Dean&Francis

ISSN 2959-6157

Research Trends in Human-AI Collaboration: A Five-Year Density Analysis Based on Domain Categorization

Yanchen Zhang

Pioneer Academy, Wayne, United States, 07470
yanchen.zhang@pioneeracademy.org

Abstract:

Human-AI teamwork is rapidly becoming one of the most energetic and promising frontiers in artificial intelligence. It drives a wide array of applications, from intelligent virtual assistants to collaborative creative platforms and interactive dialogue systems. This research offers a data-driven perspective on how this field has progressed over the past five years. By analyzing comprehensive reports such as the *AI Research Progress Report (2020–2024)* and using a specially developed keyword index for subfield classification, we identified which domains are expanding the fastest and where most research efforts are concentrated. This research integrates publication data, growth metrics, and comparative analysis across various subfields, employing visual tools like charts and linear regression models—to determine emerging hotspots and developing trends. Notably, dialogue generation systems and cooperative AI agents consistently produce the highest volume of research, while innovative AI-powered creative tools are experiencing swift growth . These findings not only provide a snapshot of the current environment in Human-AI collaboration but also suggest focal points for future breakthroughs.

Keywords: Human-AI Collaboration; Subfield Classification; Research Trend; Dialogue Systems; Quantitative Modeling

1. Introduction

As artificial intelligence (AI) continues to move increasingly forward, perhaps one of the greatest changes has been the development of human-AI collaboration. This concept goes beyond simple automation. It considers AI technologies interactive collaborators that can augment human decision-making, communication, learning, creativity, and more. From talking robots and AI assistants to tools assisting with creation and systems helping, this collaboration is reshaping technological development and end-user expectations in numerous ways. Accordingly, knowledge about the development of research in this area

can aid the forecasting of future trends and guide responsible development.

Although this subfield has gained greater prominence, it remains ill-defined and eclectic, overlapping with other subfields of human-centered AI as well as other AI subfields. Even popular applications such as ChatGPT, GitHub Copilot, and DALL·E have gained recognition, their related research covers a wide range of topics. These topics range from dialogue systems, trust, and explainability to co-creative tools, educational platforms, and decision-support systems. However, very few works have consistently tracked these subfields of time. It is therefore difficult to discern where academic interest has concentrated, where subfields have converged and stabilized, and where new research opportunities lie.

This paper fills the void with a data-driven investigation of Human-AI collaboration work from 2020-2024. We take large datasets from the AI Research Progress Report [1] and extract a scheme of classifications based on selectively selected sets of keyword groupings. These categories include Dialogue Systems, Human-AI Teaming, AI Creativity Tools, Emotion and Trust, Interactive Interfaces, Educational Agents, Governance, Autonomous Support Systems, and other related areas. Using these classifications, we compile yearly publication counts and discern growth patterns using growth measures and regression models.

By visualizing and interpreting these trends, this paper promotes knowledge of the future of Human-AI co-working research and subfields' development toward each other. It is designed to inform and guide researchers, designers, and developers to strategically position their work within this rapidly expanding body of research. The rest of the paper is organized such that Section 2 explains our methodology and scheme of keyword-based classification; Section 3 displays the trend analysis by time; Section 4 covers interpretation of data across domains; and Section 5 concludes with the implication of future research and practice. These macro-level patterns align with external benchmarks: the Stanford AI Index 2024 documents a sustained rise in AI research output and the diffusion of generative-model capabilities across application domains. [4]

2. Methodology

This paper adopts a Bibliometrics and Content Analysis methodology to study research trends in Human-AI collaboration from 2020 to 2024. We constructed the dataset primarily by using yearly publication counts from the AI Research Progress Report and supplementing it with records searched on the OpenAlex open-access scientific index.[1, 2] For the identification of domain-specific trends, we developed a custom subfield taxonomy of top-level areas of interest: Dialogue Systems, Human-AI Teaming, AI Creativity Tools, Trust & Explainability, Interactive Interfaces, Governance & Ethics, Education & Tutoring, Autonomous Systems, and Emotion & Empathy.

Each subfield was defined by a carefully crafted list of relevant keywords (e.g., "chatbot", "co-creative tools", "explainable AI"), which were used to identify and collect publication records from the source dataset.

Data was formatted into tables where rows represented single years (2020–2024) and columns represented paper counts corresponding to subfields. Neither filtering at a higher level nor data cleaning was attempted because the raw data was clean enough to accommodate higher-level quantitative computation. The study scope excludes citation analysis and full-text semantic interpretation; only publication volume is used as a proxy for research activity.

