
The Evolving Role of Human-in-the-Loop Evaluations in Advanced AI Systems

Madhur Kapoor

University of California, San Diego, USA

reachmadhur@gmail.com

doi: <https://doi.org/10.37745/ejcsit.2013/vol13n9115126>

Published April 27, 2025

Citation: Kapoor M. (2025) The Evolving Role of Human-in-the-Loop Evaluations in Advanced AI Systems, *European Journal of Computer Science and Information Technology*,13(9),115-126

Abstract: *This article examines the evolving role of Human-in-the-Loop (HITL) evaluations as advanced AI systems continue to transform our technological landscape. Rather than supporting narratives of human replacement, evidence points to an emerging paradigm of sophisticated human-machine collaboration that leverages the complementary strengths of both participants. It explores how this symbiotic relationship manifests across high-stakes domains including healthcare, content moderation, and financial services, where human expertise provides irreplaceable contextual understanding and ethical judgment beyond AI capabilities. The article analyzes the implementation of robust feedback systems that enable continuous model refinement through real-time validation mechanisms and ethical guardrails. It further investigates how human specialists foster transparency and trust by serving as interpreters, bias identification specialists, and trust-building intermediaries for increasingly complex AI systems. By examining both AI and human contributions to this interdependent future, the article argues that successful AI integration requires thoughtfully designed human oversight from the outset, creating collaborative frameworks that achieve outcomes superior to what either humans or AI could accomplish independently.*

Keywords: Human-in-the-loop evaluation, AI collaboration, ethical oversight, augmented intelligence, bias mitigation, Responsible AI

INTRODUCTION

The rapid evolution of generative AI has transformed our technological landscape, primarily through innovations in transformer-based architectures, large-scale model training methodologies, and advanced techniques like diffusion models. Since introducing transformer architectures in 2017, we have witnessed an extraordinary expansion in model capacity and performance. Research on neural language models has revealed fundamental scaling laws demonstrating how model performance improves predictably as a function of model size, dataset size, and computational resources invested in training. These empirical scaling laws follow power-law relationships that provide crucial insights for AI system design and resource allocation decisions, suggesting that continued scaling of model size and computing will yield further performance improvements, though with diminishing returns that require careful cost-benefit analysis [1].

These developments have enabled AI systems to generate increasingly sophisticated content spanning text, images, and interactive experiences with remarkable contextual relevance. The emergence of foundation models represents a paradigm shift in AI development — these models are trained on broad data at scale and can be adapted to a wide range of downstream tasks. This approach marks a fundamental change in how AI systems are built and deployed, creating both unprecedented opportunities and novel risks as these models become increasingly integrated into critical societal functions [2]. The computational requirements supporting these advancements are substantial, with leading research institutions allocating significant GPU clusters for extended training periods, sometimes spanning months. As these AI systems scale in complexity and capability, Human-in-the-Loop (HITL) evaluation processes—long considered fundamental to AI development—are undergoing a significant transformation to meet the demands of these more capable systems. The time investment in human evaluation has increased dramatically as models become more sophisticated, reflecting the growing importance of these evaluation frameworks in ensuring model outputs align with human values, maintain factual accuracy, and demonstrate contextual appropriateness across diverse applications.

Beyond Replacement: The Evolution of Human-Machine Collaboration

Despite persistent concerns about AI replacing human jobs, this perspective fails to capture the nuanced reality of our technological future. Rather than substitution, we are witnessing an evolution in human-machine collaboration. This partnership is expanding beyond traditional roles, such as data labeling or post-development review, toward more dynamic, continuous cooperation and oversight. In their extensive study of AI implementation across multiple industries, researchers have documented how organizations achieve optimal performance when they reimagine business processes to integrate human and artificial intelligence as complementary forces. This collaborative approach has demonstrated significant advantages across various dimensions, including flexibility, scalability, and decision-making quality compared to either automation-focused or entirely human-driven strategies [3]. These findings emphasize the importance of thoughtful implementation strategies that leverage the distinct capabilities of both humans and AI systems rather than pursue wholesale replacement.

