

Designing Human-Centered AI Platforms for Scalable Global Software Through Integrated Product Engineering with Market-Ready Architecture

Kshitij Mahant¹, Xiaosi Yang², Shilpi Bhattacharya³

¹Sr. Manager, Technical Marketing

²Partner at XStar Capital, San Francisco Bay Area

³Product Manager at IBM, USA

ARTICLE INFO

Received: 01 Apr 2025

Revised: 15 May 2025

Accepted: 25 May 2025

ABSTRACT

The growing demand for intelligent, scalable, and ethically aligned software has intensified the need for AI platforms that balance technological robustness with human-centered design. This study presents a holistic framework for designing human-centered AI platforms intended for global scalability by integrating product engineering principles and market-ready architecture. Employing a design-based research methodology, the platform was co-developed with users through iterative prototyping and evaluated for usability, scalability, and fairness. Usability assessments using the System Usability Scale (SUS) showed high satisfaction rates across professional developers, citizen developers, and end users. Performance benchmarks confirmed the system's ability to operate efficiently across cloud, edge, and on-premise environments, with edge deployment yielding the most balanced latency and throughput. Multiple regression analysis revealed that architectural modularity, transparency, and multilingual support significantly influenced user adoption. A fairness audit across demographic groups found equitable outcomes overall, with slight disparities observed for non-native-language users. These results emphasize the importance of integrating human-centered principles into AI development to ensure accessibility, trust, and adaptability. The study concludes with a validated design approach that positions AI platforms to succeed both technologically and socially in global software ecosystems.

Keyword: Human-centered AI, scalable software, integrated product engineering, usability, fairness, market-ready architecture, user adoption, global deployment

Introduction

Background and context

In the era of rapid digital transformation, the design and development of artificial intelligence (AI) platforms are evolving beyond technical capabilities alone, shifting toward a more human-centric focus (Ciano et al., 2025). With increasing demands for scalable, intelligent software solutions that serve a global user base, the emphasis on designing AI systems that are not only technologically robust but also socially relevant has become a key research and industrial imperative (Li & Duan, 2025). Organizations today are investing heavily in AI to optimize operations, personalize services, and innovate product development, but many solutions still lack adaptability to diverse user needs across different regions and cultures (Pan et al., 2023). Therefore, there is a critical need to align AI development with human-centered principles to ensure usability, inclusivity, and adoption at scale.

Human-centered design in AI systems

Human-centered AI involves designing intelligent systems that consider human values, cognitive capacities, and behavioral patterns throughout the lifecycle of the software (Suganya et al., 2025). It emphasizes transparency, interpretability, and user empowerment. By incorporating user feedback loops, intuitive interfaces, and ethical safeguards, human-centered AI platforms can bridge the gap between technological potential and human experience (Hwang et al., 2025). This approach does not just improve user interaction but enhances trust, satisfaction, and real-world impact. In global software environments, where variability in language, infrastructure, and user literacy is substantial, such principles become even more critical to ensure equitable access and engagement (Chrysikou et al., 2024).

Scalability and global software demands

Scalability is one of the cornerstones of modern software engineering. For AI platforms to function globally, they must be scalable across markets with differing technical, regulatory, and user contexts (Bhatnagar et al., 2023). A scalable AI platform must manage data heterogeneity, support multi-modal interfaces, adapt to varying deployment environments (on-premises, cloud, edge), and handle cross-border privacy and security constraints. The fusion of AI scalability with a human-centered design approach ensures that growth in system capabilities aligns with the diverse needs of a global user base (Meselhy & Almalkawi, 2025). Moreover, the rise of decentralized digital infrastructures calls for architectures that are both adaptable and resilient, capable of rapid iteration without compromising core design principles.

Integrated product engineering approach

To achieve scalable and usable AI solutions, integrated product engineering plays a vital role. This approach unifies software development, user experience design, hardware compatibility, and lifecycle sustainability into a single coherent process. (Schiuma & Santarsiero, 2024) Instead of operating in silos, cross-functional teams collaborate to develop market-ready products that are not only intelligent and scalable but also robust, secure, and context-aware. Integrated product engineering enables early-stage alignment of AI capabilities with end-user requirements, reduces development cycles, and fosters continuous improvement. This is essential when delivering AI software across domains such as healthcare, finance, logistics, or education, where domain-specific customization is non-negotiable (Bermejo-Martín & Rodríguez-Monroy, 2020).

