

UNVEILING SCHOLARLY NARRATIVES ON HUMAN TECHNOLOGY: A STRUCTURAL TOPIC MODELING APPROACH

Wojciech Jarecki
University of Szczecin
Poland
ORCID 0000-0002-4463-7607

Magdalena Olczyk
Gdansk University of Technology
Poland
ORCID 0000-0002-4880-5890

Przemysław Olczyk
University of Gdańsk
Poland
ORCID 0009-0001-0899-1262

Monika Imreh-Toth
Széchenyi István University
Hungary
ORCID 0000-0002-0094-5827

Abstract: *This study systematically maps the thematic evolution of human–technology research from 2020 to 2025 using Structural Topic Modelling (STM) applied to over 2,000 abstracts from the Web of Science Core Collection. The analysis identifies six dominant themes: Industry, AI and Sustainability; Education and Human-Centred Design; Public Health, Community and Equity; Robotics, HRI and Ergonomics; Philosophy and Ethics of Technology; and Clinical mHealth and Usability. The results reveal a structural realignment from pandemic-driven experimentation to institutionalised, interdisciplinary research embedded in industrial, clinical, and community systems. Thematic inequality declined while diversity stabilised, indicating a mature and balanced research ecosystem. Methodologically, the study introduces a reproducible STM-based workflow integrating Gini and Shannon indices. Empirically, it provides a data-driven map of cross-disciplinary convergence. Conceptually, it demonstrates that human–technology inquiry increasingly operationalises ethics and sustainability through design, governance, and applied practice.*

Keywords: *human technology; structural topic modeling (STM); bibliometric analysis; research mapping*

©2025 Jarecki, Olczyk, Olczyk, & Imreh-Toth, and the Centre of Sociological Research,
Poland

DOI: <https://doi.org/10.14254/1795-6889.2025.21-3.2>



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

INTRODUCTION

Technological change in the twenty-first century is rapidly reshaping healthcare, education, work, governance, and communication. The spread of artificial intelligence (AI), robotics, data-driven systems, and digital platforms not only alters how people use tools, but also how those tools configure human action, values, and institutions (Shneiderman, 2022; Tiutiunyk et al., 2021). Technologies are not neutral artefacts; they are components of socio-technical systems that embody values, distribute agency, and co-produce social norms and everyday practices (Winner, 1980; Bijker et al., 2012; Blok, 2022). A practical illustration is transport and logistics telematics, where real-time monitoring and analytics reshape operational decisions and human work practices (Szcześniak & Gorzelańczyk, 2024). Within these systems, the human role is not fixed but negotiated in practice: people appropriate, workaround, and domesticate systems through routines of use, repair, and adaptation, while algorithmic decision-making redistributes agency and accountability across humans and artefacts (Johannessen et al., 2023). These lived interactions help explain why similar technologies yield uneven benefits and harms across contexts (Schürkmann & Anders, 2025). Assessing whether ongoing change advances well-being, equity, and sustainability - or reproduces asymmetries and vulnerabilities - therefore requires inquiry that bridges social sciences, engineering, ethics, and design. As the literature expands at speed and across disciplines, even specialists struggle to retain analytical orientation, underscoring the need for systematic, cross-disciplinary mapping of human–technology studies.

Over the past two decades, a vast and heterogeneous body of human–technology research has emerged across diverse fields. In healthcare, the rise of mHealth, telemedicine, and data-driven remote care has generated thousands of studies exploring usability, accessibility, and patient engagement (Kallander et al., 2013; Zhao et al., 2024). In education, human-centred design and design thinking approaches have transformed pedagogical models and learning environments (Liedtka, 2018; Mishchuk et al., 2025). Meanwhile, robotics and ergonomics research has expanded toward collaborative human-robot systems emphasizing safety and usability (Goodrich & Schultz, 2007; Ajoudani et al., 2018), while philosophical and ethical debates have deepened around embodiment, technological mediation, privacy, and responsibility (Floridi et al., 2018; Tutar et al., 2025). Yet these communities often operate in parallel, which hampers cumulative knowledge and policy-relevant synthesis. In this context, bibliometric and scientometric approaches offer transparent tools for mapping intellectual structures. Unlike narrative reviews, they rely on systematic data collection and computational methods that reduce selection bias and enable scalability (Donthu et al., 2021). We leverage these advantages to provide a data-driven map of the human – technology literature over 2020–2025 – a period spanning the pandemic shock and the subsequent acceleration of digitalization – based on abstracts retrieved from the Web of Science Core Collection (WoSCC) using a carefully constructed query.

Our objective is to identify the field’s leading themes and trace their evolution. We apply Structural Topic Modeling (STM), which uses document-level metadata to estimate how topic prevalence varies over time and across contexts. To characterize structural balance, we complement STM with the Gini and Shannon indices to capture concentration and diversity. Using a corpus of more than 2,000 abstracts, we (i) map dominant thematic clusters for 2020

–2025, (ii) analyze their temporal dynamics before and after the pandemic, and (iii) assess whether the field is diversifying or fragmenting.

This study contributes in three ways. Methodologically, it advances a transparent, reproducible STM-based workflow for cross-disciplinary mapping. Empirically, it assembles a scalable corpus and analytic framework that synthesises dispersed subfields. Structurally, it introduces indicators and visualisations that diagnose thematic organisation and change, offering a transferable template for comparative analyses.

The remainder proceeds as follows: Section 2 describes data, corpus construction and methods; Section 3 presents results on dominant themes, temporal evolution and diversity metrics; Section 4 discusses implications for interdisciplinarity and theory; Section 5 concludes with future research directions and policy relevance.

