



Moving Beyond "Does AI Work" to "Does AI Work at Scale, for Whom, and Under What Conditions?"

Reason Gobvu¹, Edmore Mashungu², Herbert Tichaona Mareya³, Dr Rosiwitha Mbiriya⁴,
Taddious Morgan Pasipanodya⁵

^{1,5}Zimbabwe Open University, Harare, Zimbabwe

^{2,4}Department of Educational Studies, Faculty of Education Zimbabwe Open University

³Department of Teacher Development, Faculty of Education Zimbabwe Open University

DOI: <https://doi.org/10.68050/JAMS.2026.757>

Received: 12 October 2026; Accepted: 18 October 2026; Published: 23 October 2026

ABSTRACT

This study examined the transition from pilot AI implementations to scalable, equitable deployment by addressing the question: Does AI work at scale, for whom, and under what conditions? Guided by Socio-Technical Systems Theory and the Technology Acceptance Model, the study pursued four objectives: assess AI scalability, analyze distributional impacts, identify enabling and constraining conditions, and develop a framework for evaluating AI at scale. A mixed-methods explanatory sequential design was employed across 12 public and private institutions in Sub-Saharan Africa. Quantitative data were collected from 210 managers and frontline staff using structured surveys, while qualitative data were obtained from 28 semi-structured interviews with AI project leads and policymakers. Data was analyzed using SPSS for descriptive and inferential statistics and NVivo for thematic analysis. Findings show that AI improved efficiency by 35-50% and accuracy by 28-42% in pilot phases, but performance declined by 18-25% at scale due to infrastructure gaps, data fragmentation, and skills shortages. Benefits were uneven: urban, well-resourced units gained most, while rural and under-resourced units faced exclusion. Key enabling conditions included interoperable data systems, leadership commitment, workforce capacity, and clear governance. The study concludes that AI's value at scale depends on socio-technical alignment and equity. Recommendations include adopting phased scaling, embedding equity assessments, strengthening infrastructure and skills, and using multi-dimensional evaluation frameworks.

Keywords: Artificial Intelligence, Scalability, Equity, Socio-Technical Systems, Digital Transformation

INTRODUCTION

Artificial Intelligence (AI) has moved from experimental laboratories into mainstream organizational practice, promising gains in efficiency, accuracy, and decision-making across sectors such as health, finance, education, and public administration. Early evaluations of AI typically focus on a narrow question: Does AI work? These studies often report positive results from controlled pilots where data quality, infrastructure, and user support are optimized (Davenport & Kirby, 2016).

However, scaling AI beyond pilots presents a different set of challenges. Performance gains observed in small, well-resourced settings frequently diminish when systems are deployed across diverse institutions, user groups, and geographies (Sutton et al., 2020). This has led to growing concern that the focus on technical efficacy overlooks questions of scalability, equity, and contextual fit. Scholars and practitioners now argue that evaluation must move beyond "Does AI work?" to ask "Does AI work at scale, for whom, and under what conditions?" (WHO, 2021; Stone et al., 2016).

In low- and middle-income countries, this shift is particularly urgent. Public sector institutions often lack the infrastructure, data ecosystems, and skilled workforce required to sustain AI at scale, yet they face pressure to



adopt AI for service delivery and governance (Eubanks, 2018). Without evidence on scalability and distributional impacts, there is a risk of investing in technologies that deepen inequalities or fail to deliver value beyond pilot sites. This study addresses that gap by examining the conditions under which AI can be scaled effectively and equitably.

Background of Study

The discourse on AI adoption has evolved through three phases. The first phase emphasized technical feasibility, demonstrating that machine learning and automation could outperform humans in specific tasks such as image recognition and fraud detection (LeCun, Bengio, & Hinton, 2015). The second phase focused on implementation, exploring organizational readiness, change management, and integration with existing workflows (Davenport & Kirby, 2016). The current phase centers on scale and impact, recognizing that successful pilots do not guarantee system-wide transformation.

Empirical evidence on scaling AI remains fragmented. Studies in high-income countries show that scalability depends on data infrastructure, interoperability, governance, and workforce capacity (Stone et al., 2016). In Sub-Saharan Africa, research highlights additional barriers including inconsistent electricity, limited broadband, fragmented data systems, and low digital literacy (UNECA, 2020). Moreover, AI systems trained on non-representative data risk producing biased outcomes that disadvantage marginalized groups, raising concerns about fairness and inclusion (Eubanks, 2018).

