

©2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Responsible AI: Portraits with Intelligent Bibliometrics

Yi Zhang, *Senior Member, IEEE*, Mengjia Wu, Guangquan Zhang, Jie Lu, *Fellow, IEEE*

Abstract—Shifting the focus from principles to practical implementation, responsible artificial intelligence (AI) has garnered considerable attention across academia, industry, and society at large. Despite being in its nascent stages, this emerging field grapples with nebulous concepts and intricate knowledge frameworks. By analysing three prevailing concepts - explainable AI, trustworthy AI, and ethical AI, this study defined responsible AI and identified its core principles. Methodologically, this study successfully demonstrated the implementation of leveraging AI's capabilities into bibliometrics for enhanced knowledge discovery and the cross-validation of experimentally examined models with domain insights. Empirically, this study investigated 17,799 research articles contributed by the AI community since 2015. This involves recognizing key technological players and their relationships, unveiling the topical landscape and hierarchy of responsible AI, charting its evolution, and elucidating the interplay between the responsibility principles and primary AI techniques. An analysis of a core cohort comprising 380 articles from multiple disciplines captures the most recent advancements in responsible AI. As one of the pioneering bibliometric studies dedicated to exploring responsible AI, this study will provide comprehensive macro-level insights, enhancing the understanding of responsible AI while furnishing valuable knowledge support for AI regulation and governance initiatives.

Impact Statement — This study stands as one of the pioneering investigations into responsible artificial intelligence (AI), utilizing large-scale and comprehensive bibliometric studies. It develops an analytical framework that elaborates cutting-edge intelligent bibliometric models. The study analyses responsible AI from diverse aspects, including the recognition of key technological players, the identification of topical landscape, hierarchy, and evolutionary pathways, and the profile of a core cohort about responsible AI. Of particular significance is the empirical insights of the bibliometric portraits for responsible AI, which would benefit the broader AI community and relevant stakeholders to understand this emerging topic at a macro level and provide essential knowledge support for AI regulation and governance initiatives.

Index Terms—Responsible AI, Bibliometrics, Topic Analysis.

This paper was submitted on January 15, 2024.

This work was supported by the Commonwealth Scientific and Industrial Research Organization (CSIRO), Australia, in conjunction with the National Science Foundation (NSF) of the United States, under CSIRO-NSF #2303037; and the Australian Research Council (ARC) under Laureate Project FL190100149.

Y. Z., M. W., G. Z., and J. L. are with the Australian Artificial Intelligence Institute (AAII), University of Technology Sydney (UTS), Sydney, Australia.

I. INTRODUCTION

ARTIFICIAL intelligence (AI), with its incredible capabilities in learning patterns from large-scale data and training human-like intelligent models, is revolutionizing the human society from automation to digitalization with autonomous and self-awareness machines. The unprecedented success of generative AI, particular, ChatGPT, further triggers a global debate on AI and its profound societal impact, such as risks and threats to the broad society [1], and what to regulate and how [2].

Since the early beginning of the AI boom in the 2010s, in fact, global government agencies, giant IT companies, research institutions, professional societies, and non-profit organizations have started to produce a series of principles and guidelines for *AI ethics* [3], broadly touching the good of humanity, individuals, societies, and the environment¹. While AI ethics provides a visionary blueprint to describe what are the right things that we expect AI to do, the society urged an executive and tactical plan for how we can do.

This urgent need fosters the rise of *responsible AI* – a plan for both AI systems and their practitioners. This responsibility entails an AI system's fundamental functionality that enables AI to reason about and act according to human values. It further underscores the accountability and awareness of AI researchers and practitioners on the direct societal impact of AI [4]. Emphasizing a pragmatic and application-driven approach, the industry simplifies responsible AI as approaches to designing, developing, and deploying AI under a legal, ethical, and trustworthy framework [5], and has produced consultation reports and industry standards to support the implementation of responsible AI practices [6, 7].

The AI community's response to responsible AI appears to adopt a spontaneous bottom-up paradigm. There have been numerous algorithms and models, placing emphasis on their responsibility-related functionalities, e.g., the protection of privacy-sensitive data [8] and the promotion of fair decision support [9]. From a technical standpoint, impressive endeavours on conceptualizing AI's social responsibility [10]

(emails: Yi.Zhang@uts.edu.au, Mengjia.Wu@uts.edu.au, Guangquan.Zhang@uts.edu.au, Jie.Lu@uts.edu.au).