Market-ready architecture for sustainable deployment

The success of any AI platform depends heavily on the architecture that underpins it. A market-ready architecture is defined by its ability to support modularity, interoperability, and rapid deployment (Roach, 2024). It anticipates evolving business models and regulatory environments while providing flexibility for updates and enhancements. In the context of human-centered AI, the architecture must also accommodate user feedback mechanisms, transparency features, and adaptive learning capabilities (Bengler et al., 2023). Designing such an architecture requires close integration between engineering, product strategy, and market research. By embedding human values at the architectural level, developers can create AI platforms that are not only technically feasible but also socially and commercially sustainable.

Objective of the study

This research aims to present a multi-layered framework for designing human-centered AI platforms that are globally scalable, using an integrated product engineering methodology and a market-ready architecture. It explores how these components interact to produce intelligent systems that are capable of delivering value across diverse contexts and user expectations, thereby advancing the future of ethical, accessible, and impactful AI.

Methodology

Research design and framework development

The methodology adopted for this study is rooted in a multi-phase, design-based research (DBR) approach that combines qualitative user research, iterative system design, and quantitative validation. The overarching framework integrates principles from human-computer interaction (HCI), product engineering, and software architecture to formulate a scalable and human-centered AI platform. The development process was informed by existing models of integrated product engineering, where cross-functional teams—including AI developers, UX designers, data scientists, product managers, and target users—collaboratively shaped system requirements and functionalities. The objective was to ensure the AI platform could be scaled globally while maintaining relevance and adaptability across user demographics.

Designing human-centered AI platforms

To embed human-centered principles into the AI platform, the study employed ethnographic methods, user persona development, and usability testing. A diverse cohort of participants from various geographic, socio-economic, and professional backgrounds were interviewed and observed. These insights informed the creation of user personas and behavioral models that influenced the AI system's design. The interface, decision-making logic, and feedback loops were tailored to ensure inclusivity, trust, and transparency. Key features such as explainability (XAI), natural language processing (NLP) for multilingual support, and adaptive UI components were developed and tested through iterative prototyping cycles. Usability was assessed using System Usability Scale (SUS) scores and cognitive walkthroughs, with statistical analysis conducted to quantify satisfaction across different user segments.

Engineering scalable global software

The scalability of the AI platform was ensured through a combination of microservices architecture, containerized deployment, and API-first development. Simulations were run using synthetic and real-time datasets from multiple domains (e.g., healthcare, finance, logistics) to assess system performance across different scales. Load testing and stress testing were conducted using tools like Apache JMeter and Locust, and the resulting data were analyzed using descriptive statistics and performance metrics such as latency, throughput, and resource utilization. Scalability was benchmarked across three deployment environments: cloud (AWS, GCP), hybrid edge-cloud systems, and on-premise setups. Analysis of variance (ANOVA) was performed to identify significant differences in performance across deployment environments and use cases.

Integrated product engineering workflow

A central part of the methodology was the implementation of an integrated product engineering workflow, which synchronized product lifecycle phases—from concept design and system modeling to

market validation. Agile development cycles were used to iteratively refine the product, with sprint reviews focused on verifying alignment between technical output and user expectations. A requirements traceability matrix (RTM) was maintained to link user needs to engineering features and test outcomes. Design Thinking sprints enabled rapid ideation, while DevOps pipelines ensured continuous integration and delivery. Statistical tools such as Principal Component Analysis (PCA) and cluster analysis were used to reduce feature complexity and group product functionalities based on user relevance.

Validation and statistical analysis

Quantitative data were collected through structured surveys, performance benchmarks, and real-world pilot deployments. The effectiveness of the human-centered features was statistically evaluated using paired sample t-tests comparing pre- and post-deployment usability and satisfaction scores. Regression analysis was used to assess the relationship between architectural choices (e.g., modularity, interface types) and user adoption rates. Additionally, a Structural Equation Modeling (SEM) framework was applied to identify latent variables influencing trust and usability in AI systems.

Ethical considerations and bias mitigation

All user studies were conducted under institutional ethics clearance, with informed consent obtained from participants. To ensure fairness and reduce bias, fairness-aware ML models were tested, and demographic parity metrics were analyzed. Differential performance across demographic subgroups was statistically assessed using chi-square tests and fairness disparity indices.

Results

The findings from the study highlight the effectiveness and performance of the designed human-centered AI platform across usability, scalability, adoption, and fairness dimensions. Based on usability evaluations conducted with three distinct user groups—professional developers, citizen developers, and end users—the System Usability Scale (SUS) scores demonstrated consistently high levels of satisfaction. As presented in Table 1, professional developers reported the highest mean SUS score of 82.1 (SD = 5.0), followed by citizen developers (78.4, SD = 6.1) and end users (74.0, SD = 7.2). Notably, over 80% of all user categories rated their experience as satisfactory. These results are visually summarized in Figure 1, where the SUS scores with standard deviation error bars reflect the platform's usability consistency across diverse user types.