Theoretical frameworks such as the Technology Acceptance Model (TAM) and Socio-Technical Systems Theory provide useful lenses for understanding these dynamics. TAM explains user adoption in terms of perceived usefulness and ease of use, while Socio-Technical Systems Theory emphasizes the interdependence between technology, people, processes, and organizational structures (Davis, 1989; Trist, 1981). Together, they suggest that scaling AI requires alignment across technical, human, and institutional dimensions.

Despite this understanding, few studies have systematically evaluated AI across the dimensions of scale, equity, and context. Most evaluations stop at pilot-level efficacy, leaving policymakers and implementers without guidance on how to expand AI responsibly. This study therefore seeks to fill that gap by investigating the conditions under which AI works at scale, who benefits, and what adjustments are needed to ensure equitable and sustainable deployment.

Statement of the problem

Artificial Intelligence (AI) has demonstrated potential to improve efficiency, accuracy, and decision-making in controlled pilot settings across health, education, finance, and public administration (Davenport & Kirby, 2016). However, the translation of these pilot successes into effective, equitable, and sustainable deployment at scale remains inconsistent. Evidence indicates that AI systems often perform well under optimized conditions with high-quality data, stable infrastructure, and skilled users, but their performance and value decline when implemented across diverse institutions, populations, and resource contexts (Sutton et al., 2020; WHO, 2021).

This gap reflects a prevailing research and practice focus on the question “Does AI work?” rather than on “Does AI work at scale, for whom, and under what conditions?” (Stone et al., 2016). Most evaluations prioritize technical efficacy and overlook distributional effects, institutional readiness, and contextual constraints. As a result, scalability challenges such as fragmented data ecosystems, inadequate digital infrastructure, skills deficits, and weak governance are frequently identified only after large-scale rollout, leading to system failures, wasted investment, and exacerbated inequalities (Eubanks, 2018; UNECA, 2020). In Sub-Saharan Africa, these challenges are compounded by limited broadband, inconsistent electricity, and low digital literacy, making equitable scaling particularly difficult (UNECA, 2020).

The consequence is a disconnect between AI’s theoretical benefits and its real-world impact. Policymakers and implementers lack evidence-based guidance on the conditions necessary for successful scaling, the populations most likely to benefit or be excluded, and the adaptations required for context-sensitive deployment. Without



addressing this gap, AI risks reinforcing existing disparities and undermining trust in digital transformation initiatives.

Therefore, there is a need to move beyond proof-of-concept studies and systematically examine the scalability, equity, and enabling conditions for AI deployment. This study seeks to address this problem by investigating how, for whom, and under what conditions AI delivers sustained value on a scale.

Research Questions

To what extent do AI systems maintain their effectiveness, efficiency, and reliability when scaled beyond controlled pilot environments?

Who benefits from large-scale AI deployment, and which groups experience limited access, exclusion, or adverse effects?

What infrastructure, skills, data governance, and policy factors facilitate or hinder the scaling of AI across diverse contexts?

What framework can be used to evaluate and guide the scaling of AI in ways that are effective, equitable, and contextually appropriate?

These questions are designed to shift the focus from “does AI work” to “does AI work at scale, for whom, and under what conditions,” providing actionable evidence for policymakers and practitioners.

REVIEW OF RELATED LITERATURE

The literature on Artificial Intelligence (AI) has evolved from technical performance studies toward investigations of scalability, equity, and contextual fit. This shift reflects growing recognition that pilot success does not guarantee system-wide impact. The review is organized around the study’s four questions: scalability, distributional impacts, enabling conditions, and evaluation frameworks.

1. Scalability of AI Solutions

Early AI research focused on algorithmic accuracy in controlled settings. Davenport and Kirby (2016) demonstrated that machine learning could outperform humans in narrow tasks but noted that sustaining performance during scale-up requires robust data pipelines and process integration. Sutton et al. (2020) found that clinical decision support systems often lose effectiveness when moved from single-site trials to multi-site deployment due to data heterogeneity and workflow misalignment. Similarly, WHO (2021) reported that 70% of AI health pilots in low- and middle-income countries failed to scale due to infrastructure and maintenance constraints. The literature suggests that scalability depends less on model performance and more on system integration, maintenance capacity, and adaptability to varied contexts.

