

Article

What Makes AI Human-Centered? Identifying and Prioritizing the Attributes of Human-Centeredness: An Exploratory Study with Asia-Pacific Stakeholders

Aung Pyae 

International School of Engineering, Faculty of Engineering, Chulalongkorn University, Bangkok 10330, Thailand; aung.p@chula.ac.th

Abstract

Human-Centered AI (HCAI) has emerged as a guiding paradigm for designing AI systems that align with human values, needs, and well-being, yet the field lacks consensus on what constitutes human-centeredness. This study addresses that gap through a four-phase sequential mixed-methods design: (1) thematic analysis of 81 HCAI definitions from academic, institutional, and industry sources, yielding 78 keywords; (2) frequency-based statistical categorization; (3) expert evaluation producing a final inventory of 26 attributes; and (4) a cross-sectional survey ($N = 145$), predominantly drawn from the Asia-Pacific region (77.2%, with Myanmar, Singapore, and Thailand most represented), in which practitioners, academics, and students rated each attribute on a 7-point Likert scale, complemented by a reflexive thematic analysis of open-ended responses. The 26-item scale demonstrated excellent internal consistency. Trust, values, benefits, needs, and usability were rated most highly, while affective and cognitive attributes—emotions, behaviours, and empathy—were consistently rated lower, a pattern the qualitative data suggest reflects perceived intractability rather than indifference. Inter-attribute correlations revealed interpretable substructures, including an experience/usability cluster, an emotion/empathy cluster, and a participatory engagement cluster, while human control operated as a conceptually independent dimension. Five qualitative themes provided interpretive context: user needs and augmentation as design drivers, ethical foundations and value alignment, trust as a relational outcome contingent on transparency, the complexity of human experience as a design challenge, and structural barriers including corporate incentives, regulatory gaps, and resource constraints. In this predominantly Southeast Asian sample, all three stakeholder groups showed substantial agreement on which attributes matter most and least. The primary divergence ran between academics and students: academics assigned higher importance to participatory and process-oriented attributes, while students emphasized tangible outcomes. Practitioners occupied an intermediate position, with a distinctive emphasis on ethical values. These findings offer an empirically grounded vocabulary for human-centeredness, positioned as an exploratory foundation for future psychometric refinement, with implications for HCAI design practice, education, and cross-stakeholder dialogue.



Academic Editor: Jose María Merigo

Received: 22 March 2026

Revised: 2 May 2026

Accepted: 19 May 2026

Published: 26 June 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: human-centeredness; artificial intelligence; Asia-Pacific; human-centered artificial intelligence

1. Introduction

Recent advancements in artificial intelligence (AI) have accelerated the integration of AI-driven applications into virtually every domain of human activity—from healthcare

and education to transportation and creative work—necessitating a fundamental shift from purely technology-centered design toward approaches that foreground human experiences, values, and well-being [1]. In response, Human-Centered AI (HCAI) has emerged as a paradigm that prioritizes human welfare, aligns AI technologies with human values, and enhances the quality of human–AI interaction [2–4]. Drawing on the interdisciplinary traditions of AI, Human–Computer Interaction (HCI), Human-Centered Design (HCD), and User Experience Design, HCAI bridges computer science, cognitive science, and engineering to advocate for collaborative, ethical, and human-centric design practices [4–6]. Recent work has continued to expand the empirical and conceptual foundations of HCAI [1], underscoring the timeliness of consolidating its core attributes.

Despite the growing prominence of HCAI as a design philosophy and research agenda, a fundamental challenge persists: there is no consensus on what constitutes human-centeredness in AI, nor is there an empirically grounded framework that systematically identifies and organizes its core attributes [7,8]. Existing work tends to address isolated dimensions of HCAI—such as ethics [9,10], explainability [11], or trustworthiness [2]—without consolidating these into a shared, multi-attribute vocabulary that spans the full conceptual breadth of the construct. This fragmentation limits the field’s ability to develop coherent design standards, evaluation criteria, and educational curricula for HCAI.

The practical consequences of this conceptual fragmentation are significant. Without a shared vocabulary for human-centeredness, design teams lack a common framework for evaluating whether an AI system adequately addresses human needs, values, and experiences. Whereas traditional human–computer interaction offers well-established evaluation methods such as Nielsen’s heuristic evaluation [12], no analogous, broadly adopted framework exists for assessing the distinctive human-centered concerns of AI systems—including value alignment, autonomy preservation, and equitable benefit distribution. Evaluation criteria remain ad hoc, varying across organizations and projects. Moreover, recent governance instruments—including the NIST AI Risk Management Framework [13], the EU Artificial Intelligence Act [14], and the ISO/IEC 42001 standard [15]—articulate high-level requirements for trustworthy, human-centric AI but do not decompose human-centeredness into granular, assessable attributes that can bridge regulatory intent and design practice.

This paper addresses these gaps through two key objectives. First, it synthesizes existing definitions of HCAI from academic, institutional, and industry sources to conceptualize human-centeredness and extract its constituent attributes. Second, it empirically investigates how three stakeholder groups—AI practitioners, academics, and students—perceive and prioritize these attributes through both quantitative ratings and qualitative responses. The primary contribution is the development and mixed-methods empirical assessment of a 26-attribute inventory of human-centeredness, ranked by perceived importance and enriched by qualitative themes. We position this study as the exploratory first phase of a planned multi-study program. The research is guided by two questions:

- RQ1: How is human-centeredness conceptualized in HCAI based on insights from existing literature?
- RQ2: What are the perspectives of AI practitioners, academics, and students on understanding and prioritizing human-centeredness in HCAI?

We position the study as an exploratory first phase conducted with a predominantly Asia-Pacific stakeholder sample; cross-cultural replication is identified as a necessary next step. The paper makes four contributions that together distinguish it from prior HCAI scholarship. First, whereas existing frameworks operate at the level of broad dimensions—Shneiderman’s [2,3] two-dimensional model of human control and computer automation, Xu’s [4] tripartite structure of technology enhancement, human factors, and ethical design, Schmager et al.’s [8] decomposition into purpose, values, and properties, and Ozmen

Garibay et al.'s [7] six grand challenges—we derive a more granular 26-attribute inventory with operational definitions, providing a vocabulary at the construct level that can be linked to assessable design criteria. Second, where Amershi et al.'s [5] influential 18 guidelines operate at the human–AI interface level, our inventory operates at the construct level and is therefore complementary rather than overlapping. Third, prior frameworks are typically proposed without empirical assessment of how stakeholders prioritize their components; we provide such an assessment across three groups (practitioners, academics, students) and combine it with reflexive thematic analysis of open-ended responses. Fourth, we surface the human- vs. user- prefix distinction as an explicit organizing principle within the inventory, an aspect under-articulated in prior HCAI work despite its long history in the HCI/HCD literature.

The remainder of this paper is organized as follows. Section 2 reviews the related literature. Section 3 describes our four-phase mixed-methods research design. Section 4 presents the results. Section 5 discusses the implications, and Section 6 offers concluding remarks.

2. Related Work

2.1. Historical Foundations: From HCI to HCAI

The intellectual foundations of HCAI are rooted in decades of HCI research that has progressively emphasized human agency, usability, and collaborative design. Licklider's [16] seminal work on human–computer symbiosis and Engelbart's [17] vision of augmenting human intellect established the foundational premise that technology should amplify rather than supplant human capabilities. These early visions were subsequently formalized through the development of cognitive models of interaction [18], direct manipulation interfaces [19], and usability heuristics [12], which collectively established user-centric design as a core principle of computing. The ISO 9241-210 standard [20] codified a process-oriented approach to designing interactive systems, and Norman's [21] articulation of design principles extended human-centered design into everyday practice.

The resurgence of AI—driven by breakthroughs in deep learning [22] and the increasing deployment of machine learning systems in high-stakes domains—has brought renewed urgency to the question of how human-centered principles can and should be applied to intelligent systems. The emergence of dedicated research institutions, including Stanford's Human-Centered AI Institute and the AI Now Institute, signaled a field-wide recognition that AI development must be guided by considerations of human well-being, fairness, and accountability. This institutional momentum converged with growing scholarly attention to responsible and ethical AI [9,10], giving rise to HCAI as a distinct paradigm.

2.2. Defining Human-Centered AI: Perspectives and Frameworks

A defining challenge in the HCAI literature is the absence of a unified, widely accepted definition. Shneiderman [2,3] offers one of the most influential characterizations, framing HCAI through a two-dimensional model of human control and computer automation and arguing that well-designed AI systems can achieve high levels of both. Xu [4] proposed a complementary framework structured around technology enhancement, human factors, and ethical design. Xu and colleagues [6] subsequently elaborated this framework in the context of human–AI interaction, identifying unique challenges that distinguish AI systems from traditional computing. Amershi et al. [5] contributed 18 empirically derived design guidelines for human–AI interaction that operationalize HCAI principles at the interface level.