METHODS

Data

A fundamental prerequisite for any bibliometric study is the construction of a well-defined and consistently curated bibliographic database. Establishing a reliable dataset is the first critical step that determines the transparency, comparability, and reproducibility of subsequent analyses. Our database is created based on records in the Web of Science Core Collection (WoSCC) is widely regarded as the gold standard for bibliometric and scientometric research due to its rigorous curation of peer-reviewed journals and consistent metadata structure (Pranckutė, 2021). Compared to other databases such as Scopus or Google Scholar, WoSCC offers higher transparency, stability of indexing, and compatibility with advanced analytical tools for citation and text analysis. The Web of Science Core Collection includes the Science Citation Index Expanded (SCIE), which covers the natural and technical sciences since around 1900 (webofscience.help.clarivate.com); the Social Sciences Citation Index (SSCI) for the social sciences (guides.library.queensu.ca); and the Arts & Humanities Citation Index (AHCI), which includes the arts and humanities from approximately 1975 (lib.montana.edu). In addition, it contains conference proceedings, book citation indexes, and the Emerging Sources Citation Index (ESCI) (guides.lib.umich.edu).

The query in the Web of Science Core Collection (WoSCC) was conducted on 20 May 2025. The search was limited to document types including articles, proceedings papers, early access papers, book chapters, and books published from the years 2020 to 2025. Additionally, only abstracts written in English were included in created database, as they provide the most content-rich descriptions of research focus and ensure the relevance of texts for subsequent topic modeling. The query was constructed around the core expression “human technology”, including its direct orthographic variants such as “human-technology” and compound forms like “human technology interaction” or “human-technology interaction”. To capture semantically equivalent terms commonly employed in the field, the query further incorporated “human-centered/centred technology”, “human-centered/centred design”, and “human-centered/centred computing”, thereby accounting for both American and British orthography as well as syntactic variations. This design enhanced recall while maintaining precision, since all selected expressions directly relate to the conceptual nucleus of human-centred

technological development and interaction. By systematically including these lexical variants, the search minimized the risk of omitting relevant publications due to orthographic or stylistic differences, while avoiding overly broad or tangential terms that could introduce noise into the corpus. The final query was: AB (abstracts)=(`"human technology" OR "human-technology" OR "human technology interaction" OR "human-technology interaction" OR "human-centered technology" OR "human-centred technology" OR "human-centered design" OR "human-centred design" OR "human-centered computing" OR "human-centred computing"`).

The final dataset comprised 2,022 abstracts retrieved from the Web of Science Core Collection, spanning 1,223 distinct journals and authored by approximately 1,993 unique researchers. This wide coverage underscores the highly interdisciplinary nature of human–technology scholarship.

Methodology

We employ Structural Topic Modeling (STM) – a probabilistic text mining technique designed to uncover latent thematic structures in large textual corpora while simultaneously accounting for document-level metadata (Roberts et al., 2014). STM assumes that each document d is represented as a mixture of K latent topics, and each topic corresponds to a probability distribution over the vocabulary. In contrast to classical Latent Dirichlet Allocation (LDA), STM enables topic prevalence and content to vary systematically as a function of observed covariates, such as author attributes, publication year, or institutional affiliation.

Prior to model estimation, the corpus was pre-processed through tokenization, stopword removal, lowercasing, and stemming. Rare words and documents with extremely low word counts were excluded to improve convergence and interpretability. All analyses were conducted in R using the `stm` package (Roberts et al., 2019). The estimation procedure proceeds in five main steps (Roberts et al., 2016):

1) Initialization: Each document d is initially assigned random topic proportions θ_d representing a mixture over K topics. These serve as starting values for the subsequent optimization routine.

2) Expectation – Maximization Updates: STM employs a variational Expectation–Maximization (EM) algorithm to iteratively update latent assignments. During each iteration, the model compares observed word frequencies to the current estimates of topic–word distributions. Words that are strongly associated with a topic increase the posterior probability of that topic in the document’s composition.

3) Incorporation of Metadata: Topic prevalence is linked to document-level covariates through a logistic–normal regression structure. The latent variable n_d follows a multivariate normal distribution with a mean determined by the covariates

$$X_d: n_d \sim N(X_d, \Gamma, \Sigma) \quad (1)$$

where X_d is the vector of metadata, Γ is the matrix of regression coefficients, and Σ denotes the covariance matrix.

4) Transformation into Topic Proportions: The latent vector n_d is then mapped onto the simplex of valid topic proportions using the softmax transformation:

$$\theta_{dk} = \frac{\exp(n_{dk})}{\sum_{j=1}^K \exp(n_{dj})} \quad (2)$$

ensuring that all topic proportions are positive and sum to one.

5) Convergence: The iterative updates continue until convergence, yielding posterior topic proportions θ_d that jointly capture the observed word co-occurrences and the systematic variation explained by document-level metadata.

A crucial step in the Structural Topic Model approach is determining the optimal number of topics (K). This decision is not purely algorithmic but requires balancing statistical diagnostics with substantive interpretability. Following established guidelines by Roberts et al. (2014; 2016), three complementary criteria were jointly evaluated. First, held-out likelihood measures predictive accuracy on data excluded from estimation. It increases rapidly for small K and then flattens, indicating diminishing returns from additional model complexity. Second, semantic coherence captures how consistently high-probability words co-occur across documents. It is typically highest for smaller K and decreases as the number of topics grows, reflecting the fragmentation of semantically meaningful clusters. Third, exclusivity assesses topic distinctiveness through the uniqueness of top words. It tends to increase with K , as rarer terms form separate clusters, although excessive segmentation may produce overly narrow or analytically trivial topics.