To identify temporal trends, we calculated two main indices for every subfield:

·Annual publication volume, utilized for monitoring absolute research attention.

·Five-year growth rate, calculated as the relative percent change from 2020 to 2024.

These numbers were depicted in bar charts and line graphs using Microsoft Excel and Python. For example, line graphs reveal temporal growth at subfield levels, while bar graphs enable side-by-side comparison of research volumes in the endpoint (2024) by comparison. Information on rates of growth was used within domain ranking based on emergent relevance.

By using this methodology, we attempt to provide a systematic comparative account of the evolution of the diverse subfields of Human–AI collaboration of the past five years, and what subfields gained accelerating academic interest.

3. Results

This section provides a detailed view of Human–AI Collaboration subfield publication trends between 2020 and 2024. For ease of understanding and to show the findings' structure, we divide the findings into four subsections:

(1) overall patterns over the five-year span,

(2) subfields with the highest publication volume in 2024,

(3) most rapidly growing and shrinking subfields by compound growth rate, and

(4) year-by-year domain-specific trajectories.

These analyses are supported by three figures visualizing publication counts and growth patterns.

ISSN 2959-6157

3.1 Five-Year Publication Trends Overall

An analysis of publication trends from 2020 to 2024 identifies distinct trends in the development of Human-AI collaboration subfields.

Figure 1 illustrates that publication volume within all nine subfields maintains a steady upward trend within the five-year period and is demonstrating sustained research interest within the domain. While some subfields maintained steady growth rates, others exhibited fluctuations or stagnation, reflecting the varied momentum of different subfields.

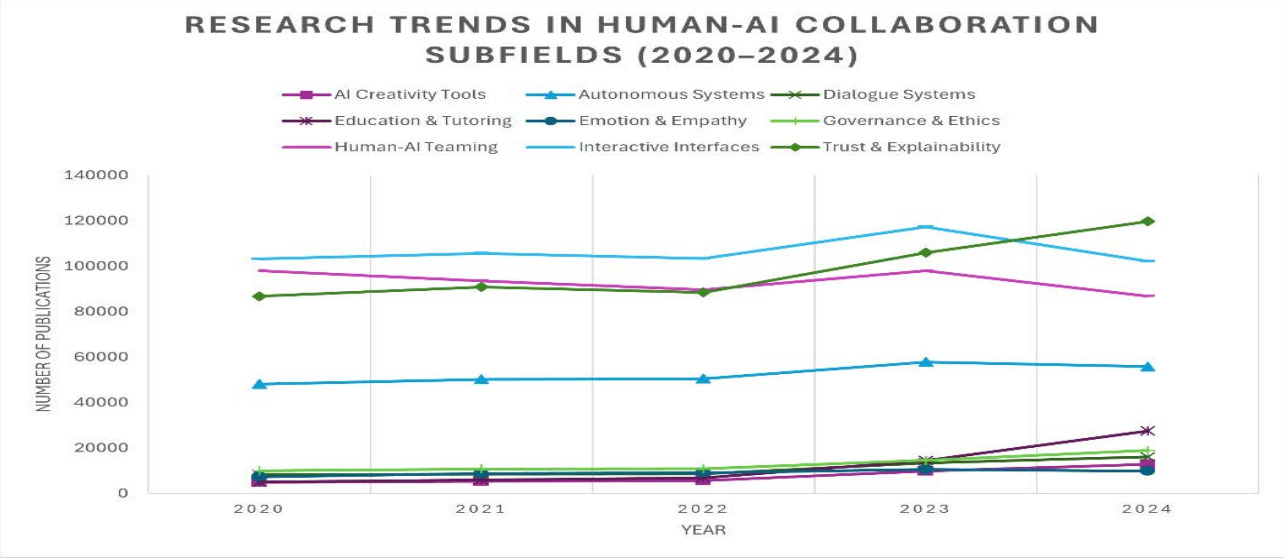

Figure 1. Yearly publication trends of nine subfields of Human–AI Collaboration from 2020 to 2024.

This line chart displays the chronological evolution of publication counts and illuminates fluctuations of research momentum and constancy within subfields.