2.3. Governance, Standards, and Evaluation Frameworks

At a broader strategic level, Ozmen Garibay et al. [7] identified six grand challenges for HCAI spanning human well-being, responsible AI, labour and economy, privacy and security, environmental sustainability, and lifelong learning. The most recent systematic review by Schmager et al. [8] decomposes HCAI into three constituent elements—purpose, values, and properties—highlighting that despite growing convergence, the field still lacks consensus on the precise boundaries of the concept. This lack of definitional consensus is further evidenced by the proliferation of AI ethics guidelines worldwide. Jobin, Ienca, and Vayena [23] mapped the global landscape of such guidelines, identifying convergence around principles of transparency, justice, non-maleficence, responsibility, and privacy, but also noting significant divergence in how these principles are operationalized.

2.4. Human-Centeredness: The Core Construct

Within the broader HCAI discourse, human-centeredness represents the foundational organizing principle—the normative commitment that AI systems should be designed for, with, and accountable to human beings. While HCAI as a paradigm encompasses technical, organizational, and policy dimensions, human-centeredness specifically refers to the design orientation that prioritizes human values, needs, capabilities, and experiences [3,20]. Despite its foundational role, human-centeredness remains an under-specified construct. As Schmager et al. [8] observe, definitions vary considerably in scope and emphasis, complicating efforts to translate human-centeredness into actionable design requirements. The challenge is compounded by the interdisciplinary nature of the concept, which draws on traditions in HCI, participatory design, value-sensitive design [24], responsible AI [9], and human factors engineering.

A conceptual distinction that warrants explicit attention is the difference between ‘human-’ prefixed and ‘user-’ prefixed constructs within the HCAI vocabulary. Human-prefixed attributes (e.g., human dignity, human cognition, human emotions) refer to broader qualities, rights, and capacities of human beings as such—dimensions that are relevant regardless of whether an individual is directly interacting with an AI system. User-prefixed attributes (e.g., user needs, user involvement, user models) refer specifically to the person engaging with an AI system in a particular context of use. This distinction mirrors the long-standing difference between human-centered design, which foregrounds universal human values and capabilities [20], and user-centered design, which focuses on the interactive context and task requirements. Both levels of analysis are essential for a comprehensive characterization of human-centeredness in AI.

2.5. Stakeholder Perspectives in HCAI

A central tenet of human-centered design is that the perspectives of diverse stakeholders must be incorporated into the design process [20,24]. In the context of HCAI, three stakeholder groups are particularly relevant: practitioners, academics, and students. Practitioners are responsible for translating HCAI principles into working systems, and their perspectives are shaped by the practical constraints of deployment [5,25]. Academics contribute theoretical depth and methodological rigor, advancing the conceptual foundations upon which design frameworks are built. Students represent the emerging workforce that will shape future AI development. Comparative studies that systematically assess how these groups differ in their conceptualization and prioritization of human-centeredness attributes remain scarce. As Xu et al. [6] argue, the successful implementation of HCAI depends on alignment between those who theorize about human-centeredness, those who implement it, and those who will carry it forward.

2.6. Gaps Addressed by This Study

The foregoing review reveals four interrelated gaps that this study addresses. First, there is no empirically grounded, multi-attribute characterization of human-centeredness that decomposes the concept into specific, assessable dimensions. Second, existing frameworks are typically proposed without empirical assessment across diverse stakeholder groups. Third, comparative analyses of how different stakeholders prioritize the attributes of human-centeredness are largely absent. Fourth, current governance instruments [13–15] and HCAI design frameworks address human-centeredness at different levels of abstraction, with limited empirical work connecting regulatory requirements to assessable design attributes. The four contributions enumerated in Section 1 map directly onto these four gaps.

We do not claim that no prior work has examined human-centeredness in AI; on the contrary, the field has produced influential conceptual frameworks [2–4,7,8] and design guidelines [5], systematic reviews of HCAI definitions [8], and large-scale analyses of AI ethics principles [23]. What is, to our knowledge, absent from the existing literature is the specific combination of (a) systematic synthesis of HCAI definitions into a granular, operationally defined attribute inventory at the construct (rather than guideline or principle) level, (b) empirical assessment of how multiple stakeholder groups prioritize these attributes using both structured Likert ratings and reflexive thematic analysis of open-ended responses, and (c) a deliberate human-/user- prefix distinction integrated into the inventory itself. Schmager et al.'s [8] recent systematic review is the closest comparator: it synthesizes definitions and proposes a purpose–values–properties decomposition, but does not empirically assess stakeholder prioritization or operationalize the construct as a multi-item inventory. Ozmen Garibay et al. [7] articulate six grand challenges through expert consultation but do not generate a granular attribute set or assess stakeholder priorities quantitatively. Amershi et al. [5] empirically derive 18 guidelines but at the interaction level rather than the construct level, and without comparative stakeholder analysis. Our study is therefore positioned as complementary to—and bridging between—these strands of work.

3. Method

3.1. Research Design

This study adopted a sequential, multi-phase mixed-methods research design [26] to identify and empirically assess the core attributes of human-centeredness in AI systems. The research proceeded through four interconnected phases: (1) systematic definition collection and thematic analysis, (2) keyword extraction and statistical frequency categorization, (3) expert evaluation and keyword refinement, and (4) a cross-sectional survey incorporating both structured Likert-scale ratings and open-ended qualitative questions administered to practitioners, academics, and students. This design reflects the methodological integration advocated by Johnson and Onwuegbuzie [26], in which qualitative insights inform quantitative instrument development and quantitative findings are enriched by qualitative contextualization. We position this study as the first exploration phase of a planned multi-study program. The overview of the research design is presented in Figure 1.

3.2. Phase 1: Definition Collection, Thematic Analysis, and Keyword Extraction

3.2.1. Definition Collection

A systematic search was conducted across academic databases (e.g., ACM Digital Library), institutional publications, and authoritative web sources—including official organizational websites such as Stanford's Institute for Human-Centered AI (HAI) and the OECD AI Policy Observatory—to identify definitions of HCAI and human-centeredness

in AI. Search terms included combinations of ‘human-centered AI,’ ‘human-centred AI,’ ‘human-centeredness,’ and ‘human-centric artificial intelligence.’ Sources were included if they (a) provided an explicit definition or characterization of HCAI or human-centeredness in AI, (b) were published between 2017 and 2025, and (c) appeared in peer-reviewed journal articles, conference proceedings, government reports, standards documents, industry white papers, or official organizational web pages from recognized research institutions and international bodies. Sources that mentioned HCAI only in passing without offering a definition were excluded. For example, definitions were drawn from Shneiderman’s [2,3] foundational work on reliable, safe, and trustworthy HCAI and from Stanford HAI’s characterization of human-centered AI as technology designed to augment human capabilities. This process yielded a total of 81 unique definitions.




Figure 1. Four-phase mixed-methods research design. Phase 1 produced 78 unique keywords from inductive thematic analysis of 81 HCAI definitions (2017–2025). Phase 2 categorized these by frequency, retaining 15 high-frequency keywords and forwarding 63 mid-frequency keywords to expert evaluation. Phase 3 yielded a 26-attribute inventory through expert rating and three rounds of structured deliberation. Phase 4 administered a survey (N = 145) and analyzed responses through integrated quantitative and qualitative methods. Methodological details for each phase are described in Sections 3.2–3.5.

3.2.2. Thematic Analysis and Keyword Extraction

Each of the 81 definitions was segmented into sub-definition units and subjected to inductive thematic coding. Sub-definition texts coded under the ‘Human-Centeredness’ theme were extracted and organized into a dedicated dataset (Step 1), yielding 195 sub-definition texts. Each sub-definition was then tagged with keywords representing the

core concepts it expressed. The raw extraction produced 77 unique keywords. These were subsequently cleaned and consolidated (Step 2): singular and plural variants were merged (e.g., ‘human experiences’ was merged with ‘human experience’), terminology was standardized, and compound extractions were separated into individual terms where appropriate. This cleaning process produced a final set of 78 unique keywords representing the conceptual landscape of human-centeredness in AI.

3.3. Phase 2: Frequency Analysis and Statistical Categorisation

The frequency of occurrence of each of the 78 keywords across the sub-definition texts was tallied. The resulting distribution exhibited extreme positive skew: the median frequency was 1 (IQR = 1), with Q1 = 1 and Q3 = 2, indicating that the vast majority of keywords appeared only once or twice while a small number recurred with substantially higher frequency.

To distinguish empirically prominent keywords from the long tail of infrequent terms, a frequency threshold of three or more occurrences was adopted. This threshold was selected on both distributional and substantive grounds. Distributionally, keywords appearing three or more times constituted the top 19.2% of the distribution (15 of 78 keywords), representing a clear inflection point above which terms recurred across multiple independent definition sources. Substantively, this cutoff distinguished genuinely cross-cutting concepts from terms specific to individual definitions. Alternative threshold approaches were considered—including top-quartile selection, natural breaks, and Tukey’s [27] boxplot upper fence ($Q3 + 1.5 \times IQR = 3.5$, which would yield 11 keywords at four or more)—but the threshold of three offered the most defensible balance between inclusiveness and empirical selectivity. The 15 high-frequency keywords were retained for the final inventory; the remaining 63 keywords were classified as mid-frequency and contributed to the subsequent expert evaluation.