Taken together, these diagnostics indicate that selecting $K = 6$ provides a balanced compromise (Figure 1)

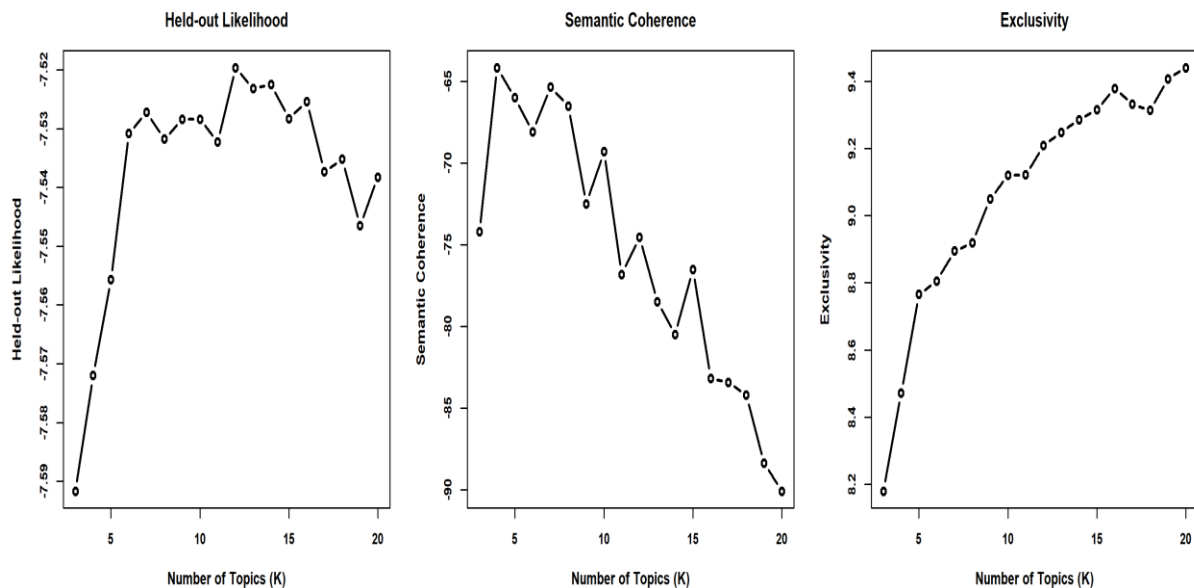


Figure 1. Diagnostic Metrics for Selecting the Optimal Number of Topics
[Authors' computation based on STM diagnostic framework by Roberts et al. (2016)]

At $K = 6$, the held-out likelihood is near its maximum, coherence remains sufficiently high for reliable human interpretation, and exclusivity is strong enough to yield clear thematic separation without undue granularity. Substantively, the six-topic configuration maps cleanly onto the major strands of the human–technology literature, avoiding both under-aggregation and over-fragmentation. Hence, $K = 6$ is both statistically defensible and substantively

vocabulary used within a topic, revealing the conceptual breadth of human–technology research. Collectively, they show a balanced distribution between applied and theoretical domains - ranging from industrial transformation and design education to public health, robotics, ethics, and clinical usability - thereby illustrating the multidimensional nature of the field. We provide table 1 with a precise, textual summary of the most representative terms, formally defining each topic according to the STM algorithm.

Table 1. Top 15 FREX terms describing the thematic composition of the 2020–2025 human–technology corpus
[Source: Authors’ computation based on STM output]

Topic	Top 15 FREX Words	Label
Topic 1	transformation; sustainable; business; sustainability; industry; manufacturing; smart; companies; consumers; framework; IoT; maturity; adoption; HCAI; innovation	Industry, AI & Sustainability
Topic 2	students; teachers; engineering; design; learning; creativity; course; problems; product; prototyping; thinking; user; testing; education; feedback	Education & Human-Centred Design
Topic 3	communities; vaccine; dementia; vaccination; equity; covid; pandemic; programs; public; rural; aging; disparities; equitable; co-design; care	Public Health, Community & Equity
Topic 4	flight; accuracy; motion; neural; vehicles; task; experiment; automation; simulation; ergonomics; human-robot; comfort; sensors; prototype; HRI	Robotics, HRI & Ergonomics
Topic 5	relations; embodied; agency; mediated; moral; philosophical; non-human; shops; animals; argue; privacy; alcohol; media; humans; animal	Philosophy & Ethics of Technology
Topic 6	patients; clinical; HIV; nurses; cancer; pain; mHealth; clinic; medication; CDS; chronic; usability; intervention; pediatric; accessibility	Clinical mHealth & Usability

Topic 1 – Industry, AI & Sustainability. This topic explores how digital transformation and AI reshape human work within industrial systems, redefining the roles of operators, engineers, and managers, and their interactions with HMIs, IoT infrastructures, and HCAI tools, as well as the implications for safety, trust, and decision quality. Research connects operational metrics (OEE, defect rates, energy intensity, carbon footprint) with human-centred indicators (usability, cognitive load, transparency, acceptance). Data from connected devices and user feedback support co-design processes and continuous improvement. Governance discussions address bias and privacy in HCAI, interoperability across platforms, and the alignment of ESG objectives with workplace constraints and employee well-being.

Topic 2, titled "Education and Human-Centred Design," explores how learners, teachers, and external users co-create technologies through human-centred practices. Instruction connects theory to application through problem framing, prototyping, and user testing, with students acting as co-designers who translate user needs into requirements and refine solutions based on feedback from studies and field trials. Programmes assess human outcomes such as usability, workload, accessibility, trust, and explainability alongside technical performance, using rubrics, evaluations, and structured test protocols. Courses include ethical safeguards for generative tools related to attribution, bias, privacy, and consent, and promote inclusion through universal design and accessibility. Stakeholders such as patients, teachers, operators, and community partners provide context, validate use cases, and help ensure that solutions function in real settings. Capstone studios and living lab projects expose learners to realistic

constraints, including safety, cost, and maintenance, while reflection and peer critique help develop responsible innovation skills. Overall, the topic highlights human agency in design and demonstrates how education connects engineering decisions with user value, equity, and social responsibility.

