin Ding, Xing He, Hongjun Zhou, Haoran Tong, Yufei Lin, Tongyi Shen, Ning Ding, and Shaoshan Liu. 2026. "AIRSPEED: An Open Source Data Production Platform for Embodied Artificial Intelligence," ACM Transactions on Cyber-Physical Systems, doi: 10.1145/3806827.

[4] Shenzhen Institute of Artificial Intelligence and Robotics for Society, "AIRSPEED: An open-source universal data production platform for embodied artificial intelligence." https://airs.cuhk.edu.cn/en/airspeed

[5] Shaoshan Liu. 2024. "Establishing standards for embodied AI," Communications of the ACM. https://cacm.acm.org/blogcacm/establishing-standards-for-embodied-ai/

[6] Shaoshan Liu. 2024. "The value of data in embodied artificial intelligence," Communications of the ACM. https://cacm.acm.org/blogcacm/the-value-of-data-in-embodied-artificial-intelligence/

[7] Mark Huson and Randall Morck. 2024. "Technology Transfer: The Rise of the Entrepreneurial University," The School of Public Policy Publications, University of Calgary, 2024.

[8] Saheed A. Gbadegeshin, Anas Al Natsheh, Kawtar Ghafel, Omar Mohammed, Ashten Koskela, Antti Rimpiläinen, Joonas Tikkanen, and Antti Kuoppala. 2022. "Overcoming the Valley of Death: A New Model for High Technology Startups," Sustainable Futures, vol. 4, 2022.

[9] Michael J. Karels. 2003. "Commercializing Open Source Software," ACM Queue, vol. 1, no. 5.

[10] Tony Bailetti. 2009. "How Open Source Strengthens Business Models," Technology Innovation Management Review.

[11] Marco Ceccagnoli, Chris Forman, Peng Huang, and D. J. Wu. 2012. "Cocreation of Value in a Platform Ecosystem: The Case of Enterprise Software," MIS Quarterly, vol. 36, no. 1, pp. 263–290.

ARTICLE

Catalyzing AI Governance and Democratic Resilience in Africa

by Chinasa T. Okolo
Technecultura

Keywords: AI Governance, Africa, Digital Democracy, Elections, Disinformation

An Introduction to AI, Elections, and the Future of Democracy in Africa

The democratic evolution of governments across Africa took prominence after a series of military dictatorships following the shift from colonial rule. This era of post-independence democratization efforts, often referred to as the "third wave of democratization" [25], began in the early 90s, with countries such as Benin, Guinea-Bissau, and Zambia transitioning to democratic rule. While organizations like the African Union, initially established as the Organization for African Unity in 1963, have been essential in promoting democratic principles, a growing shift towards authoritarian rule across the continent has resulted in democratic backsliding in countries such as Ethiopia, Côte D'Ivoire, Guinea, the Central African Republic, and Mali [2]. While the COVID-19 pandemic has accelerated this democratic backsliding [22], African countries face enduring issues such as corruption, economic instability, ethnic rivalry, religious conflict, and weak governance institutions that hinder democratic maturity [4, 7]. As political instability, conflict, and foreign intervention undermine the quality of governance, AI will introduce new complexities to these challenges. Issues with access to electrification, a lack of computing infrastructure, and a low concentration of skilled AI talent hinder AI development across the continent. Additionally, as African governments increasingly leverage AI in governance and democratic processes such as elections, AI tools could limit political awareness, intensify voter suppression, and spread disinformation.

Digital technologies, enhanced by widespread mobile penetration and expanding internet connectivity across the continent, have democratized access to information, enhanced government services, and reshaped political engagement in African countries. To improve public service delivery and transparency, several African governments have introduced e-governance initiatives over the past two decades [29]. These programs include efforts such as Ghana's development of a National Information Clearing House to improve access to public and private sector data, Kenya's development of the Electronic Graft Management (EGM) pilot project that provided a platform to investigate corruption and enabled citizens to report crimes, and Rwanda's development of a platform to enable consumers to prepay their utilities through SMS to Electrogaz, the national water and electricity company, to prevent fraud and ease bill collection in rural areas. Over the past few years, a number of African countries have started to leverage AI in their digital government services. For example, Togo's NOVISSI platform distributes emergency cash transfers via mobile money using machine learning techniques [30]. African governments such as South Africa, Eswatini, Madagascar, and Nigeria are leveraging AI to assist with voter registration, authentication, management, and engagement, as well as detecting cyber threats and aiding oversight and decision-making [11]. While AI-enabled digital governance offers many opportunities to streamline government processes, reduce corruption, and bolster public trust, these tools also pose many risks. Integrating AI into government systems raises concerns about heightened cybersecurity vulnerabilities, data privacy violations, and inadequate regulatory frameworks. Furthermore, these tools undermine democratic functions, as digital platforms are increasingly used for disinformation campaigns, political manipulation, and social control. Additionally, African countries are disproportionately adopting surveillance technologies in ways that violate civil liberties, especially during elections or periods of politically tense conflicts [24].