Table 1. Usability outcomes for human-centered AI platform (N = 180)

User Group	n	Mean SUS Score	SD	Satisfaction Rate ($\geq 80\%$)
Professional Developers	50	82.1	5.0	90 %
Citizen Developers	60	78.4	6.1	85 %
End Users	70	74.0	7.2	80 %

Scalability and performance benchmarking further supported the robustness of the system architecture under high demand. During peak load testing with 5,000 concurrent users, the cloud deployment achieved the highest throughput (2,100 transactions/second), though it exhibited the highest median latency (600 ms). Edge deployment offered the most balanced performance with a latency of 400 ms and throughput of 1,850 transactions/second. On-premise setups showed moderate results with higher latency (500 ms) and lower throughput (1,450 transactions/second), as shown in Table 2. These latency trends across user loads and environments are plotted in Figure 2, which illustrates the relative

performance efficiency of the edge environment under scaling demands. A one-way ANOVA confirmed significant differences in latency across deployment environments ($F(2, 12) = 48.7, p < 0.001$).

Table 2. Performance benchmarks under peak load (5 000 concurrent users)

Deployment environment	Throughput (tx s ⁻¹)	Median latency (ms)	CPU utilisation (%)
Cloud	2 100	600	78
Edge	1 850	400	71
On-Prem	1 450	500	82

One-way ANOVA for latency across environments: $F(2, 12) = 48.7, p < 0.001$.

To understand user adoption factors, a multiple regression analysis was conducted using architectural and design predictors. As indicated in Table 3, modularity of the architecture ($\beta = 0.42, p < 0.001$), transparency features ($\beta = 0.31, p < 0.001$), and local-language coverage ($\beta = 0.26, p = 0.004$) were all significant positive predictors of 30-day adoption rate, explaining 62% of the variance. Training-time efficiency, however, did not show a significant effect ($\beta = 0.07, p = 0.279$).

Table 3. Multiple-regression predicting 30-day user adoption ($R^2 = 0.62$)

Predictor (Standardised)	β	SE	t	p
Architectural Modularity	0.42	0.07	6.18	< 0.001
Transparency Index	0.31	0.08	3.88	< 0.001
Local-Language Coverage	0.26	0.09	2.93	0.004
Training-Time Efficiency	0.07	0.06	1.09	0.279

Finally, a fairness audit was carried out across demographic sub-groups to evaluate bias mitigation in the AI system. As shown in Table 4, all demographic groups except non-native-language speakers met the fairness threshold (equalized-odds $\Delta \leq 0.10$). The non-native-language user group showed a slightly higher disparity ($\Delta = 0.11$), suggesting a need for improvement in multilingual model tuning and interface localization. Overall, the results suggest that the human-centered AI platform is usable, scalable, and equitable, with minor adjustments required to optimize performance for all user groups.

Table 4. Fairness assessment across demographic sub-groups

Sub-group	Equalised-Odds Δ	Disparity Ratio	Pass/Fail (≤ 0.10)
Male	0.04	1.03	Pass
Female	0.06	0.97	Pass
High-income users	0.05	1.05	Pass
Low-income users	0.08	0.92	Pass
Non-native-language speakers	0.11	0.88	Fail




Figure 1: Mean SUS scores by user group




Figure 2: Latency vs Concurrent Users Across Deployment Environments

Discussion

Human-centered design enhances usability across user groups

The usability outcomes from this study underscore the strength of incorporating human-centered design into AI platform development. The high System Usability Scale (SUS) scores across all user groups, particularly professional and citizen developers, reflect a well-balanced interface that aligns with user expectations and cognitive models (Wang & Casciani, 2024). These results, as shown in Table 1 and Figure 1, validate the importance of early-stage user involvement, iterative design cycles, and contextual customization. The usability score for end users, though slightly lower, remains well above

average, indicating that the system is intuitive even for non-technical audiences. This suggests that design efforts focused on accessibility, including simplified workflows and adaptive UI components, were effective (Chrysikou et al., 2024). Future improvements can be directed at enhancing user onboarding and localization features to further support less tech-savvy users.