Topic 3 – "Public Health, Community and Equity" highlights the central role of human agency in shaping technologies for health, care, and community well-being. Research in this cluster examines how patients, caregivers, and local actors use, adapt, and reinterpret technologies in contexts such as vaccination programmes, pandemic response, dementia care, and rural health delivery. Rather than treating technology as an external intervention, studies frame it as a social interface where relationships of trust, reciprocity, and responsibility are negotiated. Human-centred and participatory design approaches enable communities to articulate priorities, co-create prototypes, and evaluate usability and fairness in real-life settings. Through these practices, people act not merely as users but as co-designers and stewards of digital and organisational systems that mediate care and distribute resources (Yakymova et al., 2022). Equity frameworks and anti-racist design principles guide how technologies are implemented, ensuring they respect cultural values, accessibility needs, and diverse forms of knowledge. Overall, this topic emphasises that the effectiveness and legitimacy of public health innovation depend on recognising people as active participants who co-produce technologies, governance processes, and ethical standards for equitable care.

Topic 4 – "Robotics, HRI & Ergonomics" . This topic presents robots and automated systems as partners that co-shape human work, rather than as standalone tools. Studies investigate how operators, pilots, and drivers allocate attention and authority with machines in flight simulators, motion-tracking tasks, and automated-vehicle testbeds, and how these interactions influence safety and performance. Human agency is central: operators monitor, negotiate handovers, and recover from edge cases; designers translate sensory and neural signals into interfaces that support transparency, predictability, and trust; participants' feedback calibrates control policies and comfort thresholds. Outcome measures connect human experience—workload, situational awareness, trust calibration, embodiment, and comfort—with technical metrics such as accuracy and precision. The literature regards automation as both an opportunity and a risk: it can reduce physical strain and improve consistency, yet it requires ergonomic principles and governance that preserve attention, appropriate reliance, and meaningful human control. Overall, the research demonstrates that future industrial and transport systems depend on HRI designs that prioritise human performance and well-being while responsibly leveraging robotic capabilities.

Topic 5 – Philosophy and Ethics of Technology. This topic explores how people and technologies co-constitute each other in everyday and institutional contexts. Contributions analyse embodied experience, agency, and mediation, drawing on postphenomenology and socio-technical network perspectives to explain how artefacts shape perception, choice, and responsibility. The literature addresses privacy, surveillance, manipulation, alienation, and the limits of intervention, showing how these issues are negotiated in ordinary settings such as media platforms and retail environments, as well as in public institutions. Designers, users, and regulators co-define acceptable practice through consent, transparency, explainability, accountability, and audit mechanisms, embedding ethics in design and governance procedures rather than treating them as stand-alone reflection. Debates also consider non-human actors and heritage contexts, where technology acts as a mediator and active participant in networks

of relations. Overall, the focus is on human agency and value conflict, and on maintaining human dignity and fairness as technologies are specified, deployed, and governed.

Topic 6, labeled "Clinical mHealth and Usability," This topic examines how mobile health tools and clinical decision support (CDS) systems are designed, implemented, and governed within real care pathways. Research focuses on patient groups such as people living with HIV, oncology patients, and those with chronic conditions, highlighting the roles of clinicians, nurses, and caregivers in selecting, configuring, and using digital tools. Studies assess human outcomes alongside clinical ones, including usability, learnability, workload, trust, explainability, adherence, and safety. Methods include co-design with patients and staff, think-aloud usability tests, workflow observations on wards and in clinics, and field trials tracking adherence and symptom reporting. A recurring theme is integration with existing infrastructures. Work addresses EHR interoperability, alert fatigue, data quality, privacy, and consent management, as well as accessibility for users with low literacy, children, and minority populations. Designs that support shared decision-making, plain-language prompts, multilingual interfaces, and inclusive visual cues improve engagement and equity. Governance is treated as part of design: teams calibrate CDS thresholds, audit algorithmic bias, and set procedures for human override. Overall, the literature shows how human-centred design translates into clinical practice by aligning digital tools with routines, responsibilities, and values in care delivery, thereby improving effectiveness while protecting dignity, safety, and fairness.

Building on the thematic composition shown in Figure 2, the next analysis examines how the prevalence of the six topics changes over time. By tracking their annual distribution from 2020 to 2025, we can identify shifts in scholarly attention and moments when specific domains gain or lose prominence. Figure 3 visualizes these temporal dynamics and shows how the field of human–technology interaction responds to social, technological, and policy developments.

The temporal evolution of topics between 2020 and 2025, as shown in Figure 3, reveals a substantial restructuring of research priorities in the human–technology domain. Topics 2 ("Education and Human-Centred Design") and 4 ("Robotics, HRI and Ergonomics"), which are most prominent in 2020–2021, record the steepest subsequent declines. Their early salience aligns with the global pivot to remote education and the rapid, experimental deployment of automation during the initial COVID-19 response. As emergency adaptation gave way to normalization, and with typical publication lags moving clinical and implementation studies into later years, research attention shifted from short-term instructional and ergonomic challenges toward systemic questions of integration, governance, and social impact. The decline of these topics therefore reflects the waning of crisis-driven innovation and the absorption of human-centred methods into broader design and policy frameworks.

By contrast, Topic 6 ("Clinical mHealth and Usability") shows the strongest growth and becomes the most prevalent theme by 2025. This expansion parallels the institutionalization of digital health infrastructure, the consolidation of telemedicine practices, and the global policy emphasis on healthcare resilience. Initiatives such as the European Health Data Space and national digital-health strategies have legitimized mHealth research as a permanent component of health-system modernization. Similarly, Topics 1 ("Industry, AI and Sustainability") and 3 ("Public Health, Community and Equity"), which were marginal in 2020, advance into the top tier by 2025. Their rise mirrors wider socio-political trends: the European Green Deal, the AI

Act, and post-pandemic recovery programs have redirected research funding and corporate strategies toward sustainable industry, ethical AI, and social well-being.