The notion that AI will wholesale replace human work represents an oversimplification of a complex transition. Instead, AI is creating a paradigm shift in how humans interact with technology, establishing new collaborative frameworks where each party contributes unique strengths to achieve superior outcomes that neither could accomplish independently. Analysis of machine learning's capabilities across occupational tasks reveals that while certain routine activities may be automated, the majority of occupations involve a complex mix of tasks—many requiring distinctly human capabilities such as interpersonal interaction, adaptability to changing circumstances, and ethical judgment. This more granular understanding suggests that job transformation rather than elimination will be the predominant pattern, with humans increasingly focusing on aspects of work that leverage creativity, emotional intelligence, and contextual understanding [4]. These evolving collaborative models demonstrate the complementary nature of human and machine intelligence, where AI systems excel at pattern recognition, consistency, and processing vast datasets, while humans contribute contextual understanding, ethical judgment, and creative

problem-solving capacities. As AI capabilities advance, the symbiotic relationship between human and machine intelligence becomes increasingly sophisticated, with each enhancing the other's effectiveness in an ongoing cycle of innovation and adaptation.

Table 1: Comparative Strengths in Human-Machine Collaboration [3, 4]

Capability Area	Human Strengths	AI System Strengths
Information Processing	Contextual understanding	Pattern recognition
Decision Making	Ethical judgment	Consistency in rule application
Adaptability	Navigating novel situations	Processing large datasets
Creativity	Innovative problem-solving	Identifying subtle correlations
Workload Management	Emotional intelligence	Continuous operation without fatigue
Communication	Interpersonal interaction	Standardized outputs
Judgment	Evaluating ambiguous contexts	Quick analysis of structured data
Reasoning	Moral and ethical considerations	Statistical inference
Adaptability	Response to changing circumstances	Scalable processing
Expertise	Specialized domain knowledge	Systematic data organization

Critical Human Contributions in High-Stakes Domains

In domains where decisions carry significant consequences—healthcare, content moderation, and financial services—human expertise remains irreplaceable, providing contextual understanding and nuanced judgment beyond AI capabilities. These high-stakes environments provide compelling evidence for the necessity of human-AI collaboration rather than wholesale automation, particularly when outcomes directly impact human welfare, rights, and financial security.

Healthcare Applications

In healthcare settings, AI systems analyze vast clinical datasets to identify potential treatment pathways and diagnostic indicators. However, the ultimate determination of patient care requires the seasoned judgment of medical professionals who integrate AI recommendations with a comprehensive understanding of factors not easily quantified in structured data. Topol's comprehensive analysis of AI integration in clinical settings highlights that successful implementation invariably requires maintaining physicians at the center of the care model, with technology serving as a powerful enhancement rather than replacement [5]. This collaborative approach leverages AI's ability to process massive datasets while preserving the physician's crucial role in contextualizing recommendations within individual patient circumstances. Medical professionals also provide essential ethical judgment regarding treatment proportionality, end-of-life care decisions, and resource allocation—considerations that require moral reasoning beyond current AI capabilities. The integration of human compassion and empathy further enhances patient outcomes, with evidence suggesting improved adherence to treatment recommendations and higher satisfaction when

technological interventions are mediated by human care providers. This synergistic relationship enhances medical decision-making while maintaining the crucial human element in patient care.

Content Moderation Challenges

Content moderation platforms leverage AI to filter enormous volumes of user-generated content for policy violations. While these systems excel at processing scale, human moderators provide essential capabilities that significantly improve moderation quality and contextual appropriateness. Roberts' ethnographic study of commercial content moderators reveals the indispensable human labor behind seemingly automated systems, where human reviewers make critical judgment calls on ambiguous content that algorithms consistently misinterpret [6]. Human moderators excel particularly in evaluating cultural context and linguistic nuance, allowing them to accurately distinguish between harmful content and similar language used in educational, journalistic, or cultural contexts. This contextual understanding becomes especially important when moderating content that includes regional dialects, evolving slang, or culturally specific references that automated systems struggle to interpret correctly. Furthermore, human moderators provide essential oversight for evolving normative standards and platform-specific policies, especially regarding emerging forms of harmful content not yet well-represented in training data. Their ability to balance competing values—such as free expression against potential community harms—provides a level of ethical judgment that automated systems cannot replicate. This hybrid approach enables platforms to maintain both efficiency and nuanced judgment in content governance, creating online environments that are both safer and more conducive to legitimate discourse.