¹ The United Nations Educational, Scientific and Cultural Organization (UNESCO) recommended four core values of AI systems: Human rights and human dignity, living in peaceful, ensuring diversity and inclusiveness, and environment and ecosystem flourishing. See the website: <https://www.unesco.org/en/artificial-intelligence/recommendation-ethics>

and informing policy development for responsible AI [11] have created prominent impact. However, in this rapidly developing and emerging area, the AI community has not reached a consensus regarding the definitions, concepts, and technologies pertaining to responsible AI, despite some roadmaps outlining its application in specific sectors, e.g., agriculture [12] and healthcare [13]. Consequently, an investigation into the topical landscape, hierarchy, and evolution of responsible AI, along with an exploration of its key technological contributors, will bring fresh and expert insights to the development of a comprehensive strategy for AI governance and regulation.

This study aims to fulfill the gap by developing an intelligent bibliometrics-based analytical framework, designed to unveil comprehensive portraits of responsible AI, and discover empirical insights. Intelligent bibliometrics refers to the development and applications of AI-empowered computational models for analysing large-scale bibliometric data such as scientific articles and patent documents [14]. Notable models include scientific evolutionary pathways, enabling the trace of technological change through streaming data analytics [15], topical hierarchical trees, employed to profile the knowledge structure of a research domain using density-based hierarchical community detection [16], and knowledge trajectory prediction using graph analytics [17]. Intelligent bibliometrics has achieved significant success in broad science, technology, and innovation studies, e.g., tracking topical disruption, evolution, and resilience of coronavirus research [18] and unveiling digital transformation enabling capabilities [19].

This study created a comprehensive set of responsible AI-related portraits, with the objectives of recognizing key players and their relationships, identifying the topical landscape and hierarchy of responsible AI, charting its evolutionary pathways, and elucidating the interplay between responsibility principles and primary AI techniques. We combined research articles collected from Scopus and the Web of Science Core Collection and cross-referenced with the widely acknowledged academic knowledge graph OpenAlex. This yielded a dataset comprising 17,799 distinct articles on responsible AI.

Through this investigation, we drew the technical backbones of responsible AI, spanning from data architectures (data science) to analytical algorithms (machine learning), and from fundamental theories (mathematics) to system applications (computer networks). With a strong focus on privacy and security studies, we observed a dominance of contributions from China, the United States, and India, collectively accounting for over 50% of total publications. Our findings also highlighted concerted efforts in addressing the issues of explainability, transparency, trustworthiness, and reliability, particularly emanating from the machine learning community. Furthermore, our focus on a core cohort of 380 articles explicitly referencing the term “responsible AI” unveiled the cross-disciplinary nature of the field. From the perspective of a complex eco-system, the chaotic connections among terms from computer science and the humanities and social science disciplines may underscore the gradual establishment of responsible AI as a new knowledge area. Top-tier universities and prominent IT companies from the West are playing pivotal

roles in driving its development.

This study stands as one of the pioneering investigations into responsible AI, utilizing large-scale bibliometric analysis. It develops an analytical framework that elaborates cutting-edge intelligent bibliometric models. The study encompasses diverse aspects, including the recognition of key technological players, and the identification of topical landscape, hierarchy, and evolutionary pathways. Of particular significance is the empirical insights of the bibliometric portraits for responsible AI. These valuable insights would benefit the broader AI community and relevant stakeholders to understand this emerging topic at a macro level and provide essential knowledge support for AI regulation and governance initiatives.

This paper is organized as follows: We conceptualized responsible AI with definitions, principles, and practices in Section II. Section III introduces the data source and outlines the methodologies of intelligent bibliometrics. Section IV presents the results of our study, drawing a series of portraits on responsible AI, regarding its key players and topical landscape, hierarchy, and evolutionary pathways. We concluded this study, discussed technical and political implications, and anticipated future directions in Section V.

II. CONCEPTUALIZATION: RESPONSIBLE AI

The pervasive adoption of AI techniques across public and private sectors has accelerated the interactions between AI and human activities, thereby amplifying the ethical dimensions of AI, e.g., imbalanced training data affecting underrepresented population groups [20] and the exposure and preservation of privacy-sensitive data [21]. Consequently, there has been a growing call for advanced algorithmic development and upgrades to proactively address these widely discussed ethical concerns, and it has garnered positive and constructive responses from the AI community.