Figure 3. Trends in topic prevalence across the human–technology corpus (2020–2025).
[Source: Authors’ computation based on STM output]

Together, these shifts mark a transition from reactive, pandemic-specific research to more strategic and institutionalized agendas that embed digital technologies within industrial, medical, and community systems. The human–technology field thus moves beyond early experimentation to address long-term priorities such as resilience, inclusiveness, and governance, reflecting how global crises, sustainability imperatives, and regulatory initiatives reshape both the content and direction of scientific inquiry.

To test whether these observed trends were statistically significant rather than random fluctuations, we conducted inferential comparisons of topic prevalence before (2020–2021) and after (2022–2025) the onset of COVID-19 using independent-sample t-tests with Benjamini–Hochberg correction for multiple testing (Figure 4). The results confirm both substantive and statistical realignment of the field. Topics T2 and T4, dominant before the pandemic (mean prevalence 0.30 and 0.22), declined significantly to 0.21 and 0.18 (BH-adjusted $p < 0.05$). Conversely, Topics T1 and T6 increased significantly (from 0.08 to 0.12 and from 0.15 to 0.20; BH-adjusted $p < 0.05$), while Topic T3 showed a moderate but robust increase (0.11 \rightarrow 0.14; $p < 0.05$). Topic T5 remained statistically unchanged (0.15 \rightarrow 0.13; n.s.), indicating its stable yet secondary role. Overall, the inferential evidence confirms a structural realignment from short-term pandemic responses to diversified, policy-oriented, and socially embedded lines of inquiry.




Figure 4. The comparison of topic prevalence before (2020–2021) and after (2022–2025).
[Source: Authors’ analysis]

In the next step we analyse the correlation matrix of topic proportions, which illustrates how frequently the six topics co-occur within the same abstracts (see Figure 5).




Figure 5. Correlation matrix of topic proportions.
[Source: Authors’ computation based on STM output]

As shown in Figure 5, the overall correlations are weak and predominantly negative, indicating a relatively distinct thematic separation between topics. The absence of strong positive correlations suggests that most papers are clearly associated with a single dominant thematic area rather than combining multiple ones.

The highest negative correlations occur between Topic 6 (Clinical mHealth and Usability) and Topics 1, 2, 4, and 5 ($r \approx -0.3$), implying that clinical and health-related studies tend to be conceptually and empirically isolated from research on industry, education, robotics, and ethics. Topic 3 (Public Health, Community and Equity), however, shows a weak positive link with Topic 6 ($r = 0.04$), reflecting their shared focus on health and social well-being. The moderate negative correlations among the remaining topics (-0.1 to -0.3) indicate specialization without fragmentation: while subfields are distinct, they coexist within a coherent intellectual structure. Overall, the correlation matrix confirms that the six-topic configuration captures meaningful thematic boundaries, supporting the internal validity of the STM model and reinforcing the interpretation that human–technology research has evolved into a diversified yet integrated domain.

To complement the descriptive analysis of topic prevalence, two indices were employed to capture the structural characteristics of the thematic distribution: the Gini index and Shannon entropy. Both indices were computed using normalized topic prevalence scores obtained from the STM posterior estimates. The Gini index (Figure 6) quantifies the degree of inequality in topic prevalence. It is derived by comparing the differences between all pairs of topic shares and normalizing them by the mean share across topics. In intuitive terms, the index measures how unevenly research attention is distributed: values close to 1 indicate strong concentration, where a few topics dominate the field, while values near 0 reflect a more balanced thematic structure.




Figure 6. Evolution of Thematic Inequality (Gini Index, 2020–2025). [Source: Authors' analysis]

The decline in the Gini index indicates a gradual equalization of topic shares between 2020 and 2025, meaning that research attention became more evenly distributed across themes. In relation to Figure 3, this trend reflects a shift from the early dominance of Topics 2 (Education and Human Centered Design) and 4 (Robotics, HRI and Ergonomics) toward the growing importance of Topics 1 (Industry, AI and Sustainability), 3 (Public Health, Community and Equity), and 6 (Clinical mHealth and Usability). The decrease in inequality corresponds with

the diversification and institutionalization of the human technology field, as studies expanded from short term responses during the pandemic to broader systemic and policy oriented agendas. The pattern is also consistent with the correlation matrix in Figure 4, where weak and mostly negative correlations indicate specialization without fragmentation. Overall, the declining Gini index shows that the field evolved toward a more balanced and mature configuration in which diverse research areas coexist within an integrated structure.

The Shannon entropy, in turn, measures the overall diversity of topic distribution, with higher values indicating a more balanced and pluralistic research landscape and lower values reflecting concentration around a few dominant themes. It is calculated by multiplying each topic's share by the natural logarithm of that share, summing the results for all topics, and taking the negative of that sum. While the Gini coefficient showed a steady decline from 0.3 in 2020 to 0.1 in 2025, indicating a gradual reduction in thematic inequality, the entropy measure reveals a more complex dynamic. (Figure 7) .




Figure 7. Changes in Research Diversity Over Time: Shannon Entropy Index, 2020–2025.

[Source: Authors' analysis]

Entropy increased from 0.512 in 2020 to a peak of about 0.530 in 2023, marking a period of growing diversification and balance across research streams. After this peak, a mild decline to 0.517 in 2025 suggests a partial reconsolidation of scholarly attention, likely around new dominant themes that emerged during the earlier expansion. Taken together, the Gini and Shannon indicators depict a research field that initially diversified and then began to stabilize into a moderately pluralistic yet more coherent configuration.