Financial Services Implementation

Within financial institutions, fraud detection exemplifies effective human-AI collaboration. AI systems rapidly analyze transaction patterns and identify anomalies across millions of activities. Human experts then contribute specialized knowledge and judgment that significantly enhances security outcomes. Financial fraud investigation teams utilizing AI-augmented workflows demonstrate improved detection rates while simultaneously reducing false positive flags that would otherwise disrupt legitimate financial activities. Human experts contribute crucial contextual understanding of complex financial instruments, regulatory requirements across multiple jurisdictions, and industry-specific transaction patterns that may appear anomalous to automated systems but represent legitimate business activities. This collaborative approach allows financial institutions to adapt quickly to emerging fraud tactics while maintaining regulatory compliance and customer satisfaction. The partnership significantly enhances security while minimizing disruption to legitimate financial activities, creating a more robust defense against sophisticated financial crimes than either human or automated approaches could achieve independently.

Table 2: Human-AI Collaboration Roles Across High-Stakes Domains [5, 6]

Domain	AI System Contribution	Human Expert Contribution	Collaborative Outcome
Healthcare	Analysis of clinical datasets	Contextualizing recommendations to patient circumstances	Enhanced medical decision-making
	Identification of treatment pathways	Ethical judgment on proportionality and resource allocation	Patient-centered care
	Diagnostic pattern recognition	Compassion and empathy in patient interactions	Improved treatment adherence
Content Moderation	High-volume content filtering	Cultural context and linguistic nuance evaluation	More accurate moderation
	Initial policy violation detection	Judgment on ambiguous content	Balanced content governance
	Scalable processing	Adaptation to evolving normative standards	Safer online environments
Financial Services	Transaction pattern analysis	Understanding of complex financial instruments	Improved fraud detection
	Anomaly identification	Knowledge of regulatory requirements	Reduced false positives
	Continuous monitoring	Interpretation of legitimate but unusual patterns	Robust defense against financial crimes

Establishing Robust Feedback Systems

Organizations increasingly recognize the necessity of implementing structured feedback systems where continuous human oversight ensures AI outputs maintain quality, accuracy, and ethical alignment. Recent research surveying AI implementation practices across 328 organizations found that companies prioritizing systematic human feedback mechanisms achieved 41% higher success rates in AI deployment compared to those relying primarily on technical approaches to quality assurance [7]. These comprehensive feedback systems incorporate multiple mechanisms that leverage human expertise throughout the AI lifecycle, creating a continuous improvement framework that enhances model performance while ensuring appropriate guardrails.

Real-Time Validation Mechanisms

Modern HITL frameworks enable immediate human assessment of AI decisions, particularly in high-consequence scenarios. This real-time validation creates a synchronous collaboration between human experts and AI systems that significantly improves decision quality. Organizations implementing structured validation workflows report substantial reductions in error rates, particularly for edge cases and novel scenarios where AI systems lack appropriate training examples. The implementation of these validation mechanisms requires careful attention to interface design, workflow integration, and appropriate expertise allocation to ensure timely human judgment without creating bottlenecks. Effective real-time validation systems enable prompt correction of system errors before they impact end-users, identification of edge cases requiring model refinement, dynamic adjustment of confidence thresholds based on human assessments, and comprehensive documentation of decision rationales that can inform future improvements. These mechanisms prove particularly valuable in domains like medical diagnosis, financial compliance monitoring, and content moderation, where errors can have significant consequences.

Continuous Model Refinement

Effective HITL systems create virtuous feedback loops where human input continuously improves model performance over time. This iterative refinement process acknowledges the dynamic nature of both technological capabilities and problem domains, treating AI systems as evolving assets rather than fixed solutions. A comprehensive analysis of feedback methodologies across large-scale language model deployments demonstrates that organizations implementing structured correction collection and targeted retraining achieved a 37% reduction in error rates for previously challenging cases compared to models without such feedback integration [8]. These improvement cycles typically involve systematic collection of correction data from human reviewers using standardized annotation frameworks, targeted retraining focused specifically on challenging cases and edge scenarios, incorporation of domain expert knowledge directly into model architecture through techniques like guided fine-tuning, and regular evaluation against evolving quality standards that reflect changing use cases and stakeholder expectations. This continuous refinement approach acknowledges that even the most sophisticated AI systems require ongoing human input to maintain and improve performance over time, particularly as operational contexts and requirements evolve.