This section is to conceptualize the definition and principles of responsible AI by distinguishing it from related concepts such as explainable AI, trustworthy AI, and ethical AI and theoretically elaborating academic debates, industry standards, and domain practices.

A. Responsible AI: Conceptual Analysis

Before the advent of responsible AI, the AI community has shown heightened interest in addressing societal concerns and creating positive societal impact through the rise of several conceptual trends. This section is to analyse three prominent concepts: explainable AI, trustworthy AI, and ethical AI.

Explainable AI highlights an AI system’s capacity to be comprehensible to humans, not only in explaining its decisions, recommendations, and predictions, but also in delineating the process of generating these actions [22]. In the AI community, there has been a concerted effort to devise novel approaches to understand feature sensitivity and influence [23] and navigate the trade-off between model performance and explainability [24]. Often referred to by alternative names such as transparent and accountable AI [25], explainability is frequently intertwined with reliability and trust in the literature [26].

Trustworthy AI emphasizes the safety, security, fairness, and

privacy protection of an AI system, and underscores the crucial role of institutional, software, and hardware mechanisms in constructing a reliable and accessible AI ecosystem [27]. The pursuit of trustworthy AI aligns with a shared interest in detecting and preventing attacks and protecting data privacy [8]. This concept has far-reaching influence in cybersecurity, internet of things, social media, etc. [28]

AI ethics (and ethical AI) transcends the boundaries of the AI community and marks AI’s initiative societal interactions, with the intensive involvement of academia, industry, government, and the public. Unlike a purely technical initiative, AI ethics encapsulates the collective understanding and societal expectation regarding AI and its applications [29]. Some of its core ethical principles include: transparency, justice & fairness, non-maleficence, responsibility, privacy, beneficence, freedom & autonomy, trust, sustainability, dignity, and solidarity [3]. The comprehensive set of ethical is in line with fundamental human values. However, some ongoing debates within the community argue that “we should not yet celebrate consensus around high-level principles that hide deep political and normative disagreement” [30].

B. Responsible AI: Principles and Practices

Responsible AI addresses this argument by transitioning from principles to practices, viewing AI as a socio-technical ecosystem encompassing a comprehensive interplay of techniques, developers, and practitioners [4, 31]. Aligned with the values inherent in explainable AI, trustworthy AI, and ethical AI, this socio-technical ecosystem requires diversity and inclusiveness to foster cross-boundary innovation. Thus, responsible AI goes beyond merely “ticking ethical boxes” or incorporating add-on features in AI systems [31].

Conceptually, the practice of responsible AI follows the management architecture of an AI system, traversing ethical guidelines, risk controls, data governance, and training and education [32]. Notably, Microsoft [6] defined six principles of responsible AI standards: accountability, transparency, fairness, reliability & safety, privacy & security, and inclusiveness. These principles are accompanied by detailed targets specifying the stakeholders to whom each standard applies, executive requirements, and stepwise instructions with milestones. Similarly, addressing specific pain points faced by practitioners, Accenture proposed four pillars for responsible AI implementations: organizational, operational, technical, and reputational [33].

In academia, Cheng, et al. [10] outlined the root causes of AI’s responsible issues, such as data bias (e.g., inconsistency, lack of transparency, and imbalance), over-simplified loss functions, inappropriate evaluation metrics, and mis-interpretation. Despite the absence of a unanimous consensus, the AI community has undertaken substantial efforts to individually address some of these issues, e.g., adversarial learning to foster fair representation by mitigating biases in features [34], deep uncertainty quantification [35], and self-

explainable graph neural networks [36].

In this study, we adhered to the definition and interpretation of responsible AI put forth by Dignum [31]. However, we augmented this perspective by integrating the Microsoft responsible AI standards [6] and the conceptualization of socially responsible AI [10]. This approach enables us to encompass the viewpoints from both industry and academia. We, thus, defined:

- **Responsible AI** is a socio-technical ecosystem with AI techniques, developers, and practitioners. The system enables AI with fundamental functionalities to reason about and act according to human values, and fosters accountability and awareness among AI developers and practitioners regarding AI’s societal impact.
- **The responsibility principles** include accountability, explainability, transparency, fairness, intelligibility, unbiased, non-discrimination, reliability, safety, privacy, security, inclusiveness, and accessibility.