DISCUSSION

This study systematically mapped the thematic evolution of human–technology research from 2020 to 2025 using a Structural Topic Modelling (STM) approach. The model identified six dominant topics: Industry, AI and Sustainability; Education and Human-Centred Design; Public Health, Community and Equity; Robotics, HRI and Ergonomics; Philosophy and Ethics of Technology; and Clinical mHealth and Usability. The findings reveal a significant structural realignment in the field after 2020. The early pandemic period was characterised by rapid,

practice-oriented responses in education and ergonomics, while subsequent years saw a marked shift towards applied, solution-oriented, and interdisciplinary research streams embedded in industrial, clinical and community systems. Thematic inequality, measured by the Gini index, declined steadily from 0.30 in 2020 to 0.10 in 2025, while Shannon entropy increased to a peak in 2023 and then mildly declined ($0.512 \rightarrow \sim 0.530 \rightarrow 0.517$), indicating a pluralistic yet increasingly balanced field. Together, these indicators suggest a maturing research ecosystem that has moved beyond disciplinary silos towards a more integrated configuration of applied subfields.

The shift from crisis-driven, short-term responses to practical and institutionalised application reflects both societal and institutional pressures shaping knowledge production. Global crises such as COVID-19 accelerated the adoption of digital tools in health care, education, and industrial organisation (Shneiderman, 2022; Zhao et al., 2024). Consequently, scholars prioritised empirical evaluation and design optimisation over abstract theorisation. The increasing prominence of Industry, AI, and Sustainability mirrors the widespread diffusion of industrial AI and the IoT as drivers of green and digital transitions (Phatthanachaisuksiri et al., 2025). This strand connects human-centred artificial intelligence (HCAI) with corporate sustainability frameworks, echoing the call for socially aligned innovation emphasised by Shneiderman (2022). This institutional embedding is also reflected in the maturation of sustainability reporting research, which maps how disclosure and accountability practices become standardized governance mechanisms (Agama & Zubairu, 2022). In parallel, Clinical mHealth and Usability has emerged as the fastest-growing domain, reflecting intensified investment in digital health infrastructure and patient-centred care (Deniz-García et al., 2023). Consistent with these shifts, Education and Human-Centred Design and Robotics, HRI, and Ergonomics declined from their early-pandemic prominence, while Public Health, Community, and Equity registered a moderate increase. Overall, these transformations indicate that the human–technology nexus has entered an era of institutional embedding, where the “human” is defined through design protocols, usability standards, and governance requirements rather than philosophical debate.

Conversely, Philosophy & Ethics of Technology did not change significantly, suggesting that value-sensitive inquiry is increasingly diffused into adjacent applied domains rather than marginalised. Ethical considerations are now integrated into AI governance, user-experience design, and data management (Batool et al., 2025). This aligns with Floridi (2023), who argues that ethics is increasingly enacted through design choices and algorithmic audits rather than stand-alone philosophical argumentation. By contrast, Education & Human-Centred Design declined from its early-pandemic prominence, indicating that while pedagogical reform remains influential (Hu & Chan, 2025), its relative share has diminished as attention shifted toward institutionalised industrial, clinical, and community agendas. Collectively, these shifts illustrate a movement from debating the “why” of human–technology relations to optimising the “how” of human-centred implementation.

The trends observed in our analysis are in line with earlier mapping efforts. García-Fernández et al. (2024) found that research on digital human resources shifted from conceptual discussion to managerial implementation, while Zhao et al. (2024) observed a similar shift in healthcare HCI. The current analysis extends these insights by quantifying not only topic prevalence but also structural balance, showing that the expansion of applied domains does not lead to fragmentation. Moderate negative correlations among topics indicate specialisation

without isolation: health-oriented studies rarely overlap with robotics or educational design, yet all coexist within a coherent network of shared methods and ethical vocabularies. At the same time, the persistence of philosophical inquiry at a stable level, together with the diffusion of ethical reasoning into adjacent applied areas, resonates with what Blok (2022) calls the “re-embedding of creativity” in human–technology relations. Similar findings appear in bibliometric analyses of AI ethics, where clusters on fairness, transparency, and accountability increasingly merge with applied computer-science venues (Qiu et al., 2025). Hence, this pattern reflects not a retreat of philosophy but its integration into pragmatic domains that internalise ethical reasoning.

Our findings contribute to broader theoretical debates on how human agency and technological mediation co-evolve. The diversity of topics supports the socio-technical systems perspective emphasizing that technologies are not external tools but nodes in networks of human and non-human actors (Bisconti et al., 2024). The rise of practice-based research in design and robotics corroborates postphenomenological accounts that view interaction as mutual shaping rather than one-way influence (Gray et al. 2025). Empirical evidence from STM shows that humans increasingly appear not as abstract users but as situated agents - operators, patients, students, and consumers - who negotiate functionality and meaning through participation and feedback loops. This evolution resonates with Johannessen et al. (2023), who describe “multi-site domestication” of technologies across institutions, illustrating how agency is redistributed through everyday appropriation and repair.

Moreover, the convergence of industrial AI and sustainability indicates a paradigmatic shift towards “responsible automation”, where efficiency gains are assessed against social and environmental outcomes (Shabur et al., 2025). By quantifying how these topics emerge and interact, this study empirically substantiates theoretical claims about the co-production of technological and normative orders. It also demonstrates that resilience in socio-technical systems depends not only on technological robustness but also on institutional capacity to integrate ethical and participatory feedback.

CONCLUSIONS

This study systematically mapped the thematic evolution of human–technology research from 2020 to 2025 using a Structural Topic Modelling approach. By analysing over 2,000 abstracts from the Web of Science Core Collection, six dominant research clusters were identified: Industry, AI and Sustainability; Education and Human-Centred Design; Public Health, Community and Equity; Robotics, HRI and Ergonomics; Philosophy and Ethics of Technology; and Clinical mHealth and Usability. Collectively, these topics delineate a field that has shifted from early, crisis-driven and experimental implementations towards empirically grounded, solution-oriented, and institutionally embedded inquiry integrating design, ethics, and governance. Ethical reasoning remains present but is increasingly operationalised within applied domains.