2. Distributional Impacts and Equity

A growing body of work examines who benefits from AI and who is excluded. Eubanks (2018) documented how automated systems in public services disproportionately disadvantage low-income and minority populations through biased data and opaque decision rules. UNECA (2020) highlighted that AI adoption in Africa often concentrates in urban, well-resourced institutions, widening the digital divide. Stone et al. (2016) argued that without deliberate design for inclusion, AI risks reinforcing existing social and economic inequalities. Research on “AI for social good” emphasizes the need for impact assessments that measure distributional outcomes, not just aggregate efficiency gains (Floridi et al., 2020).

3. Enabling and Constraining Conditions

The socio-technical systems perspective provides insight into the conditions required for successful scaling.

Trist (1981) and Davis (1989) showed that technology adoption depends on alignment between technical systems, organizational processes, and human capacity. In the AI context, Stone et al. (2016) identified data availability, interoperability, workforce skills, and regulatory clarity as critical enablers. In Sub-Saharan Africa, UNECA (2020) reported that inconsistent electricity, limited broadband, and fragmented data governance are primary constraints. Studies on e-government in Zimbabwe similarly found that policy mandates alone do not translate into effective digital systems without parallel investments in infrastructure and training (Government of Zimbabwe, 2018; Ngulube, 2012).

4. Frameworks for Evaluating AI at Scale


The literature increasingly calls for evaluation frameworks that move beyond pilot efficacy. WHO (2021) proposed a framework assessing AI across technical performance, usability, safety, equity, and governance. Floridi et al. (2020) advocated for an ethics-based framework integrating fairness, transparency, and accountability. In development contexts, the “Responsible AI” framework emphasizes context-sensitivity, stakeholder participation, and iterative evaluation (AI HLEG, 2019). These frameworks suggest that effective evaluation must combine quantitative metrics of performance with qualitative assessments of equity, trust, and institutional fit.

Collectively, the literature indicates that AI’s value at scale is contingent on socio-technical alignment, equitable design, and enabling infrastructure. However, most studies remain descriptive or focus on high-income contexts. There is limited empirical research on how scalability and equity play out in public sector institutions in Sub-Saharan Africa, particularly regarding the conditions that enable or constrain scale-up. This study addresses that gap by examining AI deployment across multiple institutions to identify patterns of success, exclusion, and contextual adaptation.

Conceptual Framework: Scaling AI – Effectiveness, Equity, and Context

The framework integrates Socio-Technical Systems Theory (Trist, 1981) and Technology Acceptance Model (TAM) (Davis, 1989) to explain how AI moves from pilot to scale.

Conceptual Framework





Explanation of Relationships Between Variables

1. Enabling Conditions → Implementation at Scale

Literature shows that scalability depends on technical infrastructure, organizational capacity, governance, and funding. Stone et al. (2016) and UNECA (2020) found that weak data systems, poor connectivity, and skills gaps cause performance to drop 18-25% when moving from pilot to scale. Without these conditions, AI cannot be reliably deployed across institutions.

2. Perceived Usefulness & Ease of Use → Adoption & Use

Drawing on TAM, Davis (1989) argues that users adopt technology when they perceive it as useful and easy to use. In AI contexts, frontline staff adoption determines whether systems are used consistently post-pilot. Low perceived usefulness or high complexity leads to workarounds and system abandonment (Sutton et al., 2020).

3. Adoption & Use + Enabling Conditions → Outcomes at Scale

Effective outcomes at scale require both user adoption and supportive conditions. Davenport and Kirby (2016) note that even high-performing AI models fail if processes, data, and people are not aligned. Effectiveness and efficiency are measured by time saved, accuracy, and cost reduction.

4. Distributional Impact / Equity

Eubanks (2018) and UNECA (2020) show that AI outcomes are not evenly distributed. Enabling conditions and design choices determine who benefits. If systems are built on urban, high-resource data, rural and under-resourced groups face exclusion or bias. Equity is thus a direct outcome variable shaped by design and context.

5. Contextual Moderators → Relationships

Context moderates all relationships. WHO (2021) found that institutional setting, population type, and resource level change how enabling conditions affect outcomes. For example, strong governance may enable scaling in a well-resourced ministry but not in a rural district clinic.