Of the 78 unique keywords, 54 (69.2%) appeared exactly once, 9 appeared twice, 5 appeared three times, 2 appeared four times, 3 appeared five times, and 5 appeared with substantially higher frequency—ranging from 7 to 34 occurrences. The five most prominent keywords were human needs ($f = 34$), human values ($f = 19$), human well-being ($f = 17$), human behaviours ($f = 12$), and human-centric design ($f = 7$). This long-tail distribution, in which a small core of high-frequency terms coexists with a large number of singleton mentions, reflects the conceptual breadth and terminological diversity of the HCAI literature.

3.4. Phase 3: Expert Evaluation and Keyword Refinement

3.4.1. Expert Panel and Evaluation Instrument

An expert evaluation was conducted on the 63 mid-frequency keywords. A panel of four domain experts was assembled to span four complementary perspectives on human-centered AI. The first panelist was a senior academic researcher specializing in HCI and human-centered AI; the second was an academic researcher focusing on AI ethics and responsible technology design; the third was an industry practitioner with substantive experience in UX design and AI product development; and the fourth was an educator specializing in computing and human–computer interaction curricula. All four panelists had at least ten years of experience in HCI and approximately three to five years of focused engagement with human-centered AI as a named paradigm—a profile consistent with the recency of HCAI’s crystallization as a distinct field built on much longer HCI foundations [2–4]. The panel was assembled as a structured deliberative body spanning these four complementary disciplinary perspectives rather than as a statistically representative sample of the HCAI community; population-level stakeholder assessment is provided

by the Phase 4 survey ($N = 145$, Section 3.5). A structured evaluation form (Step 3) was created containing all 63 keywords, each to be rated across seven evaluative dimensions—importance, relevancy, essentiality, human-centeredness, applicability, user empowerment, and future implication—using 7-point Likert scales.

3.4.2. Expert Ratings and Mean-Based Classification

Each panelist independently rated every keyword across all seven dimensions (Step 4). A composite score was computed for each keyword by averaging across all dimensions and panelists. The overall mean composite score across all 63 keywords was ($M = 5.71$). Keywords scoring at or above this threshold were classified as above-mean and retained for structured deliberation (Step 5: 37 keywords); the remaining keywords scored below the mean and were set aside (Step 6: 26 keywords). The above-mean and below-mean sets summed to the full set of 63 mid-frequency keywords ($37 + 26 = 63$).

3.4.3. Three-Round Expert Deliberation

Three rounds of structured discussion were conducted on the 37 above-mean keywords. In Round 1, the retained keywords were reviewed collectively and initial semantic groupings were proposed. In Round 2, each keyword was evaluated against relevance to current HCAI research, alignment with established frameworks [2,4], and capacity to encapsulate a distinct conceptual dimension. Of the 37 keywords, 20 were identified as candidate additions to the inventory.

In Round 3, these 20 candidates were combined with the 15 high-frequency keywords into a unified list of 35 keywords (Step 7) and scrutinized for redundancy, conceptual overlap, and completeness. This deliberation produced three types of adjustments. First, four expert-selected keywords were merged into broader constructs already represented in the inventory: ‘human thought processes’ was subsumed under Human Cognition, ‘human perspectives’ under Human Insights, ‘human intelligence’ under Human Factors, and ‘technology and human interaction converge’ was removed as too narrowly specific. Second, two keywords that had been initially excluded were reinstated as conceptually essential: Human-Centered Design Principles (which had been present in the high-frequency set but initially flagged as overlapping with Human-Centric Design) and Human Dignity (which the panel judged to be a non-negotiable attribute of human-centeredness despite its lower composite score). Third, one new attribute—User Needs—was distinguished from Human Needs as a conceptually distinct construct, recognizing that fundamental human requirements (safety, security, belonging) differ from the specific, empirically identified requirements of particular user groups in particular contexts of use.

These adjustments yielded a final inventory of 26 attributes. The inventory was subsequently reviewed and validated by the research group, including domain experts in HCAI, to confirm conceptual clarity, mutual distinctiveness, and alignment with the literature.

3.4.4. Operational Definition Development

The operational definitions for the 26 attributes were developed through a multi-stage process: (1) systematic re-examination of the 81 original definitions, (2) iterative team discussion and consensus-building, (3) alignment review against foundational HCAI [2,3,5,8], HCI [20,21], AI ethics [9,10,23], and governance [13,15] sources, and (4) external review by an independent expert with over 10 years of experience in HCI and extensive knowledge of human-centered AI research. Table 1 presents the 26 attributes and their operational definitions.

Table 1. The 26 attributes of human-centeredness and their operational definitions.

No.	Attribute	Operational Definition
1	Human-Centric Design	A design philosophy that prioritizes human values, needs, and experiences throughout the AI development lifecycle, aiming to create systems that enhance human capabilities while upholding ethical principles and fostering trust.
2	User Involvement	The active participation of end users in the design, development, and evaluation of AI systems, ensuring that their perspectives, feedback, and lived experiences inform design decisions at each stage.
3	Human Needs	The fundamental requirements—including safety, security, usability, and social belonging—that AI systems must address to support human well-being and meaningful engagement with technology [20].
4	User Needs	The specific functional, contextual, and experiential requirements of particular user groups that must be identified through empirical inquiry and translated into system features and interaction flows [20].
5	Human Values	The ethical and moral principles—such as fairness, autonomy, dignity, privacy, and justice—that AI systems must respect and uphold in their design, deployment, and operation [23,24].
6	Human Experience	The holistic quality of a person’s interaction with an AI system, encompassing perceptions, emotional responses, sense of meaning, and overall satisfaction across the full trajectory of use.
7	Usability	The extent to which an AI system enables users to achieve their goals with effectiveness, efficiency, and satisfaction in a specified context of use, following established usability engineering principles [12,20].
8	Human Factors	The scientific understanding of human cognitive, perceptual, and physical capabilities and limitations, applied to the design of AI systems to optimize performance, reduce error, and ensure safety [18].
9	Empathy	The capacity of an AI system’s design process and outputs to recognize, interpret, and respond appropriately to users’ emotional states, social contexts, and situational needs.
10	Human-Centered Design Principles	A codified set of design guidelines—including iterative prototyping, user feedback integration, accessibility, and ethical review—that structure the process of developing AI systems around human requirements [20,21].
11	User-Centric Data	The collection, management, and use of personal and behavioral data in ways that prioritize informed consent, minimize privacy risks, ensure data relevance, and give users meaningful control over their information.
12	Human Well-Being	The positive impact of AI systems on users’ physical health, psychological welfare, social connectedness, and overall quality of life, while actively mitigating potential harms.
13	Human Trust	Users’ confidence that an AI system will perform reliably, transparently, and in alignment with their expectations, developed through consistent, predictable, and accountable system behavior over time [2,3].

Table 1. *Cont.*

No.	Attribute	Operational Definition
14	Human Decision-Making	The support AI systems provide for users' informed, autonomous decision-making by presenting clear, relevant, and unbiased information while preserving the user's authority to make final judgments [2,3].
15	Feedback	The systematic processes through which user input, behavioral signals, and evaluative responses are gathered, analyzed, and incorporated into iterative improvements of AI systems.
16	Human Control	The ability of users to oversee, intervene in, adjust, and override AI system actions, ensuring that human authority is maintained over consequential decisions and automated processes [2,3].
17	Stakeholder Engagement	The inclusive involvement of diverse parties—including end users, developers, domain experts, affected communities, and policymakers—in the design and governance of AI systems to ensure broad accountability.
18	Human Cognition	The alignment of AI system design with human cognitive processes—including attention, memory, reasoning, and mental model formation—to reduce cognitive load and support effective information processing [18].
19	Human Behaviors	The observable patterns of user action, interaction, and adaptation that AI systems should accommodate and respond to, enabling personalized, context-sensitive, and intuitive system behavior.
20	User-Friendliness	The quality of an AI system that makes it accessible, learnable, and pleasant to use for people with varying levels of technical expertise, minimizing barriers to adoption and sustained engagement.
21	Human Insights	The deep, empirically grounded understanding of user needs, motivations, contexts, and latent requirements that emerges from systematic research and informs evidence-based AI design decisions.
22	Human Emotions	The affective states—such as frustration, satisfaction, anxiety, and delight—that users experience during AI interaction, which systems should recognize, respect, and avoid exacerbating.
23	Human Goals	The specific objectives, tasks, and desired outcomes that users bring to their interactions with AI systems, which should serve as the primary criteria for evaluating system effectiveness.
24	Human Benefits	The tangible and intangible positive outcomes that AI systems provide to individuals and communities, including enhanced productivity, improved access to services, and greater equity.
25	User Models	Computational representations of individual user characteristics, preferences, knowledge states, and behavioral patterns that enable AI systems to adapt and personalize interactions [18].
26	Human Dignity	The inherent worth and rights of every individual that AI systems must respect, ensuring that interactions do not diminish autonomy, perpetuate discrimination, or compromise personal integrity [24].