The findings indicate a structural realignment after the COVID-19 pandemic. Thematic inequality (Gini index) declined steadily from 0.30 in 2020 to 0.10 in 2025, while Shannon entropy peaked in 2023 before mildly declining, suggesting a pluralistic yet increasingly coherent configuration of subfields. Early practice-oriented agendas in education and

ergonomics were replaced by applied domains combining human-centred AI, digital health, and sustainability frameworks, marking a measurable maturation of the field.

Quantitatively, convergence toward lower inequality alongside sustained diversity shows that research attention is redistributing rather than collapsing into a few dominant strands, preserving pluralism without fragmentation. Substantively, the co-evolution of design protocols, usability standards, and governance practices demonstrates that the “human” dimension is increasingly specified through auditable procedures rather than abstract ethical claims. This consolidation of a shared methodological vocabulary - user studies, living labs, algorithmic audits - enables cumulative learning and methodological comparability across contexts. The study thus provides both a reproducible diagnostic for monitoring structural balance over time and a framework for identifying where integrative research is already consolidating versus where disciplinary boundaries persist.

Methodologically, the STM-based workflow combining topic modelling with diversity metrics (Gini and Shannon entropy) offers a transferable tool for assessing consolidation and fragmentation in interdisciplinary domains. Theoretically, the results support sociotechnical and postphenomenological perspectives by showing that technologies operate as co-evolving nodes in networks of human and nonhuman actors, with humans functioning as situated agents who shape systems through feedback and adaptation.

Despite its strengths, the study is limited by reliance on English-language abstracts from WoSCC and a focus on co-occurrence rather than causal mechanisms. Future work should integrate citation and co-authorship networks, embedding models, and multilingual corpora to validate cross-domain comparability and include non-Western epistemologies.

Ultimately, this research contributes a reproducible, data-driven framework for mapping structural maturity and inclusiveness in interdisciplinary knowledge fields, offering a transferable model for evaluating how ethical and sustainable design principles become institutionally embedded within the evolving human–technology nexus.

IMPLICATIONS FOR RESEARCH, APPLICATION, OR POLICY

The field’s partial reconsolidation after 2023 points to actionable next steps. For research, the STM-plus-diversity workflow should be reused to monitor thematic balance, compare subfields, and evaluate how funding reshapes trajectories. Priorities include extending beyond English, integrating citation and co-authorship networks to map epistemic communities, and testing embedding models to sharpen topic boundaries and detect drift. For application, the growing institutionalization of ethics in design, mHealth, and industrial AI argues for operationalizing ethics via usability benchmarks, governance checklists, audit trails, human-override procedures, and documentation of data stewardship, accessibility, and equity. For policy, alignment with Horizon Europe and the Digital Decade suggests favouring cross-disciplinary consortia that link engineering, social science, and ethics, require pilots or living labs, and evaluate outcomes with human-centred metrics safety, trust, inclusion, and environmental impact. Overall, these implications shift emphasis from abstract principles toward verifiable practice, measurable standards, and funding criteria that reward responsible, socially grounded innovation.

ACKNOWLEDGMENTS

The project in Poland is financed by the Ministry of Science and Higher Education in Poland under the program "Regional Excellence Initiative".



REFERENCES

- Agama, E. J., & Zubairu, U. M. (2022). Sustainability reporting: A systematic review. *Economics, Management and Sustainability*, 7(2), 32–46. <https://doi.org/10.14254/jems.2022.7-2.3>
- Ajouadani, A., Zanchettin, A. M., Ivaldi, S., Albu-Schäffer, A., Kosuge, K., & Khatib, O. (2018). Progress and prospects of the human–robot collaboration. *Autonomous Robots*, 42(5), 957–975. <https://doi.org/10.1007/s10514-017-9677-2>
- Batool, A., Zowghi, D. & Bano, M. (2025). AI governance: a systematic literature review. *AI Ethics* 5, 3265–3279. <https://doi.org/10.1007/s43681-024-00653-w>
- Bijker, W. E., Hughes, T. P., & Pinch, T. (Eds.). (2012). *The social construction of technological systems: New directions in the sociology and history of technology* (Anniversary ed.). MIT Press.
- Bisconti, P., McIntyre, A., & Russo, F. (2024). Synthetic socio-technical systems: Poiêsis as meaning making. *Philosophy & Technology*, 37, 94. <https://doi.org/10.1007/s13347-024-00778-0>
- Blok, V. (2022). The role of human creativity in human–technology relations. *Philosophy & Technology*, 35(3), 59. <https://doi.org/10.1007/s13347-022-00559-7>
- Deniz-García, A., Fabelo, H., Rodríguez-Almeida, A. J., Zamora-Zamorano, G., Castro-Fernández, M., Alberiche Ruano, M. D. P., Solvoll, T., Granja, C., Schopf, T. R., Callicó, G. M., Soguero-Ruiz, C., Wägner, A. M., & WARIFA Consortium. (2023). Quality, usability, and effectiveness of mHealth apps and the role of artificial intelligence: Current scenario and challenges. *Journal of Medical Internet Research*, 25, e44030. <https://doi.org/10.2196/44030>
- Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines. *Journal of Business Research*, 133, 285–296. <https://doi.org/10.1016/j.jbusres.2021.04.070>
- Floridi, L. (2023). *Copyright page*. In *The ethics of artificial intelligence: Principles, challenges, and opportunities*. Oxford University Press. <https://doi.org/10.1093/oso/9780198883098.002.0003>
- Floridi, L., Cowsls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People-An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. *Minds and Machines*, 28(4), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- García-Fernández, L., Aguado-Cortés, Á., Ortega-Mier, M., & Bores-Calle, N. (2024). Mapping the main research themes in digital human resources: A bibliometric analysis. *Humanities and Social Sciences Communications*, 11, 207. <https://doi.org/10.1057/s41599-024-03795-8>
- Goodrich, M. A., & Schultz, A. C. (2007). Human–robot interaction: A survey. *Foundations and Trends in Human-Computer Interaction*, 1(3), 203–275. <https://doi.org/10.1561/1100000005>
- Gray, K., Yam, K. C., Zhen’An, A. E., Dillion, D., & Waytz, A. (2025). The psychology of robots and artificial intelligence. In D. T. Gilbert, S. T. Fiske, E. J. Finkel, & W. B. Mendes (Eds.), *The handbook of social psychology* (6th ed.). Situational Press. <https://doi.org/10.70400/IKVM1257>