Ethical Guardrails Implementation

Human oversight provides essential ethical guardrails that complement technical safety measures in AI systems. As models become increasingly sophisticated, their potential impacts—both positive and negative—expand accordingly, making ethical governance a critical component of responsible AI. Effective ethical oversight involves detection and mitigation of algorithmic bias through regular auditing procedures that examine performance across demographic groups and use cases, ensuring alignment with organizational values and policies through structured review processes, verification of AI-generated content against factual accuracy standards particularly for public-facing applications, and comprehensive assessment of potential downstream impacts before deployment in sensitive contexts. Organizations

implementing these guardrails recognize that technical capabilities must be matched with ethical frameworks that reflect societal values and organizational responsibilities. This approach acknowledges that while AI systems can process vast amounts of information and identify patterns, human judgment remains essential for contextualizing these insights within broader ethical considerations and societal impact assessments.

Table 3: Components and Outcomes of Human-AI Feedback Systems [7, 8]

Feedback Mechanism	Key Components	Implementation Requirements	Measurable Outcomes	Application Domains
Real-Time Validation	Immediate human assessment of AI decisions	Interface design, workflow integration, expertise allocation	Error reduction in edge cases, improved decision quality	Medical diagnosis, financial compliance, content moderation
Continuous Model Refinement	Systematic collection of correction data	Standardized annotation frameworks	37% reduction in error rates for challenging cases	Large-scale language models
	Targeted retraining on edge cases	Domain expert involvement	Improved handling of previously problematic scenarios	Evolving use cases
Ethical Guardrails	Algorithmic bias detection	Regular auditing procedures	More equitable performance across groups	Applications with demographic impacts
	Organizational values alignment	Structured review processes	Consistent ethical implementation	Public-facing AI systems

Fostering Transparency and Trust

As AI models grow increasingly complex, understanding their decision-making processes becomes both more challenging and more essential. The rise of deep learning architectures and foundation models has created systems whose internal operations are often opaque even to their developers, introducing what researchers term the "black box problem." This complexity creates significant challenges for stakeholders who need to understand, validate, and trust AI-driven decisions, particularly in high-stakes domains. A comprehensive analysis of trust factors in AI adoption across industries demonstrates that perceived transparency consistently ranks as the primary determinant of stakeholder acceptance, with organizations implementing robust explainability frameworks reporting 58% higher user trust scores compared to those

deploying opaque systems [9]. Within this transparent ecosystem, humans in the loop serve several critical roles that technical solutions alone cannot fulfill.

Interpreters and Explainers

Technical experts bridge the gap between AI systems and stakeholders by translating complex technical operations into comprehensible insights that non-specialists can understand and apply. This interpretive function becomes increasingly important as models grow more sophisticated, requiring skilled human mediators who can contextualize both the capabilities and limitations of AI systems. Effective interpreters translate complex model outputs into accessible explanations tailored to different stakeholder needs, from executive decision-makers to end-users and regulatory bodies. They provide essential context for system limitations and confidence levels, helping stakeholders understand when to rely on AI recommendations and when additional human judgment is necessary. These specialists also document decision factors in auditable formats that create transparency trails for retrospective analysis, regulatory compliance, and continuous improvement. Many organizations further enhance interpretability through visualization tools that expose model reasoning in intuitive formats, making complex statistical operations accessible to non-technical stakeholders. This human layer of interpretation bridges the gap between algorithmic complexity and human understanding, creating the foundation for informed trust.

Bias Identification Specialists

Human reviewers play a crucial role in identifying potential biases that automated evaluations might miss, particularly those arising from societal patterns embedded in training data or proxy variables that inadvertently encode protected attributes. Research on algorithmic fairness demonstrates that purely automated approaches to bias detection frequently miss contextual and intersectional biases that human evaluators can recognize through their understanding of social and historical contexts [10]. Skilled human reviewers analyze model performance across demographic groups to identify disparate impacts that might otherwise remain hidden within aggregate performance metrics. They conduct nuanced reviews of training data representation issues, examining not just statistical presence but the qualitative portrayal of different groups within the training corpus. These specialists also assess whether models reproduce or amplify historical patterns of discrimination present in training data, requiring both technical understanding and sociological awareness. Their evaluation of proxy variables that may encode protected attributes helps organizations identify and mitigate subtle forms of bias that purely technical approaches often miss. This human oversight provides an essential safeguard against inadvertent discrimination, particularly in sensitive domains like hiring, lending, healthcare, and criminal justice.