Scalability and deployment versatility across environments

Performance benchmarks reveal that the system is highly scalable and performs reliably across cloud, edge, and on-premise deployment environments. As indicated in Table 2 and Figure 2, the cloud deployment yielded the highest throughput, though with a trade-off in latency, particularly under heavy concurrent-user loads. The edge deployment proved most efficient in balancing latency and throughput, making it the optimal choice for real-time applications in bandwidth-constrained or latency-sensitive environments (Kleinschmidt et al., 2019). The significant differences in latency performance (ANOVA, $p < 0.001$) highlight the importance of architectural modularity and deployment strategy in global-scale software engineering. These findings support the notion that AI platforms intended for global use must be designed with deployment flexibility in mind, enabling tailored optimization based on infrastructure constraints (Petrik et al., 2019).

Architectural modularity drives adoption

The regression analysis offers valuable insights into the architectural and design factors influencing platform adoption. As shown in Table 3, architectural modularity emerged as the most powerful predictor of short-term adoption, underscoring its role in enabling scalability, adaptability, and system maintainability. Transparency features, such as explainable AI modules and audit trails, also significantly predicted user engagement, reinforcing the growing importance of trust and accountability in AI adoption (Ghassemi et al., 2024). Interestingly, local-language coverage had a measurable positive impact, confirming the utility of integrating language-specific NLP models and interface localization to serve diverse linguistic communities. On the other hand, training-time efficiency had no significant impact, suggesting that users prioritized transparency and adaptability over the underlying model optimization speed (Pentelovitch & Nagel, 2022).

Fairness audit reveals opportunities for improvement

The fairness evaluation, summarized in Table 4, generally supports the equity of the platform across most demographic groups, including gender and socioeconomic status. However, the system exhibited a mild disparity in its performance for non-native-language speakers. This gap, though marginal, highlights the complexity of achieving demographic parity in multilingual AI systems (Rotondo et al., 2025). While initial efforts in language model fine-tuning and UI adaptation proved largely successful, this result points to the need for more robust natural language processing (NLP) pipelines that consider regional dialects, mixed-language inputs, and culturally contextualized semantics (Xiao et al., 2025). Incorporating adaptive learning mechanisms that personalize system outputs based on user feedback could help close these gaps over time (Caddle et al., 2025).

Implications for global AI platform design

Collectively, the results affirm that a well-integrated, human-centered AI platform can be both globally scalable and locally relevant. The study demonstrates that combining design thinking with integrated product engineering leads to systems that not only perform under load but also resonate with diverse users (Delikostidis et al., 2016). The human-centric features transparency, inclusivity, and modularity enhance usability and trust, which in turn drive adoption and long-term sustainability. For AI to be

meaningfully embedded into global software ecosystems, developers must go beyond algorithmic accuracy and prioritize design strategies that reflect the realities of different user environments (Lee, 2025). This research offers a replicable blueprint for creating AI platforms that are not only technically sound but also ethically and socially responsive.

Conclusion

This study presents a comprehensive approach to designing human-centered AI platforms that are scalable, globally deployable, and grounded in integrated product engineering. The findings clearly demonstrate that usability, performance, and fairness can coexist when system architecture is thoughtfully aligned with human values and real-world constraints. High usability scores across user groups, robust performance under varying deployment environments, and statistically significant adoption predictors underscore the platform's effectiveness. Moreover, the modest disparity identified in fairness assessments points to areas for iterative refinement, particularly in multilingual inclusivity. By integrating transparency, modularity, and user-centered design principles from the earliest stages of development, the research offers a replicable and sustainable model for building AI systems that are not only technically sound but also socially responsible and market-ready for global adoption.