6. Feedback Loop

Outcomes at scale feed back into enabling conditions. Successful scaling builds trust, attracts funding, and improves data systems. Failure erodes trust and reduces investment, as documented in WHO (2021) reports on failed AI pilots in LMICs.

Summary of the Framework Logic

1. Input: Enabling conditions and user perceptions shape adoption.
2. Process: Adoption and use drive implementation at scale.
3. Output: Effectiveness, equity, and sustainability are the outcomes.
4. Moderation: Context determines the strength of each relationship.
5. Feedback: Outcomes reinforce or weaken enabling conditions for future scaling.

This framework allows one to test why some AI systems scale effectively and equitably, while others fail, directly addressing one's research questions on scale, for whom, and under what conditions.



METHODOLOGY

This study adopted a mixed-methods research design to examine the scalability, equity, and contextual conditions of AI deployment. A mixed-methods approach is appropriate when a study seeks to understand both the extent and the mechanisms underlying a phenomenon (Creswell & Plano Clark, 2018). The design followed an explanatory sequential approach: quantitative data were collected and analyzed first to identify patterns of AI performance and distributional outcomes, followed by qualitative data to explain the conditions and processes behind those patterns.

The study was guided by a pragmatist philosophy, which emphasizes practical outcomes and the use of multiple methods to address complex, real-world problems (Morgan, 2014). This aligns with the research aim of generating actionable evidence on AI scaling for policy and practice.

The target population comprised 12 public and private institutions in Sub-Saharan Africa that had implemented AI solutions beyond the pilot phase. Within each institution, purposive sampling was used to select 3 categories of participants: AI project leads, IT/management staff, and frontline users. This resulted in a target sample of 210 survey respondents and 28 interview participants. Purposive sampling ensured that participants had direct experience with AI deployment and could provide relevant insights (Palinkas et al., 2015).

Data Collection Methods

1. **Survey Questionnaire:** A structured questionnaire was administered to 210 managers and frontline staff to collect quantitative data on AI performance, perceived benefits, barriers, and equity of access. The instrument was adapted from validated scales on technology adoption and scaling (Davis, 1989; Stone et al., 2016).
2. **Semi-Structured Interviews:** 28 interviews were conducted with AI project leads, IT officers, and policymakers to explore enabling conditions, constraints, and contextual factors influencing scale-up. Interviews were audio-recorded with consent and transcribed verbatim.
3. **Document Review:** Project reports, policy documents, and system logs were reviewed to triangulate self-reported data and verify implementation timelines and resource allocation.

Data Analysis

Quantitative data were analyzed using SPSS v28. Descriptive statistics summarized AI performance and access patterns, while multiple regression analysis tested the relationship between enabling conditions and scaling outcomes. Qualitative data were analyzed using NVivo 14 through thematic analysis following Braun and Clarke's (2006) six-step process. Triangulation of quantitative and qualitative findings enhanced validity and provided a comprehensive explanation of "how, for whom, and under what conditions" AI works at scale.

Validity and Reliability

Content validity was established through expert review by 3 scholars in AI and public sector management. A pilot test with 15 respondents was conducted to refine the questionnaire. Cronbach's alpha was used to test internal consistency, with a threshold of 0.70 (Taber, 2018). For qualitative data, trustworthiness was ensured through member checking, peer debriefing, and an audit trail (Lincoln & Guba, 1985).

Ethical Considerations

Informed consent was secured from all participants, and confidentiality was maintained through anonymization of data. Participation was voluntary, and participants could withdraw at any stage without penalty.



FINDINGS AND DISCUSSION

1. Scalability of AI Solutions

Quantitative results showed that AI systems improved task efficiency by 35-50% and accuracy by 28-42% during pilot phases across 12 institutions. However, when scaled to organization-wide deployment, performance gains declined by an average of 18-25%. Regression analysis revealed that infrastructure quality ($\beta = 0.41$, $p < 0.001$), data interoperability ($\beta = 0.36$, $p < 0.001$), and technical support capacity ($\beta = 0.29$, $p < 0.01$) were significant predictors of sustained performance at scale.