The adoption of AI has significantly shaped political systems, gradually increasing influence on aspects such as electoral integrity and administration [27]. With the integration of these technologies into electoral tools rising, risks associated with algorithmic biases, such as facial recognition and demographic predictions, can seriously impact voter equality. Moreover, AI poses new challenges to fair political competition. Given that citizen data is widely accessible in some countries, political candidates and parties can leverage insights from AI systems to engage in political manipulation through microtargeting, AI-generated propaganda, and the spread of disinformation. AI can shift power dynamics within political systems, favoring well-resourced candidates and parties who can afford sophisticated AI-driven campaign strategies [14]. These shifting power dynamics have profound implications for how emerging independent candidates can participate equitably in an AI-enhanced political landscape. Transparency, accountability, and democratic integrity are vital to ensuring that AI is ethically integrated and deployed in political systems. Given the loopholes AI has introduced across sectors, there is a need for regulatory frameworks tailored to African contexts that address data protection, digital rights, and oversight mechanisms to prevent misuse, particularly in elections and other democratic processes. AI could shape the future principles of African democracies, from reinforcing representation and accountability to exacerbating inequality and authoritarianism. However, if leveraged responsibly, AI could help strengthen democratic norms, prevent the undermining of civic freedoms, and discourage governments from shifting to autocracies.

AI's Influence on the Digital Landscape in Africa

Africa's digital economy is driven by widespread mobile adoption, which primarily emerged due to limited traditional infrastructure and access to devices such as laptops and desktop computers

[15]. Mobile technologies have enabled digital participation across the continent, allowing users to access financial, healthcare, and government services. However, there are limitations of mobile dependence, given issues with data affordability and disparities in mobile and internet access between urban and rural areas [8]. While mobile phones may provide more viable ways to expand access to AI technologies, efforts to increase AI development are hampered by uneven data availability and low data quality. These efforts are further compounded by infrastructure issues and limited access to talent who can build, deploy, and maintain AI systems. While tech giants such as Google, Meta, and Huawei have accelerated digital transformation in Africa, their business practices significantly affect local governance, political engagement, and civil liberties across the continent. Dependency on platforms developed by these foreign companies has severe implications for public information and civic engagement and could also seriously affect digital sovereignty and political autonomy in African countries. Social media platforms, such as Facebook, X (Twitter), and YouTube, continue to impact elections across the continent by enabling disinformation and voter manipulation [16], a trend that the introduction of generative AI will only amplify.

Africa is known for its linguistic, ethnic, and cultural diversity, which play a prominent role in shaping the design and implementation of AI [9]. Since prominent AI models heavily rely on the Internet, estimated to be around 60% English, as a source for training data, there are significant challenges in designing AI that recognizes Africa's diversity. These factors affect the efficacy of tools developed to improve language processing across numerous African languages and dialects, as well as tools that can generate context-specific and accurate information about African societies [20]. As AI spreads across the continent, urban-rural differences in access to digital tools can exacerbate gaps in the ability of communities to design, develop, and build these systems, creating regional power imbalances. Thus, promoting community-driven AI solutions tailored to local contexts and societal needs will be essential. More broadly, there are challenges in ensuring African representation within the AI development cycle, from data collection to deployment. Given that researchers in many countries rely on imported AI models and datasets to conduct their work, countries must promote and support African-led AI research. Biases in AI, resulting from a lack of representation, may inadvertently reinforce ethnic, linguistic, or religious divides in African societies. This could potentially lead to inaccurate representation of particular demographic groups, escalated surveillance of marginalized communities, and algorithmic exclusion in government services, social programs, and more. As AI is also leveraged in sensitive contexts such as social services, healthcare, and election management, greater efforts will be needed to integrate cultural diversity into these tools to ensure they capture the unique sociopolitical contexts in which they will inevitably be deployed. This also raises a need for ethically guided AI policies that address these risks while supporting inclusive democratic participation.