3.4.5. Note on Attribute Distinctiveness

Several attributes address related but non-identical constructs. Human Needs refers to fundamental requirements shared across all users (safety, security, belonging), whereas User Needs refers to the specific, empirically identified requirements of particular user groups in particular contexts. Human Experience captures the holistic quality of interaction

across the full trajectory of use, whereas Usability focuses on task effectiveness, efficiency, and satisfaction as defined in ISO 9241-210 [20]. Human Emotions concerns the affective states users experience during interaction, whereas Empathy concerns the system's design-level capacity to recognize and respond to those states. Human Factors addresses the broad scientific study of human capabilities and limitations [18], whereas Human Cognition focuses specifically on mental processes relevant to information processing. Human Values refers to broad ethical principles (fairness, justice, privacy) as mapped in the global AI ethics landscape [23,24], whereas Human Dignity refers specifically to the inherent worth and rights of individuals. These distinctions are maintained throughout the survey instrument to capture the multi-dimensional character of human-centeredness [2–4,8].

3.5. Phase 4: Survey Design and Administration

3.5.1. Instrument Development

A structured survey instrument was developed comprising three sections. Section 1 collected demographic and background information including age, gender, country of residence, current professional position, current working area, years of computing experience, years of AI-specific experience, and self-reported AI knowledge level. Section 2 presented each of the 26 attributes with its operational definition; respondents rated importance on a 7-point Likert scale (1 = Not at all important to 7 = Essential). Section 3 included three open-ended questions: (a) "In your opinion, what makes an AI system human-centered?"; (b) "What is the role of human-centeredness in an AI system?"; and (c) "What do you see as the biggest challenges or barriers to achieving truly human-centered AI systems, and how might these be overcome?" Section 2 ratings were subsequently analyzed using nonparametric statistics with effect sizes interpreted following Cohen [28], and Section 3 open-ended responses were analyzed using Braun and Clarke's reflexive thematic analysis [29]; the procedures for both are detailed in Sections 3.6 and 3.7.

3.5.2. Participant Recruitment

Participants were recruited through purposive sampling [30], targeting individuals with professional or educational engagement in AI, computing, or related fields. The survey was distributed through professional email networks, LinkedIn, and messaging applications. No compensation was provided. We did not pursue probability-based representative sampling. The study is positioned as the exploratory first phase of a multi-study program (see Section 5.9), and the appropriate inferential targets at this stage are (a) internal-consistency reliability of the 26-item inventory, (b) detection of large between-group differences in attribute prioritization, and (c) saturation of qualitative themes—not population-level point estimates of attribute importance. The achieved sample ($N = 145$) exceeds the conventional rule-of-thumb minimum of 5 respondents per item for exploratory psychometric work [31] and supports detection of medium-to-large group differences in nonparametric tests at $\alpha = 0.05$ with adequate power. Confirmatory factor analysis, measurement invariance testing, and population-representative sampling are reserved for the subsequent confirmatory phase.

A total of 145 usable responses were obtained. Respondents spanned 21 countries (after standardizing spelling variants), with the largest representations from Myanmar (46.9%), Singapore (14.5%), and Thailand (12.4%), reflecting a predominantly Asia-Pacific sample. The mean age was 33.4 years ($SD = 11.3$, $N = 143$ for age statistics due to two non-numeric responses), with a range from 20 to 71. Gender distribution was approximately balanced (73 male, 70 female, 2 preferred not to say). Self-reported AI knowledge was distributed across no knowledge (2.1%), basic (35.9%), intermediate (39.3%), advanced (18.6%), and expert (4.1%).

3.5.3. Stakeholder Classification

Respondents were classified into three groups based on a combined analysis of their self-reported current position and working area. Those identifying as students (by position or sector) were classified accordingly. Respondents listing academic titles (researcher, lecturer, professor) or reporting their working area as including academia were classified as academics. The remaining respondents were classified as industry experts and practitioners.

3.6. Quantitative Data Analysis

Cronbach's alpha (α) was computed for the 26-item scale across the full sample and within each group to assess internal consistency. Means, standard deviations, medians, and coefficients of variation were computed for each attribute. Given the ordinal nature of Likert-scale data and unequal group sizes, nonparametric inferential statistics were employed [28]. The Kruskal–Wallis H test was used for between-group comparisons on each attribute, with follow-up pairwise Mann–Whitney U tests for significant results. Friedman's test assessed within-subject priority structures, and Spearman rank-order correlations (ρ) measured cross-group agreement on attribute ordering. A significance threshold of $p < 0.05$ was adopted, with effect sizes interpreted following Cohen [28].

3.7. Qualitative Data Analysis

Open-ended responses were analyzed using Braun and Clarke's [29] six-phase reflexive thematic analysis. Of the 145 survey respondents, 144 provided at least one qualitative response. After reviewing for substantive content, 112 valid responses were retained for question (a), 109 for question (b), and 115 for question (c). The analysis proceeded through familiarization, systematic code generation, theme collation, theme review, refinement, and production of the analytic narrative. Consistent with the reflexive approach, coding is treated as an interpretive act shaped by the researchers' theoretical commitments rather than as a reliability exercise requiring inter-rater agreement metrics [29]. Two researchers independently coded a randomly selected subset of 30 responses as a reflexivity and calibration exercise; disagreements were resolved through discussion.

3.8. Ethical Considerations

The study was conducted in accordance with the ethical guidelines of the host institution. Prior to participation, informed consent was obtained from all participants, who received clear explanations of the study's purpose, procedures, expected duration, and their rights, including the right to withdraw at any time without consequence. No compensation was offered for participation. Participants' confidentiality and anonymity were ensured throughout the study, and all collected data were handled securely and used solely for research purposes.

4. Results

4.1. Participants

A total of 145 respondents completed the survey with no missing data across any variable. Participants were classified into three stakeholder groups: industry experts and practitioners ($n = 67$, 46.2%), students ($n = 47$, 32.4%), and academics ($n = 31$, 21.4%). The sample achieved near-parity in gender representation (50.3% male, 48.3% female, 1.4% undisclosed).

Respondents represented 21 countries, though the distribution was concentrated in Southeast Asia (77.2%), particularly Myanmar (46.9%), Singapore (14.5%), and Thailand (12.4%). The remaining participants came from Europe (5.5%), the Americas (4.8%), East Asia (4.1%), Oceania (4.1%), South Asia (3.4%), and Africa (0.7%). Myanmar-based re-

spondents were distributed across all three groups (53.2% of students, 22.6% of academics, 53.7% of industry), while Singapore contributed disproportionately to the academic (22.6%) and industry (20.9%) subsamples. This geographic concentration means the findings predominantly reflect Southeast Asian perspectives on human-centered AI. Accordingly, the attribute rankings, group comparisons, and qualitative themes reported below should be interpreted as Asia-Pacific stakeholder perspectives on human-centeredness in AI; we return to the implications of this scope in Section 5.9.

The three groups differed substantially in age, experience, and self-assessed AI knowledge. Students were predominantly aged 18–25 (70.2%), with 53.2% reporting 0–1 year of AI experience and 55.4% reporting basic or no AI knowledge. Academics had the most mature profile: 41.9% were aged 46–55, 64.5% had 16+ years of computing experience, and 45.2% reported advanced or expert-level AI knowledge. Industry professionals clustered in mid-career age ranges (26–45), with substantial computing backgrounds but comparatively limited AI-specific experience—40.3% reported only 0–1 year of AI experience despite many having 6+ years in computing. This computing-to-AI experience gap is a defining characteristic of the subsample of industry. Table 2 summarizes key demographic characteristics.

Table 2. Participant profile by respondent group.

Characteristic	Student (n = 47)	Academia (n = 31)	Industry (n = 67)	Total (n = 145)
Gender: Female	46.8%	61.3%	43.3%	48.3%
Gender: Male	53.2%	35.5%	55.2%	50.3%
Mean age (years)	~24	~48	~34	33.4
Computing exp. 16+ yr	6.4%	64.5%	23.9%	26.9%
AI exp. 0–1 yr	53.2%	12.9%	40.3%	38.6%
AI exp. 2–5 yr	42.6%	45.2%	53.7%	48.3%
AI knowledge: Advanced/Expert	12.7%	45.2%	19.4%	22.8%
AI knowledge: Basic/None	55.4%	12.9%	38.8%	37.9%

Across the full sample, 86.9% of respondents reported five or fewer years of AI experience, reflecting the recency of widespread AI adoption. The most common professional roles were student (31.7%), researcher/lecturer/professor (16.6%), software developer (12.4%), project manager (6.9%), and machine learning developer (6.2%). The presence of UX/UI designers (4.1%), though modest, is particularly relevant to a study on human-centeredness.

4.2. Attribute Importance Rankings

The 26-item scale demonstrated excellent internal consistency ($\alpha = 0.929$, $N = 145$), with sub-group coefficients equally robust: Students ($\alpha = 0.917$), Academia ($\alpha = 0.943$), and Industry/Expert ($\alpha = 0.925$). This reflects internal consistency; construct validity through factor analysis is deferred to the next phase of this research program. Table 3 presents the 26 attributes ranked by mean importance. On the 7-point scale (1 = Not at all important, 7 = Essential), the grand mean across 3770 ratings was 5.29 (SD = 1.28). Means ranged from 5.83 (Human trust) to 4.53 (Human emotions)—a spread of 1.30 points. All 26 means exceeded the scale midpoint of 4.0, indicating that every attribute was considered meaningfully important. The ranking reflects relative prioritization among valued attributes, not a distinction between important and unimportant ones.