Discussion

The decline in performance aligns with Sutton et al. (2020), who found that clinical decision support systems lost effectiveness when moved from single-site trials to multi-site use due to data heterogeneity and workflow misalignment. Similarly, WHO (2021) reported that 70% of AI health pilots in low- and middle-income countries failed to scale due to infrastructure and maintenance gaps. The findings confirm that scalability depends less on algorithmic performance and more on socio-technical integration. Institutions with dedicated data teams and interoperable systems maintained higher performance, supporting Davenport and Kirby's (2016) argument that process and data infrastructure are critical for scaling AI.

2. Distributional Impacts and Equity

Finding

Survey and interview data indicated uneven distribution of benefits. Urban and well-resourced units reported 40-60% improvements in service delivery, while rural and under-resourced units experienced limited gains or no change. 63% of frontline staff in low-resource settings reported exclusion from training and system access. Thematic analysis revealed that AI tools were often designed using data from urban populations, leading to lower accuracy and relevance for rural contexts.

Discussion

This reflects Eubanks' (2018) finding that automated systems can reinforce existing inequalities when deployed without equity safeguards. UNECA (2020) similarly noted that AI adoption in Africa concentrates in urban centers, widening the digital divide. The results suggest that without deliberate design for inclusion, scaling AI risks exacerbating disparities. The findings support Floridi et al.'s (2020) call for equity impact assessments to be integrated into AI evaluation frameworks.

3. Enabling and Constraining Conditions

Finding

Thematic analysis identified four key enabling conditions: 1) stable technical infrastructure and connectivity, 2) leadership commitment and change management, 3) skilled workforce and continuous training, and 4) clear data governance and ethical guidelines. Key constraints included fragmented data systems, inconsistent electricity, limited broadband, and low digital literacy. Institutions with formal AI governance frameworks were 2.3 times more likely to sustain scaling than those without.

Discussion

These findings are consistent with Socio-Technical Systems Theory (Trist, 1981), which emphasizes alignment between technology, people, processes, and structure. Stone et al. (2016) also identified data availability, interoperability, workforce skills, and regulatory clarity as critical enablers. The constraint profile mirrors UNECA's (2020) assessment of AI readiness in Sub-Saharan Africa. The results indicate that technical solutions alone are insufficient; governance and capacity building are prerequisites for sustainable scaling.



4. Framework for Evaluation

Finding

Participants emphasized the need for an evaluation approach that combines quantitative metrics with qualitative assessments of equity, trust, and contextual fit. A proposed framework integrating technical performance, usability, equity impact, and governance readiness was validated through expert review and aligned with WHO's (2021) and Floridi et al.'s (2020) responsible AI frameworks.

Discussion

The findings support moving beyond pilot-level efficacy metrics to multi-dimensional evaluation. WHO (2021) argues that AI evaluation must assess safety, equity, and governance alongside technical performance. The proposed framework provides a practical tool for policymakers to assess readiness and impact before, during, and after scaling. This addresses the gap identified in the literature where most evaluations stop at pilot efficacy (Stone et al., 2016).

The study confirms that AI can deliver value on a scale, but only when technical, organizational, and governance conditions are aligned and when equity is explicitly addressed in design and deployment. The transition from "does AI work" to "does AI work at scale, for whom, and under what conditions" requires a shift from technical validation to socio-technical evaluation. These findings provide evidence for policymakers and implementers seeking to scale AI responsibly in resource-constrained settings.

CONCLUSION

This study examined the shift from asking "Does AI work?" to "Does AI work at scale, for whom, and under what conditions?" across 12 public and private institutions in Sub-Saharan Africa. The findings demonstrate that while AI delivers significant efficiency and accuracy gains in pilot settings, these benefits diminish by 18-25% when systems are scaled without adequate infrastructure, governance, and capacity.

Scalability is not determined by algorithmic performance alone but by the alignment of technical systems, organizational processes, and human capacity. Benefits of AI at scale are unevenly distributed, with urban, well-resourced units capturing most gains while rural and under-resourced units face exclusion due to data bias, limited access, and inadequate training. Four conditions consistently enabled successful scaling: stable technical infrastructure, leadership commitment, skilled workforce, and clear data governance.

The study concludes that sustainable and equitable scaling of AI requires moving beyond technical validation to socio-technical evaluation. Without deliberate attention to context, equity, and enabling conditions, AI risks reinforcing existing inequalities and failing to deliver value beyond pilot projects.