The Impact of AI on Electoral Processes in Africa

AI technologies are increasingly used in election monitoring, management, and oversight. Over the past years, researchers and governments have leveraged these tools for real-time election monitoring, including anomaly detection in voting patterns, incident tracking, and fraud detection [19]. AI has also had a significant impact on news dissemination and the delivery of voter information. African newsrooms steadily leverage AI to analyze data, gain audience insights, and create consumer-focused content [13]. However, there are limitations in the ability of AI to deliver timely and accurate information, and efforts to integrate fact-checking can be hampered by the rapid diffusion of information, especially as it moves to offline forms of media like newspapers or radio. Additionally, there are challenges of biased and opaque algorithms embedded within social media platforms that may inadvertently suppress or amplify certain types of news, impacting voter perception and awareness. In African countries, where elections are often plagued by high rates of disinformation, social unrest, internet blackouts, and ballot tampering, AI has strong potential to support monitoring efforts by civil society and international organizations, promoting transparency and credibility in electoral processes. Given the high rate of electoral fraud in African countries [12, 28], AI-enabled automated tools could also be leveraged for post-election audits, ensuring that votes are accurately counted and recorded. While AI tools could significantly refine electoral processes, governments must develop resources to enhance citizen education on civic engagement with AI in tandem with their AI integration efforts. Voters should be aware of AI tools and the potential for AI chatbots, virtual assistants, and online resources to provide misinformation on issues, candidate platforms, and voting procedures. While AI can make voter education materials available in multiple languages and formats to reach diverse communities, prominent AI translation and speech recognition tools are primarily trained on English and other Western languages and lack a solid grasp of low-resource languages [21], particularly those indigenous to African countries.

AI poses many risks to democratic integrity, including the high potential for voter manipulation, the amplification of disinformation, and privacy concerns stemming from data exploitation. AI-driven microtargeting in political campaigns through personalized ads, emails, and messages can manipulate voter perceptions and decisions [26]. Political campaigns can use AI to identify and exploit ideological vulnerabilities in specific voter groups, influencing elections through manipulative tactics. This potential for exploitation raises several ethical implications of using AI to manipulate voter sentiment and clarifies how persuasion differs from undue influence. A significant amount of research has shown that the introduction of generative AI tools has increased the production and dissemination of disinformation and “deepfake” content [5, 31]. This content can target political candidates or parties, misleading voters and negatively affecting electoral outcomes. Social media algorithms often amplify false information, creating echo chambers in which consumers' beliefs are reinforced, thereby impacting voter trust and influencing public opinion during elections [23]. As campaigns increasingly rely on personal data for targeting, there are elevated privacy risks associated with AI-driven voter profiling. Issues of consent and transparency are prominent, given that consumers are rarely informed of how their data is collected, processed, and analyzed by political parties and third-party providers [3]. Without clear oversight, this abuse of citizen data poses negative consequences for democratic integrity. AI also poses threats to election integrity, and AI-enabled automated tools can be used to hack digital voting systems and manipulate electronic vote counting, undermining trust in election outcomes [10]. This highlights the need for governments to strengthen cybersecurity measures and implement robust oversight mechanisms to counter AI-driven election tampering, alongside efforts to combat the spread of disinformation.

The rapid adoption of predictive and generative AI tools has also impacted African electoral processes. While generative AI applications emerged in part with the release of ChatGPT in November 2022, the impact of these tools on elections is still evolving. It is estimated that 36 African countries will hold general or presidential elections from 2023 to 2025, and AI use has amplified the spread of disinformation [17]. Countries like Nigeria, Senegal, and South Africa have seen a proliferation of AI-generated misinformation, including images of fake endorsements from domestic and international celebrities, manipulated audio of political candidates planning to rig elections, and AI-translated videos of speeches by politicians. However, there has also been a rise in efforts by fact-checking organizations to combat the deluge of AI-generated content intended to mislead voters. These include organizations like Africa Check, which operates across South Africa, Kenya, Nigeria, and Senegal. Other organizations are PesaCheck in Kenya, Dubawa in Nigeria, Zim Fact in Zimbabwe, Moz Check in Mozambique, 211 Check in South Sudan, Beam Reports in Sudan, FactSpace West Africa in Ghana, and Congo Check in the Democratic Republic of Congo. Governments have also been leveraging AI to combat systemic issues, with countries like Nigeria, Kenya, South Africa, and Eswatini incorporating these technologies to refine biometric voter registration and voter accreditation [1]. Civil society organizations also leverage AI to improve the information environment during elections. For example, the MAPEMA consortium in Kenya has identified over half a million negative Facebook posts and developed monitoring solutions to mitigate hate speech and disinformation in online political environments [6]. As AI compounds in influence over the coming years, it will be crucial for education initiatives to scale across African countries to progress media and digital literacy for citizens of all backgrounds.