The highest-ranked attributes emphasize foundational principles and outcomes—trust, values, benefits, needs, and usability. Human trust and Human values tied for the highest endorsement rate (57.2% rating 6 or 7) with the lowest variability (CV = 19.3% and 20.5%). The lowest-ranked attributes relate to cognitive, affective, and behavioural dimensions.

Human emotions was the most divisive attribute: highest CV (31.2%), highest proportion of low ratings (22.1%), and the only median of 4.

Table 3. Descriptive statistics for 26 human-centeredness attributes (N = 145).

Rank	Attribute	M	SD	Mdn	CV%	Top-2 %
1	Human trust	5.83	1.13	6	19.3	57.2
2	Human values	5.71	1.17	6	20.5	57.2
3	Human benefits	5.67	1.12	6	19.7	53.8
4	Human needs	5.64	1.24	6	22.0	54.5
5	Human-centric design	5.61	1.21	5	21.6	49.0
6	Usability	5.57	1.11	6	19.9	51.0
7	User-centric data	5.54	1.27	6	22.9	51.7
8	HC design principles	5.52	1.23	6	22.2	50.3
9	Human control	5.43	1.45	6	26.7	51.0
10	User needs	5.39	1.13	5	21.0	43.4
11	User-friendliness	5.39	1.21	5	22.5	45.5
12	Human decision-making	5.35	1.21	5	22.6	40.0
13	Human dignity	5.34	1.42	5	26.6	44.8
14	Human goals	5.32	1.17	5	22.0	41.4
15	Feedback	5.30	1.28	5	24.2	44.8
16	Human insights	5.21	1.15	5	22.1	40.7
17	Human well-being	5.21	1.37	5	26.4	40.7
18	User involvement	5.19	1.22	5	23.4	32.4
19	Human experience	5.17	1.22	5	23.6	38.6
20	Human factors	5.15	1.16	5	22.6	33.1
21	Stakeholder engagement	4.97	1.27	5	25.6	33.8
22	Human cognition	4.97	1.18	5	23.8	32.4
23	Empathy	4.94	1.40	5	28.4	33.1
24	Human behaviours	4.82	1.30	5	27.0	29.7
25	User models	4.77	1.15	5	24.1	22.1
26	Human emotions	4.53	1.41	4	31.2	26.2

Note. Scale: 1 = Not at all important; 7 = Essential. M = Mean; SD = Standard Deviation; Mdn = Median; CV% = Coefficient of Variation; Top-2 % = percentage rating 6 or 7.

The strongest inter-attribute correlations (Table 4) suggest latent substructures: an experience/usability factor, an emotion/empathy factor, a benefit/goal factor, and a participatory engagement factor. Notably, Human control showed near-zero correlations with User-friendliness ($r = 0.047$) and Human emotions ($r = 0.081$), suggesting it operates as a conceptually independent dimension. These clusters are consistent with the multi-dimensional structure of the inventory, in which conceptually related attributes (e.g., Human insights and Human emotions, both reflecting understanding of users' internal states) correlate more strongly with each other than with functionally distinct attributes (e.g., Human control). Formal factor analysis in future work can test whether these clusters represent stable latent dimensions.

Table 4. Strongest inter-attribute correlations.

Attribute Pair	r
Human experience ↔ Usability	0.613
Human insights ↔ Human emotions	0.613
Empathy ↔ Human emotions	0.582
Human goals ↔ Human benefits	0.574
User-friendliness ↔ Human insights	0.557
User involvement ↔ Stakeholder engagement	0.533

Between-Group Patterns

Per-respondent averages across all 26 attributes revealed a systematic gradient: academics rated highest overall ($M = 5.50$, $SD = 0.71$), followed by industry professionals ($M = 5.37$, $SD = 0.74$), then students ($M = 5.03$, $SD = 0.72$). This ordering held across the full attribute set, indicating a generalized elevation rather than selective emphasis on particular items. Each group nonetheless exhibited a distinctive priority structure (Table 5).

Table 5. Top 5 attributes by respondent group.

Rank	Student (n = 47)	Academia (n = 31)	Industry/Expert (n = 67)
1	Human Trust (5.92)	Human Needs (5.97)	Human Values (5.94)
2	Human Benefits (5.49)	User Needs (5.90)	Human Benefits (5.84)
3	Usability (5.43)	Human Values (5.90)	User-Centric Data (5.75)
4	Human Goals (5.38)	HC Design Principles (5.87)	Human Trust (5.75)
5	Human Control (5.36)	Human Trust (5.87)	Human-Centric Design (5.73)

Note. Values are group means on the 1–7 scale.

Human trust was the only attribute appearing in all three groups' top five, confirming its universal importance. Human values appeared in the top five for both academics and industry but not students, suggesting that value alignment becomes a more prominent concern with professional experience. Conversely, all three groups placed Human emotions and User models in their bottom five, indicating cross-group agreement on which attributes are least central.

Kruskal–Wallis H-tests identified significant between-group differences ($p < 0.05$) for 9 of 26 attributes (34.6%); the remaining 17 (65.4%) showed no significant variation. In every significant case, the same directional pattern held: academics rated highest, industry occupied the middle, and students rated lowest (Table 6).

Table 6. Kruskal–Wallis results for attributes with significant group differences ($df = 2$).

Attribute	H	<i>p</i>	Student	Acad.	Industry	Δ (Acad – Stu)
User needs	13.51	0.001	4.96	5.90	5.46	+0.95
User involvement	12.68	0.002	4.81	5.77	5.19	+0.97
Stakeholder engagement	10.70	0.005	4.64	5.55	4.94	+0.91
Human values	10.44	0.005	5.26	5.90	5.94	+0.65
Human experience	8.56	0.014	4.85	5.68	5.16	+0.83
Human well-being	8.45	0.015	4.72	5.55	5.39	+0.83
Human decision-making	7.55	0.023	4.98	5.71	5.45	+0.73
HC design principles	7.05	0.030	5.13	5.87	5.64	+0.74
Human factors	6.34	0.042	4.83	5.55	5.19	+0.72

Note. Group means on the 1–7 scale. Δ = Academia mean minus Student mean. All $p < 0.05$.

The largest divergences appeared on participatory and process-oriented attributes: User involvement ($\Delta = 0.97$), User needs ($\Delta = 0.95$), and Stakeholder engagement ($\Delta = 0.91$)—nearly full-point differences on a 7-point scale. Follow-up Mann–Whitney U tests confirmed that the Student–Academia comparison drove the largest differences in all nine cases (all $p < 0.01$ for the top three). The Student–Industry gap was significant for five attributes (User needs, Human values, HC design principles, Human well-being, Human decision-making), and the Academia–Industry gap reached significance for four (User involvement, User needs, Human experience, Stakeholder engagement).

4.3. Qualitative Findings

Reflexive thematic analysis of three open-ended questions yielded five themes (Table 7). Of 144 qualitative respondents, 112 (77.8%) responded to Q1 (what makes AI

human-centred), 109 (75.7%) to Q2 (the role of human-centeredness), and 115 (79.9%) to Q3 (challenges and barriers). No systematic non-response bias was detected. The distribution of themes across questions reflects the different foci of each question rather than a designed mapping; challenges-oriented themes (Themes 4 and 5) naturally emerged primarily from Q3, which asked specifically about barriers.

Table 7. Qualitative themes.

	Theme	Description	Reach
1	User needs and augmentation	AI should prioritize user requirements and enhance rather than replace human capabilities.	Q1, Q2
2	Ethical foundations and value alignment	Ethical principles—fairness, responsibility, rights—were framed as prerequisites for HCAI.	Q1, Q2
3	Trust through transparency	Trust was seen as contingent on system transparency, interpretability, and human oversight.	Q1, Q2, Q3
4	Complexity of human experience	Emotional, cultural, and behavioural complexity were identified as the hardest dimensions for AI.	Q3
5	Structural and systemic barriers	Corporate incentives, regulatory gaps, cost, and education deficits constrain HCAI.	Q3

Note. Q1 = what makes AI human-centered; Q2 = role of human-centeredness; Q3 = challenges/barriers.

4.3.1. Theme 1: User Needs and Augmentation

User/human needs was the most cited theme in Q1 (34.8%) and Q2 (22.9%), while augmentation—AI should enhance rather than replace human capabilities—was third in both (Q1: 20.5%; Q2: 18.3%). Respondents emphasized that “technology should be adapted to need . . . not the other way around” [R119]. Students foregrounded needs alongside ethics and augmentation (each 25.7%); academics showed the strongest emphasis on human-centered design as a formal methodology (29.2%); industry professionals exhibited the broadest spread, with no single theme exceeding 34%.

4.3.2. Theme 2: Ethical Foundations

Ethics was the second most cited Q1 theme (22.3%) and appeared as both a role (Q2: 16.5%) and a leading challenge (Q3: 20.0%). Students and industry cited ethics prominently (25.7% and 24.5% in Q1), while academics cited it less as a standalone theme (12.5%)—likely subsuming it within their design methodology framework. Industry showed the strongest emphasis on ethics as a challenge (22.2% in Q3).

4.3.3. Theme 3: Trust Through Transparency

Trust followed a distinctive trajectory: rare in definitions (Q1: 7.1%), more prominent as a functional role (Q2: 15.6%), and significant as a challenge (Q3: 13.9%). Respondents conceptualized trust as contingent on transparency and explainability rather than as a standalone property. One participant stated that AI “will be a deciding factor whether or not AI will be overall beneficial or downright bad for humanity” [R51]. Students showed particular concern with trust as a barrier (22.9% in Q3).