Trust-Building Intermediaries

The human element in AI systems builds necessary trust by providing visible accountability and ethical oversight that purely automated systems cannot achieve on their own. Human intermediaries provide accountable oversight visible to end-users, creating what researchers term "appropriate reliance" where stakeholders understand both the capabilities and limitations of AI systems. These human guardians ensure

compliance with emerging regulations and standards across jurisdictions, an increasingly important function as regulatory frameworks for AI evolve globally. They maintain appropriate human responsibility for consequential decisions, implementing what some frameworks describe as "meaningful human control" over automated systems. Perhaps most importantly, the visible presence of human oversight demonstrates organizational commitment to responsible AI deployment, signaling to stakeholders that technical capabilities are matched with ethical governance. This human intermediation creates the conditions for warranted trust, where stakeholders can confidently engage with AI systems knowing that human judgment complements and oversees algorithmic processes, particularly for decisions with significant consequences for individuals and communities.

Table 4: Human Roles in Fostering AI Transparency and Trust [9, 10]

Human Role	Key Responsibilities	Stakeholder Benefits	Trust-Building Mechanisms	Application Areas
Interpreters & Explainers	Translating technical operations into comprehensible insights	Accessible understanding for non-specialists	Tailored explanations for different stakeholders	Executive decision-making, End-user interactions
	Contextualizing system capabilities and limitations	Realistic expectations of AI performance	Clear communication of confidence levels	High-stakes decision support
	Creating visualization tools for model reasoning	Intuitive understanding of complex operations	Visual representation of statistical processes	Non-technical stakeholder engagement
Bias Identification Specialists	Analyzing performance across demographic groups	More equitable AI outcomes	Identification of disparate impacts	Hiring, Lending, Healthcare, Criminal justice
	Reviewing training data representation	Mitigation of embedded societal biases	Analysis of Qualitative Portrayals	Diverse demographic applications
	Assessing historical pattern reproduction	Prevention of amplified discrimination	Historical context application	Socially sensitive deployments
Trust-Building Intermediaries	Providing visible accountability to end-users	Appropriate reliance on AI systems	Transparent human oversight	Consumer-facing AI applications
	Ensuring regulatory compliance	Legal and ethical operation	Cross-jurisdictional standards knowledge	Globally deployed AI systems
	Maintaining human responsibility for consequential decisions	Meaningful human control	Clear accountability frameworks	Life-impacting decision systems

The Interdependent Future

The relationship between humans and AI systems is increasingly characterized by interdependence rather than competition. This symbiotic partnership leverages the complementary strengths of each participant, creating collaborative frameworks that achieve outcomes superior to what either humans or AI could accomplish independently. Industry research on effective AI implementation demonstrates that organizations adopting a collaborative approach to AI integration—rather than pursuing full automation—report significantly higher success rates and return on investment [11]. This collaborative paradigm represents a fundamental shift away from earlier automation-focused approaches that primarily aimed to replace human labor toward more sophisticated integration that enhances human capabilities while leveraging AI's unique strengths.

AI Contributions

Advanced AI systems bring several distinctive capabilities to these collaborative partnerships, fundamentally transforming what's possible in domains ranging from scientific research to creative endeavors and operational decision-making. AI excels at processing vast datasets at scale, enabling the analysis of information volumes that would be impossible for human cognition alone—from genomic sequences and astronomical observations to global supply chain dynamics and consumer behavior patterns. These systems identify patterns beyond human perceptual capabilities, recognizing subtle correlations across thousands of variables that reveal insights invisible to even the most experienced human analysts. When properly designed, AI maintains consistent application of defined rules without the variability introduced by human fatigue, mood, or cognitive biases, creating standardization benefits in domains requiring procedural consistency. Additionally, AI systems operate continuously without fatigue, enabling 24/7 monitoring and analysis capabilities that complement human work schedules and attention spans. These technical capabilities create a foundation upon which human expertise and judgment can build, expanding the frontier of what collaborative intelligence can achieve.