As a conclusion of the conceptual study, we presented the comparison between responsible AI (RAI) and related concepts such as explainable AI (XAI), trustworthy AI (TAI), and ethical AI (EAI) in Table I. We emphasised that responsible AI acts like an umbrella of existing concepts but offers an extension synergising conceptual manifestos and practical guidelines.

TABLE I
CONCEPTUAL COMPARISON BETWEEN RESPONSIBLE AI AND RELATED CONCEPTS

Con. ¹	Keywords	Dis. ²	Highlights
XAI	Explainability; Transparency	Inter	Model capabilities to explain decisions
TAI	Trust; Safety; Privacy; Security	Inter	Reliable and secure models to preserve privacy
EAI	Ethical principles	Cross	Manifestos engaging cross-disciplinary and societal principles
RAI	Responsible practices; Safe	Cross	Manifestos and practical guidelines across models, implementations, and governance

Note: 1) Con. refers to concepts; 2) Dis. refers to disciplinary studies, where inter and cross represent interdisciplinary and cross-disciplinary studies.

III. DATA AND METHODOLOGY

This study is to investigate the AI community’s efforts on responsible AI, including proposing principles and standards, developing responsible models, and implementing applications to tackle real-world responsibility-related concerns. To achieve this, we applied a search strategy for collecting and cross-referencing relevant research articles and developed an intelligent bibliometrics-based analytical framework to analyse the data and generate a comprehensive set of portraits on responsible AI.

A. Data

Recognized as two primary sources of scientific databases, the *Web of Science* (WoS) Core Collection² and *Scopus*³ offer a comprehensive coverage of journal articles and conference

² <https://clarivate.com/products/scientific-and-academic-research/research-discovery-and-workflow-solutions/webofscience-platform/web-of-science-core-collection/>

³ <https://www.scopus.com/home.uri>

papers, each with distinct emphases. The WoS Core Collection highlights curating high-quality journals indexed by the Science Citation Index and Social Science Citation Index, while Scopus enables a broader scope by encompassing a diverse array of conference proceedings, a platform where the AI community actively disseminates research findings. Therefore, we decided to leverage both data sources to gain their complementary strengths. Referring to the definition of responsible AI and its principles, we built the skeleton of the search string as follows:

responsib OR accountab* OR explainab* OR transparan* OR fair* OR intelligib* OR bias OR discriminat* OR reliab* OR safety OR privacy OR security OR inclusive* OR accessib**

Despite the extensive history of AI, we chose to commence our search from January 1, 2015, for the following reasons:

- Deep learning triggered the AI boom around 2011, but societal concerns regarding AI ethics did not materialize until +at least 2015 [37].
- The earliest literature on responsible AI appears to emerge around 2017 [38].

We conducted our search on October 24, 2023, and collected 12,952 articles from Scopus and 10,043 articles from the WoS Core Collection. Subsequently, we merged the two datasets and eliminated duplicates with the same Digital Object Identifier (DOI) – an article’s unique ID. While we retained articles without a DOI here, the process resulted in a refined collection of 18,298 distinct articles. The detailed search strategy is in Table II, with the following notes:

- In contrast to social sciences and humanities, the titles of research articles in computer science are typically direct and straightforward, devoid of rhetorical devices, artistic references, or embellishments. However, abstracts may introduce extraneous information such as background details and practical implications, potentially adding noise to the search results. Thus, our search focused exclusively on titles and keywords.
- Scopus utilizes the All Science Journal Classification (ASJC) system to categories journals and conference proceedings, linking them to related articles. The WoS Core Collection features its own subject category system. While both systems contain an AI-related category, they apply distinct classification criteria for journals and conferences. This motivated our decision to utilize both databases to ensure a comprehensive coverage.
- We concentrated on research articles published in journals and conference proceedings and excluded other source types (e.g., books, book series, and reports) and document types (e.g., editorials, review articles, notes, and letters).
- We only focused on English publications.