Towards a Democratic AI Future in Africa

While AI promises many benefits to African countries, these tools are fraught with numerous risks. To address concerns about the intensifying adoption of AI technologies across Africa, governments must promote good governance by enabling robust AI governance. With at least 25 out of 55 AU Member States having a policy, strategy, plan, or guideline on AI, there is a need for stronger regional and continental-wide collaboration on AI governance frameworks. Cross-sector collaboration among African governments, the private sector, civil society, and international bodies will be essential in shaping data and AI governance for democratic outcomes [18]. The African Union and Regional Economic Communities (RECs) like the Economic Community of West African States (ECOWAS), Southern African Development Community (SADC), and Arab Maghreb Union (AMU), amongst others, should play a role in leading cooperation to set common standards for AI in elections, public services, and political campaigns. These standards should be leveraged to develop regulatory measures that prevent AI misuse by integrating robust requirements for data protection, transparency, and accountability in AI systems. There is also a need to establish independent regulatory bodies and ombudsman offices to implement regulations and ensure accountability. The aforementioned institutions will be helpful in developing oversight mechanisms to monitor AI applications in political contexts, building frameworks for redress in cases of algorithmic discrimination or harm, and enacting enforcement measures, such as penalties for non-compliance.

Furthermore, governments must prioritize AI upskilling and education to expand opportunities for African communities to contribute to AI development and to benefit equitably from these technologies. Given Africa's low concentration of AI researchers and developers, it will be necessary for countries to focus on building local expertise in various aspects of AI, from AI ethics and governance to technical skills in machine learning and data science. Large tech companies like Google, Microsoft, and NVIDIA have launched online training programs and initiatives, such as startup incubators, to support the budding ecosystem on the continent. To enhance academic AI education and research, African governments must invest in supporting accessible AI education, particularly for women, rural populations, and underrepresented groups. Along with initiatives that promote entrepreneurship and small business development, accessible AI education can help create high-quality job opportunities that support local economies and reduce youth unemployment. African governments, private-sector companies, and multilateral institutions should develop partnerships to increase funding for AI research tailored to African contexts. As AI continues to scale across the continent, there is a need for interdisciplinary research that considers the social, economic, and cultural impacts of AI in African contexts. This research can then lead to AI innovations that are sensitive to African realities and knowledgeable of the diverse array of cultures, languages, and values. While a number of foreign tech companies have established AI research labs in countries like Ghana, Kenya, and South Africa, there are opportunities to develop Africa-focused AI research centers and labs to promote knowledge exchange and thought leadership in AI ethics, governance, and democratic integrity.

Institutional challenges within many African governments, including limited technical expertise, a lack of resources to develop and sustain AI initiatives, and bureaucratic inertia, which impacts the ability of governments to adopt novel practices to facilitate change, pose many roadblocks to regulating AI. Across the continent, there is a need to develop ethical guidelines and best practices to ensure AI tools are implemented in ways that support rather than undermine good governance. Principles critical to public policy in Africa, such as transparency, accountability, inclusivity, and public trust, will also be essential to enacting robust AI governance measures. Additionally, participatory approaches that involve a diverse array of stakeholders in developing AI policies will be necessary to ensure that AI governance reflects African values and priorities. Civil society actors will be crucial in shaping culturally relevant AI policies and ethical guidelines that prioritize citizen welfare and promote democratic integrity. With the high risk of AI being used to entrench political power, suppress dissidents, and manipulate electoral outcomes across the continent, African governments must progress efforts to build sufficient leadership and technical capacity to support democratically centered AI development and governance. These efforts could include leadership development programs that equip African policymakers with the skills and knowledge to develop regulations for AI technologies and secondment programs that augment government agencies with expertise from professionals working in AI, human rights, data privacy, and economic development. If adopted responsibly and inclusively, AI technologies could support African leaders in confronting governance challenges. However, there must be systematized efforts to counter the risks of these technologies while working towards implementing comprehensive frameworks that promote ethical AI innovation and functional regulations that enable responsible implementation.