4.3.4. Theme 4: Complexity of Human Experience

The complexity of human nature was the third most cited challenge (Q3: 18.3%), and the emotion/empathy challenge appeared in 11.3% of Q3 responses. One respondent noted that “to understand and simulate human emotion might be the most difficult challenge” [R2]. Industry professionals showed the highest emphasis on this challenge (16.7%), reflecting the operational difficulty of modelling human complexity.

4.3.5. Theme 5: Structural Barriers

Regulation/governance (14.8%), cost/resources (13.9%), education gaps (13.0%), and corporate profit priorities (8.7%) were identified as systemic constraints operating beyond design methods. Industry emphasized regulation (18.5%); academics uniquely highlighted cost and education as top-tier challenges (each 15.4%). One respondent observed that “current AI development is driven mostly by large companies for whom profit is the primary incentive” [R119].

4.4. Convergence Between Quantitative and Qualitative Findings

The quantitative attribute rankings and qualitative themes converge substantively, strengthening confidence in the overall findings. This section traces the alignment across four dimensions.

4.4.1. Needs and Ethics

The qualitative centrality of user needs—the most frequently cited theme across Q1 (34.8%) and Q2 (22.9%)—aligns with the quantitative finding that Human needs ($M = 5.64$), Human benefits ($M = 5.67$), and Usability ($M = 5.57$) all ranked within the top eight attributes. Similarly, the qualitative prominence of ethical foundations corresponds to Human values ranking second overall ($M = 5.71$) with the joint-highest endorsement rate (57.2%). The constructs participants rated most important in the structured survey are the same ones they articulated most readily in open-ended responses.

4.4.2. Trust as Outcome, Not Definition

Human trust received the highest quantitative mean ($M = 5.83$), yet the qualitative data reveal a nuanced picture. Respondents rarely defined human-centered AI through trust (Q1: 7.1%); instead, they described trust as something human-centeredness produces (Q2: 15.6%) and as something difficult to achieve (Q3: 13.9%). The quantitative rating captures trust’s perceived importance: the qualitative data clarify its conceptual role as a relational outcome contingent on transparency and explainability—not a design input.

4.4.3. Affective Attributes: Intractability, Not Indifference

Human emotions ranked last quantitatively ($M = 4.53$) with the highest variability ($CV = 31.2\%$). The qualitative data suggest this low ranking reflects perceived intractability rather than dismissal. Respondents acknowledged the importance of emotional understanding while identifying it as “the most difficult challenge” [R2]. Industry professionals showed the highest qualitative emphasis on the emotion/empathy barrier (16.7%). The quantitative and qualitative data thus converge on a shared interpretation: affective dimensions are seen as important but currently beyond AI’s reliable capacity.

4.4.4. The Implementation Gap

Theme 5 surfaces systemic factors—corporate incentives, regulatory gaps, resource constraints, education deficits—that lie entirely outside the 26-attribute framework. These structural barriers were not measured by the quantitative instrument but were spontaneously identified by participants as significant constraints. This reveals a gap between

recognizing the importance of human-centeredness attributes and the practical difficulty of implementing them. The demographic profile reinforces this interpretation: the computing-to-AI experience gap among industry professionals (40.3% with one year or less of AI experience) suggests that many practitioners bring technological maturity but limited exposure to the human-centered AI principles that academics more readily articulate.

5. Discussion

This study sets out to address how human-centeredness is conceptualized in HCAI (RQ1) and how practitioners, academics, and students perceive and prioritize its attributes (RQ2). Through a four-phase mixed-methods design, we derived a 26-attribute inventory and assessed it empirically across three stakeholder groups in a predominantly Asia-Pacific sample, complemented by reflexive thematic analysis. In this section, we discuss the findings in relation to existing HCAI frameworks, the patterns of stakeholder convergence and divergence, design implications, and the exploratory boundaries of the work.

5.1. Human-Centeredness as a Multi-Dimensional Construct

The thematic analysis of 81 definitions and the subsequent extraction of 78 keywords confirm that human-centeredness in AI resists reduction to any single principle within the conceptual landscape covered by the source definitions, as perceived by the Asia-Pacific stakeholders surveyed. The conceptual landscape spans ethical commitments, functional design properties, experiential qualities, and affective dimensions—corroborating the breadth identified in recent reviews [4,8] and reinforced by ongoing analyses of how AI ethics principles translate into practice [32,33]. Our 26-attribute inventory offers a more granular operationalization than existing high-level frameworks. Where Shneiderman [2,3] foregrounds two dimensions (human control and computer automation) and Xu [4] proposes a tripartite structure (technology enhancement, human factors, ethical design), our inventory decomposes these broad categories into specific, nameable attributes that can serve as a shared vocabulary for design teams, educators, and researchers. It integrates constructs from ethical AI and value-sensitive design [9,24], classical HCI and usability engineering [12,20], and affective computing—responding to longstanding calls for interdisciplinary approaches to HCAI [4,7].

The inter-attribute correlation structure reveals that these 26 attributes are not independent atoms but cluster into interpretable groupings—an experience/usability cluster, an emotion/empathy cluster, a benefit/goal cluster, and a participatory engagement cluster. At the same time, attributes such as human control showed near-zero correlations with affective constructs, suggesting that it operates as a conceptually independent dimension. This pattern is consistent with the multi-faceted nature of the construct: human-centeredness encompasses both rights-based concerns (control, dignity) and experiential concerns (emotions, empathy) that, while both essential, may be psychometrically separable. Future factor-analytic work can formally test whether these emergent clusters represent latent dimensions of human-centeredness.

5.2. Positioning Within Established HCAI Frameworks

Building on this multi-dimensional view, we now position the inventory against the most influential existing HCAI frameworks. Shneiderman's [2,3] emphasis on human control and trustworthiness finds empirical support: both human trust and human control ranked highly in our sample, confirming that stakeholders regard these as foundational to human-centered AI. However, the inventory extends Shneiderman's model by revealing that control and trust, while important, do not exhaust the construct; attributes

spanning ethics, well-being, cognition, and affect constitute additional dimensions that a two-dimensional model does not capture.

Xu's [4] tripartite framework—technology enhancement, human factors, and ethical design—finds clear correspondence in the ranked attribute structure. Design-oriented attributes (usability, human-centric design, HC design principles) ranked highly, as did ethical foundations (human values, human trust), while human factors and experiential attributes occupied the lower portion of the ranking. Schmager et al.'s [8] decomposition of HCAI into purpose, values, and properties finds empirical content across our ranked attributes—purpose in human benefits and human goals, values in human values and human dignity, properties in usability and feedback—while the ranked structure adds a prioritization layer that their taxonomy does not provide. Amershi et al.'s [5] 18 design guidelines for human–AI interaction operate at the interface level, whereas our inventory operates at the construct level; the two are complementary. More recent applications of HCAI as an evaluative framework in domain-specific contexts [34] further reinforce the need for empirically grounded, multi-attribute characterizations of human-centeredness that can be adapted to particular deployment settings.

5.3. *The Primacy of Trust, Values, and Needs*

Beyond positioning the inventory against prior frameworks, the data reveal a clear priority structure within it. Across both quantitative and qualitative data, trust, values, and needs emerged as the most central dimensions of human-centeredness. Human trust received the highest mean importance rating, yet the qualitative findings revealed a nuanced picture: respondents rarely defined human-centered AI through trust; instead, they described trust as something human-centeredness produces and as something difficult to achieve. This relational conceptualization of trust—as contingent on transparency, explainability, and meaningful oversight rather than as a standalone design property—resonates with the sociotechnical framing of human-centered explainable AI [11] and aligns with Shneiderman's [3] argument that trustworthiness requires simultaneous high levels of human control and system capability.

The convergence between quantitative primacy and qualitative salience strengthens interpretive confidence. The constructs that respondents rated most highly in the structured survey (trust, values, needs, benefits) were the same ones they articulated most readily in open-ended responses. This alignment, observed across methods and not prompted by the survey structure, suggests that these attributes occupy a genuinely central position in stakeholders' mental models of what makes AI human-centered. The finding aligns with the broad consensus around trust, transparency, and ethical alignment observed in global AI ethics guideline analyses [23].

5.4. *The Affective Attribute Gap: Intractability, Not Indifference*

If trust, values, and needs anchor the top of the priority structure, the bottom of the structure tells an equally informative story. Human emotions ranked lowest across all stakeholder groups, with the highest variability of any attribute. The qualitative data provide a corrective to any reading of this as indicating unimportance: respondents repeatedly identified emotional understanding as among the most difficult frontiers for AI, with industry professionals—who would face this challenge most directly in product development—showing the strongest emphasis on the emotion/empathy barrier. The quantitative and qualitative data converge on a shared interpretation: affective attributes are perceived as important but currently beyond AI's reliable capacity.