TABLE II
SEARCH STRATEGY

<i>Constraints</i>	<i>Scopus</i>	<i>WoS Core Collection</i>
--------------------	---------------	----------------------------

Search fields	titles and keywords	
Language	English	
Time	01/01/2015 – 24/10/2023	
Subject area/category	computer science – artificial intelligence	computer science, artificial intelligence
Document type	article OR conference paper	article OR proceeding paper
Source type	journal OR conference proceeding	n/a
Raw results	12,952	10,043
Combined results	18,298	

Academic knowledge graphs, with structured bibliographical attributes (e.g., titles, abstracts, authors, and references), disambiguated name entities (e.g., country, institution, and author names), and predefined document-topic linkages, prove highly advantageous in the pre-processing of bibliometric data. OpenAlex⁴ is an open and free-access academic knowledge graph, with a comprehensive integration of the Microsoft Academic Graph (MAG; one representative academic graph produced by Microsoft), Crossref, and some other scientific data sources (e.g., PubMed). OpenAlex inherits the hierarchical topical system known as the Field of Study (FoS) from MAG. This system consists of a series of topical tags created by hierarchical topic models [39], with each article containing at least one FoS tag.

In this study, we leveraged OpenAlex APIs to match the 18,298 combined articles from Scopus and WoS with their DOIs, returning 17,799 articles, with the missing ones primarily attributed to the absence of DOIs and the early publications of 2024. We, then, retrieved the complete bibliographical information of the 17,799 articles and extracted key entities, including 49,792 authors, 6,068 affiliations, 137 countries and regions, and 8,523 FoS tags.

B. Intelligent Bibliometrics

Intelligent bibliometrics [14] refers to the development and applications of computational models that elaborate AI and data science techniques with bibliometric data (e.g., scientific articles and patent documents) and indicators (e.g., citations, keywords, and authorships).

Intelligent bibliometrics primarily targets real-world science, technology, and innovation (ST&I) challenges, and it includes diverse methodologies, such as embedding-based ST&I topic extraction [40], network analytics and graph learning for ST&I measurements [41, 42] and relationship discovery [16, 43], and prediction models for technological forecasting [17, 44].

This study employed the following intelligent bibliometric models, in conjunction with descriptive statistics and co-occurrence analyses, which identify similar items (e.g., keywords and authors) by assessing their frequency of occurrence together.

⁴ <https://openalex.org/about>

i. Hierarchical Topic Tree

The Hierarchical Topic Tree (HTT) model is a network-based method of hierarchical community detection, designed to unveil research topics and their inherent hierarchical relationships within a corpus of research articles [16]. This model operates on a weighted co-term network as its input and identifies a collection of community anchors based on two distinct topological characteristics: 1) a higher density than its neighbours and 2) a larger distance with other high-density nodes than the distance to its up-level anchor node. Each anchor subsequently forms the nucleus of a community, with the remaining nodes allocated to the community based on the shortest topological distance. This iterative process continues, generating sub-communities until no discernible anchors remain. The sub-communities established at each iteration contribute to the construction of topical layers, with the anchors serving as representative labels.

While the number of first-layer topics is determined by a super parameter k – the top k nodes with the highest density and can maximise the average density of its related communities since they do not have any up-level anchors. The following topics and layers are decided automatically based on the two given characteristics without additional parameters preset.

Ultimately, the outcome is a hierarchically partitioned co-term network that reflects the layered intellectual structure within a field of knowledge.

ii. Scientific Evolutionary Pathways

The innovation literature posits that scientific innovation is an accumulative process driven by the progression of knowledge [45]. The Scientific Evolutionary Pathways (SEP) model was designed to track this accumulative process by identifying a predecessor-descendent relationship between a newly created topic and previous ones [15]. The SEP model organizes articles with the same publication year into one time slice, treating the entire dataset as a bibliometric stream. Each topic is defined as a collection of articles, and the model analyses each article in the stream by classifying it into its most similar topic (i.e., predecessor). It determines whether the article contributes new knowledge by assessing its semantic distance from the centroid of the topic, and then groups articles identified as contributing new knowledge into several new topics, i.e., descendants. Similar to the HTT, the only super parameter of the SEP model is the initial number of topics, i.e., the starting points of the evolution, but a K-means clustering algorithm with the elbow method will help identify the most optimal k in a given interval.