References

[1] Dennis Amata. 2024. Exploring the Opportunities and Risks of AI in African Democracy. Dataphyte. https://www.dataphyte.com/latest-reports/exploring-the-opportunities-and-risks-of-ai-in-african-democracy/

[2] Nicole Beardsworth, Hangala Siachiwena, and Sishuwa Sishuwa. 2022. Autocratisation, electoral politics and the limits of incumbency in African democracies. Journal of Eastern African Studies, 16(4), 515-535.

[3] Colin J. Bennett and David Lyon. 2019. Data-driven elections: implications and challenges for democratic societies. Internet policy review, 8(4).

[4] John Campbell and Nolan Quinn. 2021. What's Happening to Democracy in Africa? Council on Foreign Relations. https://www.cfr.org/article/whats-happening-democracy-africa

[5] Chun Chu-Ke and Yujie Dong. 2024. Misinformation and literacies in the era of generative artificial intelligence: A brief overview and a call for future research. Emerging Media, 27523543241240285.

[6] Code for Africa. 2023. Unmasking hate speech in Kenyan elections with AI and collaboration. https://medium.com/code-for-africa/unmasking-hate-speech-in-kenyan-elections-with-ai-and-collaboration-576e37d4ccb5

[7] Dengiyefa Angalapu. 2023. Africa's democratic journey: A tale of progress, setbacks, and hopes. Centre for Democracy and Development. https://www.cddwestafrica.org/blog/africa-s-democratic-journey-a-tale-of-progress-setbacks-and-hopes

[8] GSMA. 2023. The Mobile Economy Sub-Saharan Africa 2024. GSM Association.

https://www.gsma.com/solutions-and-impact/connectivity-for-good/mobile-economy/sub-saharan-africa/
[9] Arthur Gwagwa, Erika Kraemer-Mbula, Nagla Rizk, Isaac Rutenberg, Jeremy De Beer. 2020. Artificial intelligence (AI) deployments in Africa: benefits, challenges and policy dimensions. The African Journal of Information and Communication, 26, 1-28.
[10] Ben Hawes, Dame Wendy Hall, and Matt Ryan. 2023. Can artificial intelligence be used to undermine elections?. Web Science Trust. DOI:10.5258/SOTON/WSI-WP005.
[11] Samson Itodo. 2024. Artificial Intelligence and the integrity of African elections. International IDEA. https://www.idea.int/news/artificial-intelligence-and-integrity-african-elections
[12] Obaridomuru Emmanuel Lale. 2014. Democracy and electoral fraud in Africa: A comparative study of presidential elections in Ghana and Nigeria from 2000 to 2012 (Master's thesis, Eastern Mediterranean University (EMU)-Doğu Akdeniz Üniversitesi (DAÜ)).
[13] International Media Support (IMS). 2022. How African newsrooms are using AI to analyse data and produce good journalism. https://www.mediasupport.org/blogpost/how-african-newsrooms-are-using-ai-to-analyse-data-and-produce-good-journalism/
[14] Katja Muñoz. 2023. The Transformative Role of AI in Reshaping Electoral Politics. (DGAP Memo, 4). Berlin: Forschungsinstitut der Deutschen Gesellschaft für Auswärtige Politik e.V. https://dgap.org/en/research/publications/transformative-role-ai-reshaping-electoral-politics
[15] Bitange Ndemo and Tim Weiss. 2017. Making sense of Africa's emerging digital transformation and its many futures. Africa Journal of Management, 3(3-4), 328-347.
[16] Martin N. Ndlela. 2020. Social media algorithms, bots and elections in Africa. Social media and elections in Africa, Volume 1: Theoretical perspectives and election campaigns, 13-37.
[17] Chinasa T. Okolo. 2024. African Democracy in the Era of Generative Disinformation: Challenges and Countermeasures against AI-Generated Propaganda. Building a Just AI Ecosystem in Africa Conference. Research ICT Africa.
[18] Chinasa T. Okolo. 2024. Operationalizing Comprehensive Data Governance in Africa. African Journal of Sustainable Development (AJSD): Redefining African Futures: The State, Resilience, and Pathways to Progress. 2024.
[19] Movine Omondi, Adetayo A. Adesanya, George Otieno, Moses Kioko, and Eweka Y. Ogieomo. 2024. Elections in Africa: Trends, Diversity, and Emerging Technologies. Co-creation Hub. https://govote.ng/wp-content/uploads/2024/09/Election-in-Africa-I.pdf
[20] Sandeep Ravindran. 2023. AI often mangles African languages. Local scientists and volunteers are taking it back to school. Science. https://www.science.org/content/article/ai-often-mangles-african-languages-local-scientists-and-volunteers-are-taking-it-back
[21] Thomas Reitmaier, Electra Wallington, Dani Kalarikalayil Raju, Ondrej Klejch, Jennifer Pearson, Matt Jones, Peter Bell, and Simon Robinson. 2022. Opportunities and challenges of automatic speech recognition systems for low-resource language speakers. In Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (pp. 1-17).
[22] Sarah Repucci and Amy Slipowitz. 2020. Democracy under lockdown: the impact of COVID-19 on the global struggle for freedom. Freedom House. https://freedomhouse.org/report/special-report/2020/democracy-under-lockdown
[23] Samuel C. Rhodes. 2022. Filter bubbles, echo chambers, and fake news: How social media conditions individuals to be less critical of political misinformation. Political Communication, 39(1), 1-22.
[24] Tony Roberts, Judy Gitahi, Patrick Allam, Lawrence Oboh, Oyewole Oladapo, Gifty