- Hu, W., & Chan, C. K. Y. (2025). From user needs to AI solutions: A human-centered design approach for AI-powered virtual teamwork competency training. *International Journal of Educational Technology in Higher Education*, 22, 52. <https://doi.org/10.1186/s41239-025-00551-z>
- Johannessen, L. E. F., Nordtug, M., & Haldar, M. (2023). Multi-site domestication: taming technologies across multiple institutional settings. *Information, Communication & Society*, 27(11), 2077–2093. <https://doi.org/10.1080/1369118X.2023.2255644>
- Källander, K., Tibenderana, J. K., Akpogheneta, O. J., Strachan, D. L., Hill, Z., ten Asbroek, A. H. A., Conteh, L., Kirkwood, B. R., & Meek, S. R. (2013). Mobile health (mHealth) approaches and lessons for increased performance and retention of community health workers in low- and middle-income countries: A review. *Journal of Medical Internet Research*, 15(1), e17. <https://doi.org/10.2196/jmir.2130>
- Liedtka, J. (2018). Why design thinking works. *Harvard Business Review*, 96(5), 72–79.
- Mishchuk, H., Samoliuk, N., Krol, V., & Račka, I. (2025). Digital competences of university graduates: Implications for career success. *Human Technology*, 21(2), 274–292. <https://doi.org/10.14254/1795-6889.2025.21-2.2>
- Phatthanachaisuksiri, L., Cai, M., Karam, H., et al. (2025). Exploring the role of AI in achieving the SDGs: A systematic review of academic studies on AI in the industry sector [Version 1; peer review: 2 approved with reservations]. *Open Research Europe*, 5, 160. <https://doi.org/10.12688/openreseurope.20407.1>
- Pranckutė, R. (2021). Web of Science (WoS) and Scopus: The Titans of Bibliographic Information in Today's Academic World. *Publications*, 9(1), 12. <https://doi.org/10.3390/publications9010012>
- Qiu, J., Cheng, L., & Huang, J. (2025). Charting the landscape of artificial intelligence ethics: A bibliometric analysis. *International Journal of Digital Law & Governance*, 3(1), 135–167. <https://doi.org/10.1515/ijdlg-2025-0007>
- Roberts, M. E., Stewart, B. M., & Tingley, D. (2016). A model of text for experimentation in the social sciences. *Journal of the American Statistical Association*, 111(515), 988–1003. <https://doi.org/10.1080/01621459.2016.1141684>
- Roberts, M. E., Stewart, B. M., & Tingley, D. (2019). stm: An R package for structural topic models. *Journal of Statistical Software*, 91(2), 1–40. <https://doi.org/10.18637/jss.v091.i02>
- Roberts, M. E., Stewart, B. M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S. K., Albertson, B., & Rand, D. G. (2014). Structural topic models for open-ended survey responses. *American Journal of Political Science*, 58(4), 1064–1082. <https://doi.org/10.1111/ajps.12103>
- Schürkmann, C., & Anders, L. (2025). Postphenomenology unchained: Rethinking human–technology–world relations as enroulement. *Human Studies*, 48, 367–390. <https://doi.org/10.1007/s10746-024-09746-1>
- Shabur, M., Shahriar, A., & Ara, M. A. (2025). From automation to collaboration: Exploring the impact of Industry 5.0 on sustainable manufacturing. *Discover Sustainability*, 6, 341, 1–16. <https://doi.org/10.1007/s43621-025-01201-0>
- Shneiderman, B. (2022). *Human-centered AI*. Oxford University Press. <https://doi.org/10.1093/oso/9780192845290.001.0001>
- Szcześniak, J., & Gorzelańczyk, P. (2024). Analysis of the use of logistics-telematics systems for cost reduction in a transport company. *Journal of Sustainable Development of Transport and Logistics*, 9(1), 6–18. <https://doi.org/10.14254/jsdtl.2024.9-1.1>
- Tiutiunyk, I., Reznikova, N., & Sotnyk, I. (2021). The impact of digital transformation on macroeconomic stability: Evidence from EU countries. *Journal of International Studies*, 14(3), 220–234. <https://doi.org/10.14254/2071-8330.2021/14-3/14>
- Tutar, H., Bilan, S., & Şentürk, Ümit. (2025). Is there consistency in ethical sensitivity in artificial intelligence? A review of language models. *Human Technology*, 21(2), 317–338. <https://doi.org/10.14254/1795-6889.2025.21-2.4>
- Winner, L. (1980). Do Artifacts Have Politics? *Daedalus*, 109(1), 121–136. <http://www.jstor.org/stable/20024652>

Yakymova, L., Novotná, A., Kuz, V., & Tamándl, L. (2022). Measuring industry digital transformation with a composite indicator: A case study of the utility industry. *Journal of International Studies*, 15(1), 169–186. <https://doi.org/10.14254/2071-8330.2022/15-1/11>

Zhao, X., Zhang, S., Nan, D., Han, J., & Kim, J. H. (2024). Human–computer interaction in healthcare: A bibliometric analysis with CiteSpace. *Healthcare*, 12(23), 2467. <https://doi.org/10.3390/healthcare12232467>

Authors' Note

All correspondence should be addressed to
Magdalena Olczyk
Gdańsk University of Technology, Faculty of Management and Economics
Narutowicza 11/12, 80-233 Gdańsk, Poland
Email: Magdalena.Olczyk@pg.edu.pl

Human Technology
ISSN 1795-6889
<https://ht.csr-pub.eu>