The SEP model has been empirically evaluated with expert knowledge through a wide range of cases, and its performance on tracking the historical trends of scientific and technological topics has been widely recognised, e.g., the 243,606 academic proposals granted by the US National Science Foundation between 1976 and 2016 [15], the 142,877 research articles on gene-related cardiovascular diseases with a domain expert-led

cross-examination between 2008 and 2019 [46], and 2,657 articles published by Knowledge-based Systems between 1991 and 2016 [47].

Utilizing Gephi [48], the SEP model visualizes predecessor-descendent relationships in a network:

- A node represents a topic with its size proportional to the number of articles in the topic.
- An edge represents the predecessor-descendent relationship between its connected nodes. Its weight is determined by the similarity between the two topics.
- Nodes in the same colour signify their belonging to the same community. The assignment of colours is based on a modularity-based community detection algorithm integrated into Gephi.

IV. BIBLIOMETRIC PORTRAITS OF RESPONSIBLE AI

We constructed the bibliometric portraits of responsible AI from four aspects: Key technological players and their interactions with research communities and the responsibility principles, the topical hierarchy of responsible AI and its historical dynamics, the interplay between primary AI techniques and the responsibility principles, and the profile of a core collection of the responsible AI research.

Given a bird's eye view, Fig. 1 illustrates the publication trends of the 17,799 articles related to responsible AI and contributed by the AI community since 2015. As anticipated, a discernible upward trend emerges starting in 2015 and accelerating after 2020. The incomplete data for 2023 would be an exclusion in this chart.

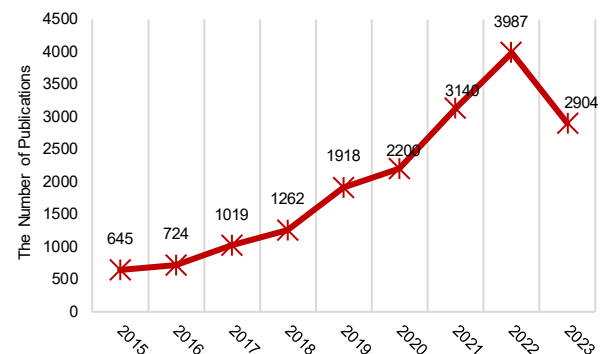


Fig. 1. Publication Trends of Responsible AI-related Articles since 2015.

A. Key Technological Players and Interactions

We retrieved organisation names linked with the 17,799 research articles and measured their productivity based on the total number of their published research articles. Research productivity is a mature indicator and its direct connection with the quantity⁵ of research publications has been widely recognised in bibliometrics [14, 49].

Defining key technological players as the most productive countries/regions, institutions, and research communities in a specific scientific or technological area, Fig. 2 paints a

pros and cons, e.g., citation counts vary with the time, and survey articles naturally attract higher citations. This study thus skipped off the quality measurement to avoid unnecessary arguments.

⁵ Comparably, measuring research quality has been debated for decades. Related indicators could be related to citations (for research impact), semantic similarities with previous articles (for research novelty), etc., but each contains

geographical portrait of the key players in the responsible AI research. At the forefront, China and the USA lead the game as dominant players, collectively contributing to over 40% of the total publications. India, the UK, and Germany follow, each contributing more than 1000 publications to the field. The 10 most productive countries collectively account for 86.5% of the research articles on responsible AI, aligning with the Pareto principle.

Table III lists the top 15 most productive institutions in responsible AI research. In conjunction with Fig. 2, an observation is drawn regarding the prevalence of universities in contributing to the topic, but three national research institutions stand out: the French National Centre for Scientific Research (France), the Chinese Academy of Sciences (China), and the Commonwealth Scientific and Industrial Research Organization (CSIRO, Australia). This presence suggests potential national interests in responsible AI research. For example, as a possible response to China’s Ministry of Science and Technology’s release of the *Governance Principles for a New Generation of Artificial Intelligence: Develop Responsible Artificial Intelligence* in June 2019⁶, Chinese universities have notably surged ahead in Table III.

TABLE III
TOP 10 MOST PRODUCTIVE INSTITUTIONS IN RESPONSIBLE AI RESEARCH

	<i>Institutions</i>	<i>Countries</i>	<i>#P</i>
1	French National Centre for Scientific Research	France	183
2	Xidian University	China	156
3	Chinese Academy of Sciences	China	151
4	Carnegie Mellon University	USA	122
5	University of Technology Sydney	Australia	121
6	University of Oxford	UK	108
7	Tsinghua University	China	106
8	Delft University of Technology	Netherland	105
9	University of Electronic Science and Technology of China	China	105
10	University of Science and Technology of China	China	105

Notes. #P: The number of publications.