Appiah-Adjei, Amira Galal, Jimmy Kainja, Sam Phiri, Kiss Abraham, Sebastian Klovig Skelton, and Anand Sheombar. 2023. Mapping the Supply of Surveillance Technologies to Africa: Case Studies from Nigeria, Ghana, Morocco, Malawi, and Zambia, Brighton: Institute of Development Studies, DOI: 10.19088/IDS.2023.027
[25] Peter J. Schraeder. 1995. Understanding the" Third Wave" of Democratization in Africa.
[26] Almog Simchon, Matthew Edwards, and Stephan Lewandowsky. 2024. The persuasive effects of political microtargeting in the age of generative artificial intelligence. PNAS nexus, 3(2), pgae035.
[27] Carol Soon and Samantha Quek. 2024. Safeguarding Elections from Threats Posed by Artificial Intelligence. IPS Working Papers No. 56. Institute of Policy Studies.
[28] Elvis Bisong Tambe and Moletsane Monyake. 2023. The impact of corruption and clientelism on voter turnout in Africa. Crime, Law and Social Change, 80(3), 329-352.
[29] Carlos A. Thomas, Victor W. Mbarika, Rufus Nwogu, Philip F. Musa, and Peter Meso. 2009. Facilitating better governance through E-government initiatives: Successful case in sub-Saharan Africa.
[30] World Bank. 2021. Prioritizing the poorest and most vulnerable in West Africa: Togo's Novissi platform for social protection uses machine learning, geospatial analytics, and mobile phone metadata for the pandemic response. World Bank Group. https://www.worldbank.org/en/results/2021/04/13/prioritizing-the-poorest-and-most-vulnerable-in-west-africa-togo-s-novissi-platform-for-social-protection-uses-machine-l
[31] Danni Xu, Shaojing Fan, and Mohan Kankanhalli. 2023. Combating misinformation in the era of generative AI models. In Proceedings of the 31st ACM International Conference on Multimedia (pp. 9291-9298).

Every Computing Project Shapes the World

by Samuel Mann, Ruth Myers, Oliver Bates, and Doug Schuler

Samuel Mann, Ruth Myers, Oliver Bates, Doug Schuler

ACM Computers and Society
Chinasa T. Okolo, Editor-in-Chief
Volume 53 ● Number 2 ● Winter 2025 ● www.sigcas.org
ACM Computers & Society (ISSN 0095-2737) is published by the Association for Computing Machinery (ACM), 1601 Broadway, 10th Floor, New York, NY 10019-7434.

Want to contribute to Computers and Society? The door is open, please walk through! We're all volunteers. Don't have time for a column? Send your ideas, comments, or suggestions to: editors_sigcas@acm.org.