3.2 Top Subfields by Volume in 2024

For 2024, Dialogue Systems, Human–AI Teaming, and AI Creativity Tools had the highest absolute publication counts. Figure 2 is a horizontal bar chart of subfield publication counts for 2024 and makes it obvious that this is a hierarchy. Dialogue Systems led the list with the greatest publication count and did so consistently throughout the observation period. Next was Human–AI Teaming, demonstrating increasing interest in co-operative agent construction and team enhancement. AI Creativity Tools, encompassing systems of co-creative poetry aid, music composition aid, and design aid, was also toward the top and demonstrates increasing interest in human-centered creativity enhancement.

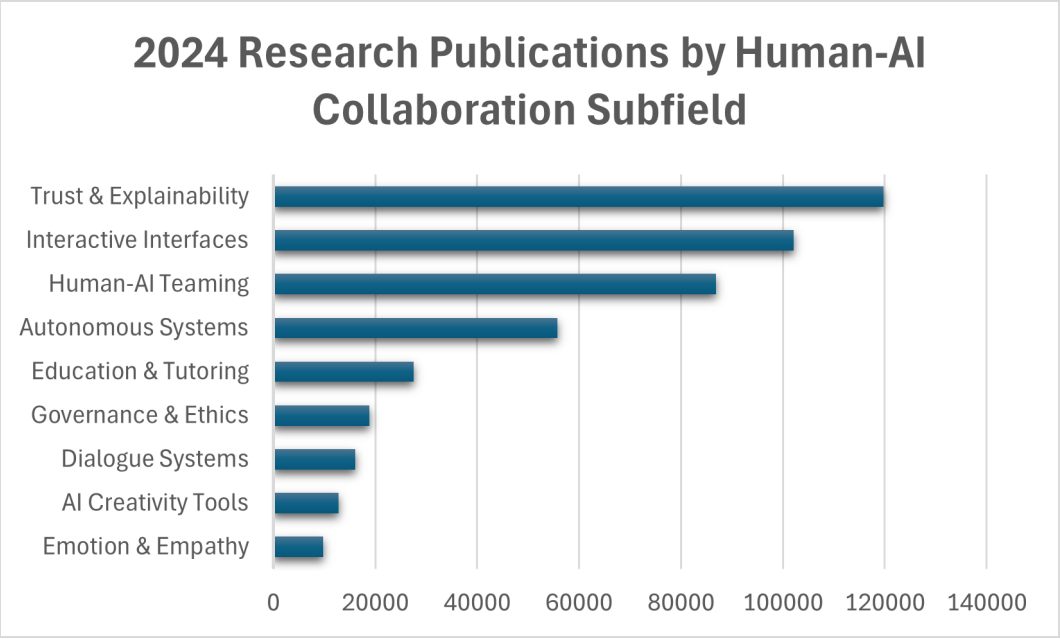

Figure 2. Publications by subfield in 2024 (ranked in descending order).

This horizontal bar chart shows the absolute publication counts of every subfield in the year 2024 and displays what are the areas which received maximum research interest.

3.3 Fastest Growing and Declining Subfields

Other than raw volume, growth rates provide insights into the strength of every subfield. Figure 3 is the percentage change in volume of papers between 2020 and 2024. Highest growth was recorded by AI Creativity Tools, Education & Tutoring, and Emotion & Empathy. It is notable that AI Creativity Tools recorded an upward compound growth of 150%, making it a subfield with both high publication volume and rapid growth. Similarly, Education & Tutoring and Emotion & Empathy, which started at low bases, recorded enormous relative growth and documented growing scholastic interest in affective computing and pedagogy with AI assistance.