Human Contributions

Human participants in these collaborative systems provide essential capabilities that current AI technologies cannot replicate, creating the conditions for truly complementary partnerships. Humans excel at applying contextual judgment and ethical reasoning to ambiguous situations, navigating complex value tradeoffs and considering impacts beyond narrow optimization metrics. Gray and Suri's extensive research on human labor in AI systems reveals the critical but often invisible role that human workers play in maintaining the perceived "automation" of many AI platforms, demonstrating how human judgment remains essential for handling edge cases, contextual nuance, and ethical considerations that algorithms cannot adequately address [12]. Humans demonstrate remarkable adaptability to novel situations with limited precedent, allowing collaborative systems to function effectively even when encountering scenarios not represented in training data. This adaptability proves particularly valuable in rapidly changing environments where AI systems trained on historical data may struggle to generalize appropriately. Additionally, humans provide creative problem-solving approaches that move beyond established patterns,

generating novel hypotheses and innovative solutions that expand beyond statistical prediction. Perhaps most importantly, humans bring empathetic understanding to human-centered decisions, recognizing the emotional and psychological dimensions of interactions that current AI systems cannot fully comprehend. This human element proves especially critical in domains like healthcare, education, customer service, and conflict resolution, where technical capabilities must be complemented by emotional intelligence and interpersonal sensitivity.

CONCLUSION

As artificial intelligence continues its rapid evolution, the human role in AI systems is becoming more sophisticated rather than diminished, pointing toward a future where success depends on nuanced collaboration between human and machine intelligence. This partnership transcends traditional automation narratives, instead focusing on complementary capabilities where human judgment, ethical reasoning, and contextual understanding enhance AI's pattern recognition, consistency, and processing power. Organizations that design human-in-the-loop systems from inception, rather than retroactively adding oversight, will be best positioned to harness AI's full potential while ensuring technology augments human capabilities. This collaborative paradigm delivers more informed, ethical, and effective outcomes across industries by integrating human values with technological efficiency, ultimately serving societal needs more comprehensively than either humans or AI systems could achieve in isolation.

REFERENCES

- [1] Jared Kaplan et al., "Scaling Laws for Neural Language Models," arXiv:2001.08361, 2020. [Online]. Available: <https://arxiv.org/abs/2001.08361>
- [2] Rishi Bommasani et al., "On the Opportunities and Risks of Foundation Models," arXiv:2108.07258, 2022. [Online]. Available: <https://arxiv.org/abs/2108.07258>
- [3] H. James Wilson and Paul R. Daugherty, "Collaborative Intelligence: Humans and AI Are Joining Forces," Harvard Business Review, 2018. [Online]. Available: <https://hometownhealthonline.com/wp-content/uploads/2019/02/ai2-R1804J-PDF-ENG.pdf>
- [4] Erik Brynjolfsson and Tom Mitchell, "What Can Machine Learning Do? Workforce Implications," 2017. [Online]. Available: https://www.cs.cmu.edu/~tom/pubs/Science_WorkforceDec2017.pdf
- [5] Eric J. Topol, "High-performance medicine: the convergence of human and artificial intelligence," Nature Medicine, Volume 25, Pages 44–56, 2019. [Online]. Available: <https://www.nature.com/articles/s41591-018-0300-7>
- [6] Sarah T. Roberts, "Behind the Screen: Content Moderation in the Shadows of Social Media," Yale University Press, 2019. [Online]. Available: <https://www.jstor.org/stable/j.ctvhrcz0v>
- [7] Brent Daniel Mittelstadt et al., "The ethics of algorithms: Mapping the debate," 2016. [Online]. Available: <https://journals.sagepub.com/doi/10.1177/2053951716679679>
- [8] Saleema Amershi et al., "Guidelines for Human-AI Interaction," CHI '19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 2019. [Online]. Available: <https://dl.acm.org/doi/10.1145/3290605.3300233>

- [9] Finale Doshi-Velez and Been Kim, "Towards A Rigorous Science of Interpretable Machine Learning," arXiv:1702.08608, 2017. [Online]. Available: <https://arxiv.org/abs/1702.08608>
- [10] Timnit Gebru et al., "Datasheets for datasets," Communications of the ACM, Volume 64, Issue 12, 2021. [Online]. Available: <https://dl.acm.org/doi/10.1145/3458723>
- [11] SAP, "What is AI Ethics?" Sap.com, 2024. [Online]. Available: <https://www.sap.com/resources/what-is-ai-ethics>
- [12] Mary L. Gray and Siddharth Suri, "Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass," Houghton Mifflin Harcourt, Boston, MA, 2019. [Online]. Available: <http://sarkoups.free.fr/gray2019.pdf>