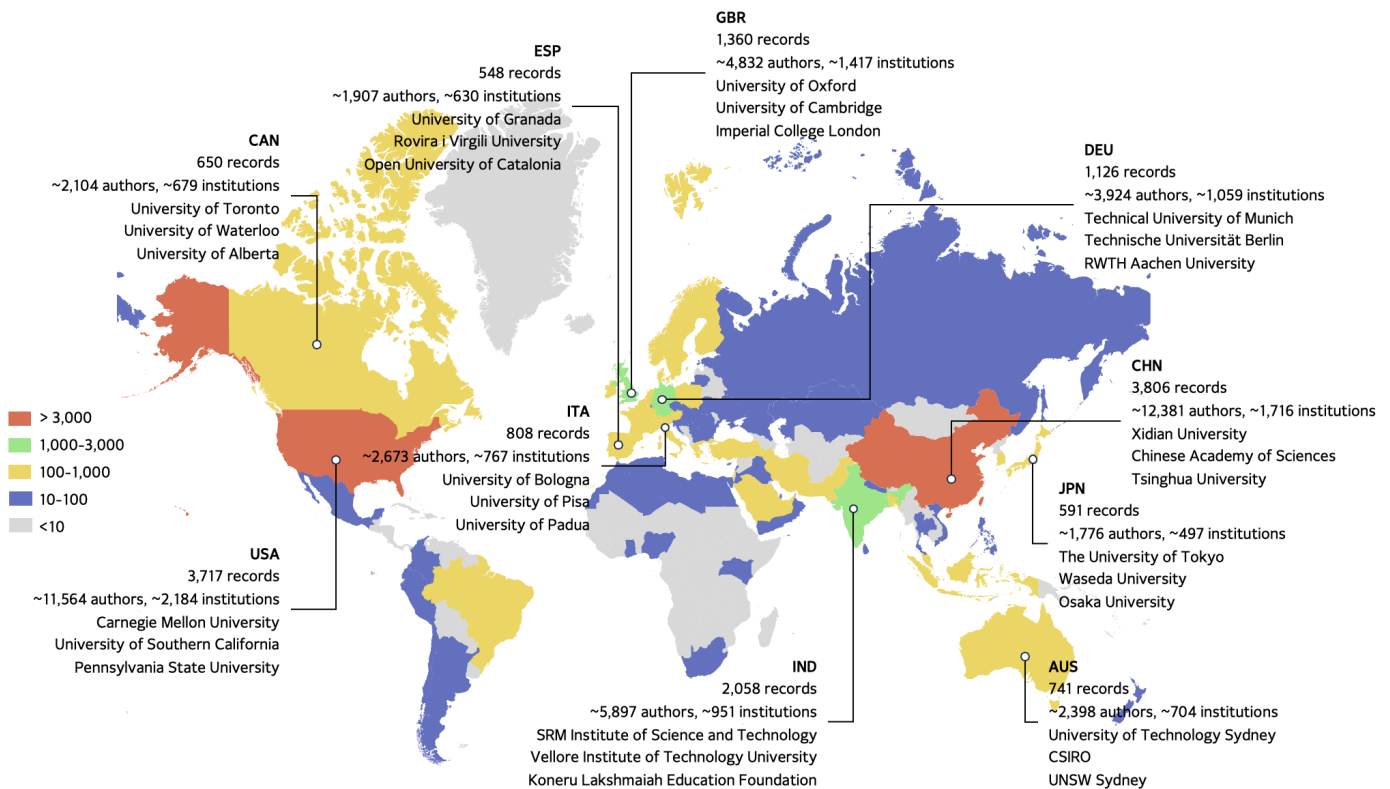


Fig. 2. Geographical Portrait of the Key Technological Players in Responsible AI Research.