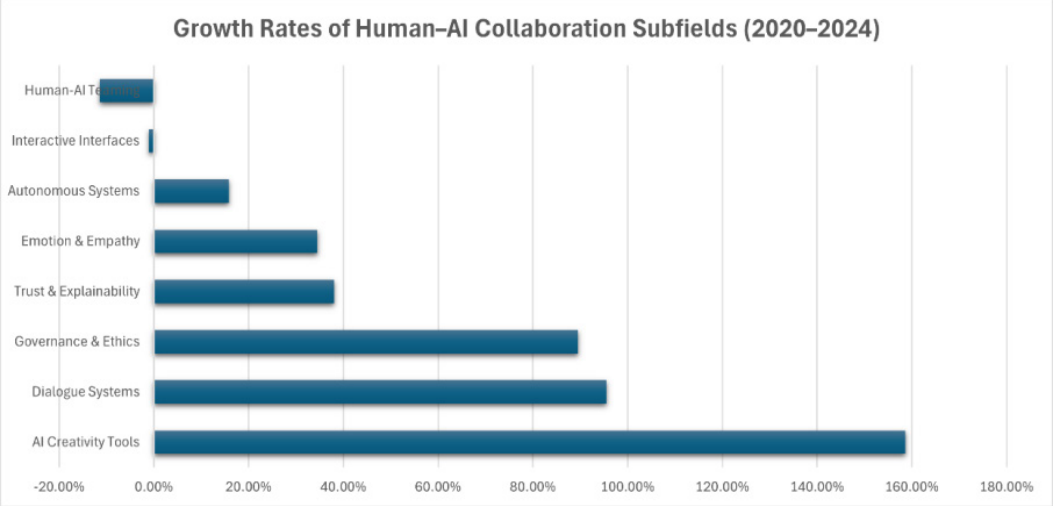

Figure 3. Subfield compound growth rates from 2020 to 2024.

This is a comparative subfield growth chart of each of the subfields within the five-year period and shows areas of very strong growth such as AI Creativity Tools and Education & Tutoring and those of negative or zero growth. Conversely, Figure 3 highlights areas of shrinkage or decline compared with other subfields' growth. Subfield Trust & Explainability had low negative growth over the five years and may have seen a shift of attention toward downstream use or adoption within other subfields such as governance. The Interactive Interfaces subfield remained

ISSN 2959-6157

stable, which may indicate maturity or a saturation of research interest.

3.4 Year-wise Evolution Patterns

Year-wise evolution trends further reinforce these findings. As detailed in Figure 1, subfields such as Dialogue Systems and Human-AI Teaming exhibited steady linear growth. while others such as AI Creativity Tools experienced a remarkable exponential growth after 2022. This temporal divergence helps highlight the varying levels of maturity and research pace within the Human-AI collaboration ecosystem.

These visualizations together give a composite picture of the research landscape. Figure 1 places longitudinal growth within context, Figure 2 establishes a baseline of currently attained activity levels, and Figure 3 highlights points of strategic growth or decline. They collectively highlight both the strength of mature subfields and the emergence of new research frontiers in Human–AI collaboration.

4. Discussion

This paper presents results that offer valuable insights regarding transformations that have occurred within the Human–AI collaboration space within the recent five-year span. From subfield-level publication evolution examined in our work, we have a clearer understanding of priorities that have changed, new prospects that are emerging, and possible points of saturation within the research space.

Arguably the most powerful trend is the growth explosion of the subfield of Education & Tutoring itself, where publication volume surged 449% between 2020 and 2024. This explosive growth can likely be traced back to the global shift toward online education triggered by the COVID-19 pandemic, which has since driven the adoption of AI-powered tutoring platforms and adaptive learning systems. Also, with the widespread availability of large language models has come the promise of scalable smart education tools that continue to galvanize interest.

Likewise, the sudden spike in AI Creativity Tools is a sign of increased interest in generative AI, co-creation systems, and human-centered design when it comes to content creation. This is parallel with larger societal developments—the emergence of AI-aided music, art, and writing—to which new questions about the collaborative creativity of human and AI system combinations have been raised. A proximate technical driver of this surge is the foundation-model paradigm, whose broad transferability and instruction-following abilities enable rapid adaptation to co-creative and educational settings [5].

While both Trust & Explainability and Interactive Interfaces remain one of the best-explored subfields by raw numbers, their relative growth has been moderate. This may perhaps be a sign of maturity or saturation of these subfields where underlying work has mostly been put in place and follow-up work is incrementally insignificant. But steady volume is a testament to steady interest in interpretable and user-centered AI systems where high stakes or regulation are concerned.

Unexpectedly, Human–AI Teaming, once a flagship area of focus, reduced slightly in terms of publication output within the five-year period. This reduction might not necessarily indicate fading interest but rather take-up of new terms or incorporation within broader interdisciplinary areas such as cooperative agents, social robots, or human-centered reinforcement learning. It is just as concerning whether theories are moving at a pace that is commensurate with rapidly advancing capabilities.

Gradual but consistent development of subfields such as Dialogue Systems, Emotion & Empathy, and Governance & Ethics implies the rising interest in the social and psychological side of Human–AI interaction. The socially embedded nature of the role of AI systems renders these subfields increasingly prominent when it comes to mitigating bias, adjusting trust, and emotional alignment.

These findings together suggest a clear pattern: a shift from structural usability (e.g., teaming and interfaces) to relational intelligence (e.g., creativity, education, and empathy). This transformation likely indicates a broader philosophical shift in AI research itself: from performance optimization toward meaningful human-AI interaction.