This intractability interpretation carries implications for research investment. As AI systems are increasingly deployed in emotionally consequential contexts—mental health

support, elder care, educational tutoring—the gap between recognized importance and perceived feasibility will become increasingly urgent. The current low prioritization may reflect rational assessment of present-day capabilities rather than a stable valuation, and future iterations of this inventory should track whether affective attributes rise in the priority structure as the technology matures.

5.5. Stakeholder Convergence and the Academic–Student Divergence

A notable finding is the substantial convergence across all three stakeholder groups on which attributes matter most and least. This broad agreement suggests that, at least within the Asia-Pacific context studied, practitioners, academics, and students share a language for identifying the priorities of human-centered AI. This convergence is encouraging for the development of shared design standards and educational curricula.

The primary axis of divergence ran between academics and students. Academics assigned higher importance to participatory and process-oriented attributes—user involvement, user needs, stakeholder engagement—and showed the strongest emphasis on human-centered design as a formal methodology, likely reflecting their deep engagement with participatory methodologies and the theoretical foundations of HCI [20,24]. Students showed a more pragmatic orientation emphasizing trust, benefits, usability, and control—attributes that are directly observable in user-facing systems. This pattern likely reflects distinct intellectual socialization rather than fundamental disagreement about what human-centeredness means.

These findings should be interpreted with appropriate caution given the small academic sub-group ($n = 31$). Nonetheless, the pattern carries potential implications for HCAI education. If students entering the AI workforce do not yet appreciate the importance of participatory design processes and well-being considerations, targeted curricular interventions may be warranted. The 26-attribute inventory itself may serve as a pedagogical resource for discussing human-centeredness across its full breadth.

5.6. Practitioners as Mediators

Practitioners occupied an intermediate position between academics and students on most attributes. Their distinctive emphasis on ethical values—placing human values at the top of their priority structure—is noteworthy given ongoing debates about the gap between AI ethics principles and practice. Prior work has documented the challenges practitioners face in operationalizing ethical commitments within organizational and technical constraints [25]. Our finding suggests a degree of attitudinal internalization of ethical principles among the practitioners surveyed. Practitioners may serve as natural mediators in cross-stakeholder dialogues, translating between the more theoretically grounded perspective of academics and the more outcome-oriented perspective of students and junior professionals.

5.7. Structural Barriers and the Implementation Gap

The qualitative findings surfaced a critical dimension not captured by the attribute inventory: structural and systemic barriers to HCAI realization. Participants identified corporate profit incentives, regulatory gaps, resource asymmetries, cost pressures, and education deficits as constraints operating beyond design methods. Industry professionals emphasized regulatory challenges; academics uniquely highlighted cost and education as leading barriers. This implementation gap represents a significant finding. This pattern aligns with recent analyses documenting the persistent gap between AI ethics principles and operational practice, which identify practitioner awareness, organizational constraints, and the absence of concrete implementation tools as recurring barriers [32,33].

The structural barriers theme suggests potential connections between the attribute inventory and governance frameworks such as the NIST AI Risk Management Framework [13] and the EU AI Act [14]. The high importance attributed to trust, control, and transparency in our sample resonates with the EU AI Act's requirements for human oversight, while the NIST framework's Map and Measure functions could draw on the ranked attribute structure to identify context-specific risks. However, a systematic mapping between the inventory and these governance instruments is beyond the scope of this exploratory study and represents an important direction for future work.

5.8. Implications for Design Practice and Education

The findings suggest several implications, offered as exploratory propositions given the study's scope. First, the 26-attribute inventory provides a structured vocabulary that design teams can use to evaluate the human-centeredness of AI systems under development. Second, the qualitative finding that trust is conceptualized as a relational outcome suggests that design efforts aimed at building trust should focus on transparency and oversight mechanisms rather than treating trust as a direct design target. Third, the consistent low prioritization of affective attributes, combined with qualitative recognition of their importance, points to a need for interdisciplinary collaboration between AI engineers and researchers in affective computing, psychology, and cultural studies. Fourth, the academic–student divergence suggests that HCAI education may benefit from explicit instruction on participatory design and well-being considerations. Fifth, the human-/user-prefix distinction identified in this study may serve as a useful organizing principle for design teams seeking to ensure their systems address both universal human qualities and interaction-specific user requirements, aligning with the longstanding distinction between human-centered design and user-centered design [20].

Drawing on the attributes with the strongest convergence across quantitative ratings, qualitative themes, and cross-group agreement, we propose nine preliminary design considerations for HCAI development, organized into three layers (Figure 2). The first layer—core foundations—comprises the five attributes that were rated most highly, supported by qualitative themes, and agreed upon across all three stakeholder groups: human trust (to be achieved through transparency and explainability rather than treated as a direct design target), human values (ethical principles embedded from the earliest design stages), human needs (fundamental requirements prioritized before system capabilities), human benefits (outcomes evaluated by positive impact on individuals and communities), and usability (effectiveness, efficiency, and satisfaction in context of use). The second layer—participatory processes—comprises three attributes that showed the largest academic–student divergence and reflect how human-centered design should be conducted: user involvement at every stage, engagement of diverse stakeholders in design and governance, and preservation of human control over consequential decisions. The third layer—non-negotiable integrity—comprises human dignity, which the expert panel reinstated as conceptually essential despite its lower composite score, recognizing that respect for inherent human worth is a boundary condition that must not be compromised regardless of other design priorities. We emphasize that these considerations are exploratory propositions derived from a single study with a predominantly Southeast Asian sample; they require empirical testing in actual AI development contexts before they can be considered validated design guidelines.




Figure 2. Preliminary design considerations for human-centered AI based on attributes with strongest cross-method convergence.

5.9. Limitations and Future Work

Several limitations should be acknowledged. The sample is geographically concentrated in Southeast Asia (77.2%), with Myanmar alone representing 46.9% of respondents. All findings should be interpreted within this context, and cross-cultural replication is essential before the attribute rankings can be considered generalizable. The academic subgroup ($n = 31$) is smaller than ideal, reducing statistical power for detecting nuanced pairwise differences; the academic–student divergence should therefore be treated as hypothesis-generating pending replication with larger, more balanced samples.

This study represents the exploratory first phase of a planned multi-study program. We did not conduct exploratory or confirmatory factor analysis, measurement invariance testing, or formal content validity indexing. The inter-attribute correlation structure suggests interpretable clusters, but these require formal factor-analytic testing with a larger, more diverse sample. The potential for construct overlap among semantically related attributes (e.g., human needs vs. user needs; usability vs. user-friendliness) was addressed through expert review, and the inter-attribute correlation analysis confirmed that flagged pairs exhibited only moderate correlations. Nonetheless, formal discriminant validity testing remains an important future direction.

The qualitative component is constrained by the survey format; open-ended questions elicited relatively brief responses that support thematic identification but limit interpretive depth. Future work should complement these findings with in-depth interviews and participatory design workshops. The connection between the attribute inventory and governance frameworks [13,14] was noted as suggestive rather than systematically developed; a dedicated mapping study represents an important next step. Finally, domain-specific adaptations for high-stakes contexts such as healthcare AI, educational AI, or autonomous systems represent a priority for future applied research, as the relative importance of specific attributes may shift considerably depending on the deployment context and the populations affected. In this study, the $N = 145$ sample, while sufficient for exploratory psy-

chometric assessment and the detection of large between-group effects, was not designed to support population-representative inference. Future confirmatory work should employ probability-based or stratified sampling with predetermined power calculations targeting specific effect sizes of interest.

More concretely, we identify three priority directions for the next phase of this research program. First, exploratory and confirmatory factor analyses on a larger sample (target $N \geq 300$) should test whether the four interpretable clusters identified here—experience/usability, emotion/empathy, benefit/goal, and participatory engagement—represent stable latent dimensions, with measurement invariance testing across stakeholder groups and cultural contexts. Second, replication studies should be conducted in at least three contrasting regional contexts (e.g., Western Europe, North America, Sub-Saharan Africa) to evaluate whether the priority structure observed in this Asia-Pacific sample generalizes or shifts in culturally meaningful ways. Third, domain-specific adaptations of the inventory should be developed and validated for high-stakes deployment contexts including healthcare AI, educational AI, and autonomous decision systems, where the relative weight of specific attributes is likely to vary substantially.

6. Conclusions

This study offers an empirically grounded characterization of human-centeredness in AI through a 26-attribute inventory assessed across practitioners, academics, and students in a predominantly Asia-Pacific sample, enriched by five qualitative themes. The findings converge on a clear priority structure: trust, values, needs, and benefits occupy the center of stakeholders' understanding of what makes AI human-centered, while affective and cognitive dimensions—though recognized as important—are perceived as the most difficult frontiers for current AI capabilities. The primary stakeholder divergence lies between academics and students, centering not on whether ethical foundations matter but on whether the participatory processes needed to realize them are understood and valued.

The qualitative findings reveal that stakeholders conceptualize human-centeredness through relational, experiential, and structural lenses. Trust is understood not as a design input but as a relational achievement contingent on transparency and meaningful oversight. The complexity of human emotion and culture is recognized as a genuine epistemological barrier, not merely a technical one. And the systemic constraints of corporate incentives, regulatory gaps, and resource limitations underscore that human-centered AI is as much a governance and organizational challenge as it is a design one.