References

- [1] Bengler, K., Harbauer, C. M., & Fleischer, M. (2023). Exoskeletons: A challenge for development. *Wearable Technologies*, 4, e1.
- [2] Bermejo-Martín, G., & Rodríguez-Monroy, C. (2020). Design thinking methodology to achieve household engagement in urban water sustainability in the city of Huelva (Andalusia). *Water*, 12(7), 1943.
- [3] Bhatnagar, T., Bandukda, M., Bell, D., & Holloway, C. (2023, November). It's a Much Harder Journey: Scaling Assistive Technology Innovations to New Markets in Africa. In *Proceedings of the 4th African Human Computer Interaction Conference* (pp. 190-200).
- [4] Caddle, X. V., Qadir, S., Hughes, C., Sweigart, E. A., Park, J. K., & Wisniewski, P. J. (2025). Building a Village: A Multi-stakeholder Approach to Open Innovation and Shared Governance to Promote Youth Online Safety. *arXiv preprint arXiv:2504.03971*.
- [5] Chrysikou, E., Biddulph, J. P., Loizides, F., Savvopoulou, E., Rehn-Groenendijk, J., Jones, N., ... & Tziraki, C. (2024). Creating resilient smart homes with a heart: Sustainable, technologically advanced housing across the lifespan and frailty through inclusive design for people and their robots. *Sustainability*, 16(14), 5837.
- [6] Chrysikou, E., Biddulph, J. P., Loizides, F., Savvopoulou, E., Rehn-Groenendijk, J., Jones, N., ... & Tziraki, C. (2024). *Creating Resilient Smart Homes with a Heart: Sustainable, Technologically Advanced Housing across the Lifespan and Frailty through Inclusive Design for People and Their Robots*. *Sustainability* 2024, 16, 5837.
- [7] Ciano, C., Chrszon, P., Fischer, P. M., Amato, F., & Gerndt, A. (2025, April). Adaptive AI in Concurrent Engineering: A Paradigm Shift in Design and Integration. In *International Conference on Advanced Information Networking and Applications* (pp. 19-30). Cham: Springer Nature Switzerland.
- [8] Delikostidis, I., van Elzakker, C. P., & Kraak, M. J. (2016). Overcoming challenges in developing more usable pedestrian navigation systems. *Cartography and geographic information science*, 43(3), 189-207.
- [9] Ghassemi, N., Goldani, A., Whishaw, I. Q., & Mohajerani, M. H. (2024). Moo-ving Beyond Tradition: Revolutionizing Cattle Behavioural Phenotyping with Pose Estimation Techniques. *arXiv preprint arXiv:2408.06336*.

- [10] Hwang, B. N., Puntha, P., & Jitanugoon, S. (2025). AI-Driven Circular Transformation: Unlocking Sustainable Startup Success Through Co-Creation Dynamics in Circular Economy Ecosystems. *Sustainable Development*.
- [11] Kleinschmidt, S., Peters, C., & Leimeister, J. M. (2019). How to scale up contact-intensive services: ICT-enabled service innovation. *Journal of Service Management*, 31(4), 793-814.
- [12] Lee, S. M. (2025). Advances in shape memory alloy textiles for wearable applications: A critical review. *Journal of Industrial Textiles*, 55, 15280837251324899.
- [13] Li, L., & Duan, L. (2025). Human centric innovation at the heart of industry 5.0—exploring research challenges and opportunities. *International Journal of Production Research*, 1-33.
- [14] Meselhy, A., & Almalkawi, A. (2025). A review of artificial intelligence methodologies in computational automated generation of high performance floorplans. *npj Clean Energy*, 1(1), 2.
- [15] Pan, C., Banerjee, J. S., De, D., Sarigiannidis, P., Chakraborty, A., & Bhattacharyya, S. (2023, February). ChatGPT: A OpenAI platform for society 5.0. In *Doctoral Symposium on Human Centered Computing* (pp. 384-397). Singapore: Springer Nature Singapore.
- [16] Pentelovitch, N., & Nagel, J. K. (2022). Understanding the use of bio-inspired design tools by industry professionals. *Biomimetics*, 7(2), 63.
- [17] Petrik, D., Springer, V., Strobel, G., Möller, F., & Schoormann, T. (2024). The price is right: Exploring pricing of digital industrial platforms. *Information Systems Management*, 41(2), 151-180.
- [18] Roach, D. C. (2024). Systems of Innovation. In *The Innovation Approach: Overcoming the Limitations of Design Thinking and the Lean Startup* (pp. 3-21). Emerald Publishing Limited.
- [19] Rotondo, B., Bakker, C., Balkenende, R., & Arquilla, V. (2025). Integrating Circular Economy Principles in the New Product Development Process: A Systematic Literature Review and Classification of Available Circular Design Tools. *Sustainability*, 17(9), 4155.
- [20] Schiuma, G., & Santarsiero, F. (2024). Open Innovation Labs. *Springer Books*.
- [21] Suganya, R. V., Vetrivel, M., & Kumar, A. S. (2025). Cultivating Innovation Ecosystems to Drive Sustainable Market Success. In *Effective Marketing and Consumer Behavior Tactics for High-End Products* (pp. 295-342). IGI Global Scientific Publishing.
- [22] Wang, W., & Casciani, D. (2024, November). Bridging the Gaps: Exploring Interdisciplinary Design and Development Processes in E-Clothing. In *2024 International Conference on the Challenges, Opportunities, Innovations and Applications in Electronic Textiles (E-Textiles)* (pp. 98-108). IEEE.
- [23] Xiao, S., Li, J., Fushimi, T., & Ochiai, Y. (2025). Generative artificial intelligence-guided user studies: an application for air taxi services. *Behaviour & Information Technology*, 1-22.