Also, this work points out the merit of domain-level quantitative monitoring toward discerning research gaps and eventual future research directions. Although publication volume is far from being an ideal indicator of progress, it is a handy tool for observing interest, focus of funds, and attention of scholarship over the long term.

The present study provides a comprehensive quantitative analysis of research trends in the field of human-AI collaboration. However, this study still has limitations. This analysis uses publication output as the sole proxy for research activity. This approach does not include key factors such as citation impact, publication quality, or author collaboration networks, which would help more deeply reveal the influence and maturity of each subfield.

At the same time, the classification of papers into subfields relies on domain-specific keyword matching. Despite careful manual curation, this method may still introduce ambiguity and overlaps between categories.

The time frame of this study is limited to 2020-2024. Although this period is practically relevant, it may be insufficient to fully reflect longer-term historical evolution or deeper disciplinary shifts. Emerging subfields with rapid

recent growth may appear disproportionately prominent, whereas more mature and stable fields may be underestimated in growth comparisons.

In future research, these limitations can also be addressed by considering more dimensions, such as citation-based impact measures, author-level productivity metrics, and conference acceptance rates. In addition, applying natural language processing techniques to paper abstracts or full texts may yield a more fine-grained classification than relying solely on keyword mapping. For example, Westergaard et al. compared 15 million full texts with 16.5 million abstracts and found that text mining on full texts consistently outperformed using only abstracts across multiple information-extraction tasks.[3]

5. Conclusion

This paper presents a comprehensive data-driven overview of Human-AI collaboration research trends from 2020 to 2024 and identifies major developments across nine explicitly defined subfields. Adopting publication count as a coarse indicator of research activity and absolute and relative growth metrics, we recognize both the top-active areas at present and the new growth leaders.

Our findings reveal a dynamic and shifting research landscape. While mature areas of concern such as Trust & Explainability and Interactive Interfaces continue very popular research subfields, newer or previously overlooked subfields such as AI Creativity Tools and Education & Tutoring have experienced skyrocketing growth. These shifts reflect larger technological and societal trends, such as the rise of generative AI, growth of big language models, and rising interest in personalized and emotionally intelligent AI systems.

Note that the quantitative methodology used here facilitates comparative study across fields and over time and provides a reproducible format for trend observation and strategic planning of AI research. Although publication volume is examined alone and citation impact and semantics of contents are not considered within this study, it is still useful to obtain a rough high-level overview of where the Human-AI collaboration research is likely trending.

Future research can incrementally build on this work using finer-grained proxies such as citation numbers, grant spending, or local contributions and by studying subfield interactions or intersections. Ultimately, it is essential for informing innovation toward more human-aware, reliable, and effective AI technologies.

In future research, this work can be expanded by using more dimensions, such as citation impact, author productivity, and conference selectivity. At the methodological level, more fine-grained modeling based on full texts or abstracts can improve subfield classification and reduce overlap. Network-based methods (including co-authorship and citation-flow analyses) can further reveal interdisciplinary patterns. By linking academic trends with external signals such as patents, open-source activity, and regulatory developments, it may also be possible to provide a more comprehensive perspective for understanding the co-evolution of human–AI collaboration across research and real-world domains.

References

[1] National Science and Technology Council. (2024). 2020–2024 progress report: Advancing trustworthy artificial intelligence research and development. NITRD/OSTP. https://www.nitrd.gov/pubs/AI-Research-and-Development-Progress-Report-2020-2024.pdf

[2] Priem, J., Piwowar, H., & Orr, R. (2022). OpenAlex: A fully-open index of scholarly works, authors, venues, institutions, and concepts. arXiv. https://arxiv.org/abs/2205.01833

[3] Westergaard, D., Stærfeldt, H.-H., Tønsberg, C., Jensen, L. J., & Brunak, S. (2018). A comprehensive and quantitative comparison of text-mining in 15 million full-text articles versus their corresponding abstracts. PLOS Computational Biology, 14(2), e1005962. https://doi.org/10.1371/journal.pcbi.1005962

[4] AI Index Steering Committee. (2024). AI Index Report 2024. Stanford Institute for Human-Centered Artificial Intelligence (HAI). Available at Stanford HAI.

[5] Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., … Liang, P. (2021). On the Opportunities and Risks of Foundation Models. arXiv:2108.07258 / Center for Research on Foundation Models, Stanford. https://arxiv.org/abs/2108.07258