As an exploratory study, this work maps the conceptual terrain of human-centeredness and generates hypotheses for confirmatory investigation. The 26-attribute inventory, while requiring psychometric refinement through factor analysis and measurement invariance testing, offers immediate practical value as a shared vocabulary for cross-stakeholder dialogue, a pedagogical resource for HCAI education, and a starting point for developing assessable design criteria. We hope it provides a useful organizing structure for the multi-disciplinary community of researchers, designers, educators, and policymakers working to ensure that AI systems are designed for, with, and accountable to the people they are intended to serve.

Funding: This research received no external funding.

Institutional Review Board Statement: Formal ethics committee review was not required for this study in accordance with Thailand's Ethical Guidelines for Human Research in Social Sciences and Humanities (2024), issued by the Institute for the Development of Human Research Protections (IHRP), Health Systems Research Institute (HSRI), Thailand (available online: <https://ihrp.hsri.or.th/blog/1752042777>, accessed on 15 April 2026), as the research involved no clin-

ical interventions, no collection of sensitive personal data, and no procedures that posed risk to participants. The study was conducted using voluntary participation, anonymized data, and standard usability and interaction evaluation methods, in line with these national guidelines and the principles of the WMA Declaration of Helsinki (World Medical Association, Declaration of Helsinki—Ethical Principles for Medical Research Involving Human Subjects, 2013; available online: <https://www.wma.net/policies-post/wma-declaration-of-helsinki-ethical-principles-for-medical-research-involving-human-subjects/>, accessed on 15 April 2026).

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: The data supporting the findings of this study are not publicly available due to participant privacy and confidentiality restrictions. Requests for access may be directed to the corresponding author.

Acknowledgments: The author gratefully acknowledges everyone who contributed to data collection, data cleaning, and preliminary analysis during the course of this study. During the preparation of this manuscript, the author used generative AI tools (Claude Opus 4.7) for the purposes of English language editing and grammar refinement. The author has reviewed and edited all output and takes full responsibility for the content of this publication.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Régis, C.; Denis, J.-L.; Axente, M.L.; Kishimoto, A. (Eds.) *Human-Centered AI: A Multidisciplinary Perspective for Policy-Makers, Auditors, and Users*, 1st ed.; Chapman and Hall/CRC: Boca Raton, FL, USA, 2024. [CrossRef]
2. Shneiderman, B. Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy. *Int. J. Hum.-Comput. Interact.* **2020**, *36*, 495–504. [CrossRef]
3. Shneiderman, B. *Human-Centered AI*; Oxford University Press: Oxford, UK, 2022.
4. Xu, W. Toward Human-Centered AI: A Perspective from Human-Computer Interaction. *Interactions* **2019**, *26*, 42–46. [CrossRef]
5. Amershi, S.; Weld, D.; Vorvoreanu, M.; Fournery, A.; Nushi, B.; Collisson, P.; Suh, J.; Iqbal, S.; Bennett, P.N.; Inkpen, K.; et al. Guidelines for Human-AI Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, Glasgow, Scotland, UK, 4–9 May 2019; Paper 3; ACM: New York, NY, USA, 2019; pp. 1–13. [CrossRef]
6. Xu, W.; Dainoff, M.J.; Ge, L.; Gao, Z. Transitioning to Human Interaction with AI Systems: New Challenges and Opportunities for HCI Professionals to Enable Human-Centered AI. *Int. J. Hum.-Comput. Interact.* **2023**, *39*, 494–518. [CrossRef]
7. Ozmen Garibay, O.; Winslow, B.; Andolina, S.; Antona, M.; Bodenschatz, A.; Coursaris, C.; Falco, G.; Fiore, S.M.; Garibay, I.; Grieman, K.; et al. Six Human-Centered Artificial Intelligence Grand Challenges. *Int. J. Hum.-Comput. Interact.* **2023**, *39*, 391–437. [CrossRef]
8. Schmager, S.; Pappas, I.O.; Vassilakopoulou, P. Understanding Human-Centred AI: A Review of Its Defining Elements and a Research Agenda. *Behav. Inf. Technol.* **2025**, *44*, 3771–3810. [CrossRef]
9. Dignum, V. *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way*; Springer: Cham, Switzerland, 2019.
10. Floridi, L.; Cows, J. A Unified Framework of Five Principles for AI in Society. *Harv. Data Sci. Rev.* **2019**, *1*. [CrossRef]
11. Ehsan, U.; Riedl, M.O. Human-Centered Explainable AI: Towards a Reflective Sociotechnical Approach. In *HCI International 2020—Late Breaking Papers: Multimodality and Intelligence*; Stephanidis, C., Kurosu, M., Degen, H., Reinerman-Jones, L., Eds.; Springer: Cham, Switzerland, 2020; pp. 449–466.
12. Nielsen, J. Heuristic Evaluation. In *Usability Inspection Methods*; Nielsen, J., Mack, R.L., Eds.; John Wiley & Sons: New York, NY, USA, 1994; pp. 25–62.
13. NIST. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*; NIST AI 100-1; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2023. [CrossRef]
14. European Parliament and Council of the European Union. Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act). *Off. J. Eur. Union* **2024**, *L*, 2024/1689.
15. ISO/IEC 42001:2023; Information Technology—Artificial Intelligence—Management System. International Organization for Standardization: Geneva, Switzerland, 2023.
16. Licklider, J.C.R. Man-Computer Symbiosis. *IRE Trans. Hum. Factors Electron.* **1960**, *HFE-1*, 4–11. [CrossRef]
17. Engelbart, D.C. *Augmenting Human Intellect: A Conceptual Framework*; SRI Report AFOSR-3223; Stanford Research Institute: Menlo Park, CA, USA, 1962.

18. Card, S.K.; Moran, T.P.; Newell, A. *The Psychology of Human-Computer Interaction*; Lawrence Erlbaum Associates: Hillsdale, NJ, USA, 1983.
19. Shneiderman, B. Direct Manipulation: A Step Beyond Programming Languages. *Computer* **1983**, *16*, 57–69. [[CrossRef](#)]
20. ISO 9241-210:2019; Ergonomics of Human-System Interaction—Part 210: Human-Centred Design for Interactive Systems. International Organization for Standardization: Geneva, Switzerland, 2019.
21. Norman, D.A. *The Design of Everyday Things*, Revised and Expanded ed.; Basic Books: New York, NY, USA, 2013.
22. LeCun, Y.; Bengio, Y.; Hinton, G. Deep Learning. *Nature* **2015**, *521*, 436–444. [[CrossRef](#)] [[PubMed](#)]
23. Jobin, A.; Ienca, M.; Vayena, E. The Global Landscape of AI Ethics Guidelines. *Nat. Mach. Intell.* **2019**, *1*, 389–399. [[CrossRef](#)]
24. Friedman, B.; Hendry, D.G. *Value Sensitive Design: Shaping Technology with Moral Imagination*; MIT Press: Cambridge, MA, USA, 2019.
25. Holstein, K.; Wortman Vaughan, J.; Daumé, H., III; Dudík, M.; Wallach, H. Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need? In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, Glasgow, Scotland, UK, 4–9 May 2019; ACM: New York, NY, USA, 2019; pp. 1–16. [[CrossRef](#)]
26. Johnson, R.B.; Onwuegbuzie, A.J. Mixed Methods Research: A Research Paradigm Whose Time Has Come. *Educ. Res.* **2004**, *33*, 14–26. [[CrossRef](#)]
27. Tukey, J.W. *Exploratory Data Analysis*; Addison-Wesley: Reading, MA, USA, 1977.
28. Cohen, J. *Statistical Power Analysis for the Behavioral Sciences*, 2nd ed.; Lawrence Erlbaum Associates: Hillsdale, NJ, USA, 1988.
29. Braun, V.; Clarke, V. Using Thematic Analysis in Psychology. *Qual. Res. Psychol.* **2006**, *3*, 77–101. [[CrossRef](#)]
30. Palinkas, L.A.; Horwitz, S.M.; Green, C.A.; Wisdom, J.P.; Duan, N.; Hoagwood, K. Purposeful Sampling for Qualitative Data Collection and Analysis in Mixed Method Implementation Research. *Adm. Policy Ment. Health* **2015**, *42*, 533–544. [[CrossRef](#)] [[PubMed](#)]
31. Nunnally, J.C.; Bernstein, I.H. *Psychometric Theory*, 3rd ed.; McGraw-Hill: New York, NY, USA, 1994.
32. Pant, A.; Hoda, R.; Tantithamthavorn, C.; Turhan, B. Ethics in AI through the practitioner’s view: A grounded theory literature review. *Empir. Softw. Eng.* **2024**, *29*, 67. [[CrossRef](#)]
33. Corrêa, N.K.; Santos, J.W.; Galvão, C.; Pasetti, M.; Schiavon, D.; Naqvi, F.; Hossain, R.; De Oliveira, N. Crossing the principle–practice gap in AI ethics with ethical problem-solving. *AI Ethics* **2025**, *5*, 1271–1288. [[CrossRef](#)]
34. Allen, K.; Yawson, A.K.; Haggenmüller, S.; Kather, J.N.; Brinker, T.J. Human-centered AI as a framework guiding the development of image-based diagnostic tools in oncology: A systematic review. *ESMO Real World Data Digit. Oncol.* **2024**, *6*, 100077. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.