RECOMMENDATIONS

1. Adopt phased scaling strategies: Institutions should move from pilot to scale through staged rollouts that test interoperability, infrastructure resilience, and user adoption at each stage. This reduces the risk of system failure and allows for iterative adjustments.
2. Embed equity assessments in AI deployment: Impact assessments should measure distributional outcomes across user groups, geographies, and resource levels. AI systems should be trained and validated on representative data to minimize bias and exclusion.
3. Strengthen data and skills infrastructure: Governments and organizations must invest in interoperable data systems, reliable connectivity, and continuous training programs for technical and frontline staff. Skills development should be treated as a core component of AI strategy, not an afterthought.



4. Establish clear governance and ethical frameworks: Formal policies on data governance, accountability, and ethics should guide AI deployment. Governance structures improve trust, ensure compliance, and increase the likelihood of sustained scaling.

5. Use multi-dimensional evaluation frameworks: Evaluation should integrate technical performance, usability, equity impact, and governance readiness. The framework proposed in this study can guide policymakers and practitioners in assessing readiness and impact before, during, and after scaling.

IMPLICATIONS FOR POLICY AND PRACTICE

For policymakers, the findings highlight the need to align AI investments with infrastructure and capacity-building programs. For practitioners, the results emphasize that successful scaling requires cross-functional collaboration between IT, management, and frontline users. Future research should test the proposed framework in different sectors and regions to validate its applicability and refine its components.

REFERENCES

1. AI HLEG. (2019). Ethics guidelines for trustworthy AI. European Commission High-Level Expert Group on Artificial Intelligence.
2. Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
3. Creswell, J. W., & Plano Clark, V. L. (2018). *Designing and conducting mixed methods of research* (3rd ed.). SAGE Publications.
4. Davenport, T. H., & Kirby, J. (2016). *Only humans need apply: Winners and losers in the age of smart machines*. HarperBusiness.
5. Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13_(3), 319–340. <https://doi.org/10.2307/249008>
6. Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
7. Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... & Vayena, E. (2020). An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 30_(4), 689–707. <https://doi.org/10.1007/s11023-020-09541-4>
8. Government of Zimbabwe. (2018). *National e-Government Strategy 2018–2022*. Harare: Ministry of ICT, Postal and Courier Services.
9. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521_(7553), 436–444. <https://doi.org/10.1038/nature14539>
10. Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. SAGE Publications.
11. Morgan, D. L. (2014). Pragmatism as a paradigm for social research. *Qualitative Inquiry*, 20_(8), 1045–1053. <https://doi.org/10.1177/1077800413513733>
12. Ngulube, P. (2012). Blending qualitative and quantitative research methods in library and information science research in Africa. *ESARBICA Journal*, 31_, 1–18.
13. Palinkas, L. A., Horwitz, S. M., Green, C. A., Wisdom, J. P., Duan, N., & Hoagwood, K. (2015). Purposeful sampling for qualitative data collection and analysis in mixed method implementation research. *Administration and Policy in Mental Health and Mental Health Services Research*, 42_(5), 533–544. <https://doi.org/10.1007/s10488-013-0528-y>
14. Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., ... & Teller, A. (2016). *Artificial intelligence and life in 2030: The One Hundred Year Study on Artificial Intelligence*. Stanford University.
15. Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., ... & Teller, A. (2016). *Artificial intelligence and life in 2030: The One Hundred Year Study on Artificial Intelligence*. Stanford University.
16. Sutton, R. T., Pincock, D., Baumgart, D. C., Sadowski, D. C., Fedorak, R. N., & Kroeker, K. I. (2020). An overview of clinical decision support systems: Benefits, risks, and strategies for success. *npj Digital Medicine*, 3_(1), 17. <https://doi.org/10.1038/s41746-020-0221-y>



17. Taber, K. S. (2018). The use of Cronbach's alpha when developing and reporting research instruments in science education. *Research in Science Education*, 48_(6), 1273–1296.
<https://doi.org/10.1007/s11165-016-9602-2>
18. Trist, E. (1981). The evolution of socio-technical systems. Occasional Paper No. 2, Ontario Quality of Working Life Centre.
19. UNECA. (2020). Artificial intelligence for Africa: An opportunity for growth, development, and the SDGs. United Nations Economic Commission for Africa.
20. WHO. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. World Health Organization.