Our investigation into research communities extends to the analysis of involved journals and conference proceedings to responsible AI. Table IV presents the top 15 journals and conference proceedings, which have published the largest

number of articles on responsible AI. Considering AI’s subareas, there is a notable representation of journals spanning various facets of AI: *IEEE TNN* and *Artificial Intelligence* for theoretical developments in AI, particularly in machine

⁶ See the details from the website of the US Library of Congress: <https://www.loc.gov/item/global-legal-monitor/2019-09-09/china-ai-governance-principles-released/>

learning studies; *Information Sciences* and *Knowledge-based Systems* for AI’s technological innovation and system applications; and *IEEE TKDE* and *Pattern Recognition* for broad topics related to data mining and knowledge discovery. Despite our primary focus on the AI community, the inclusion of the journal *AI & Society* highlights the cross-disciplinary nature of responsible AI, showcasing active engagements with studies in humanities and social sciences.

In terms of conference proceedings, with their annual nature, the number of publications linked to one proceeding is comparable to as that of journals. However, a discernible surge in responsible AI research is observed after 2021, with the computer vision and communications societies emerging as the most active research communities for this new topic. This trend sheds light on the rising prominence of journals in machine learning and intelligent systems.

TABLE IV
TOP 15 JOURNALS AND CONFERENCE PROCEEDINGS WITH THE LARGEST NUMBER OF RESPONSIBLE AI-RELATED PUBLICATIONS

Journals	#Pub.	Conferences	#Pub.
Information Sciences	362	2021 IEEE BigData	55
AI & SOCIETY	280	2023 IEEE/CVF WACV	43
Expert Systems with Applications	234	2022 IEEE GLOBECOM	40
Knowledge Based Systems	179	2021 IEEE GLOBECOM	40
Neurocomputing	160	2022 IJCNN	39
IEEE Transactions on Knowledge and Data Engineering (TKDE)	109	2022 ICASSP	36
Journal of Intelligent and Fuzzy Systems	99	2021 IEEE TrustCom	30
Journal of Ambient Intelligence and Humanized Computing	92	2022 IEEE/CVF CVPR	30
Soft Computing	91	2021 IEEE SSCI	26
IEEE Transactions on Neural Networks and Learning Systems (TNN)	87	2021 IEEE/CVF ICCV	22
Artificial Intelligence	86	2022 ACM WSDM	21
Applied Intelligence	84	2022 ACM/IEEE HRI	20
Pattern Recognition	83	2022 GECCO	20

Frontiers in Artificial Intelligence	80	14	2022 ACM MM	19
SN Computer Science	76	15	2022 IEEE ICDE	18

To portray the core research community of responsible AI, we linked journals (Fig. 3) and conference proceedings (Fig. 4) with a set of selected responsibility related FoS tags in OpenAlex. While the distribution of involved countries and publication venues mirrors that of the entire dataset, a distinct pattern emerges - most articles in this collection cluster around topics “privacy” and “security”, with only a few articles scattered across topics like “explainability”, “transparency”, and “trustworthiness”. This is consistent with the involvement of the communications society in the realm of cybersecurity and privacy, suggesting that these issues may have triggered the initial interaction between the AI community and real-world concerns related to responsibility. Simultaneously, this finding motivates us to delve deeper into the profiling of responsible AI through a series of topic analyses.

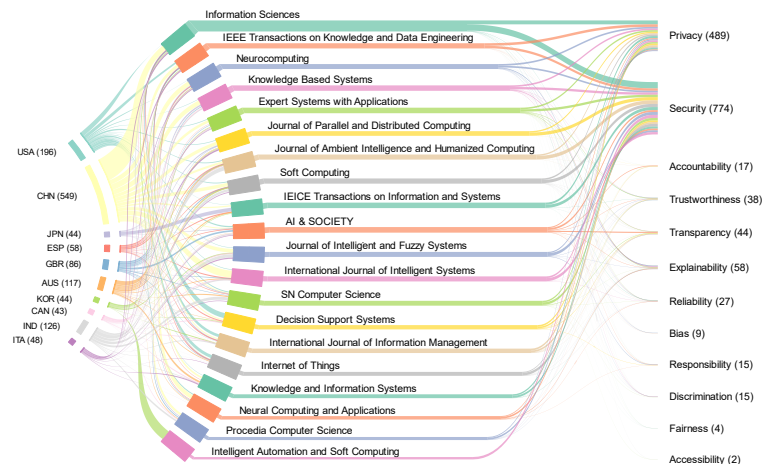


Fig. 3. Portraits of the Responsible AI Community with Journals.

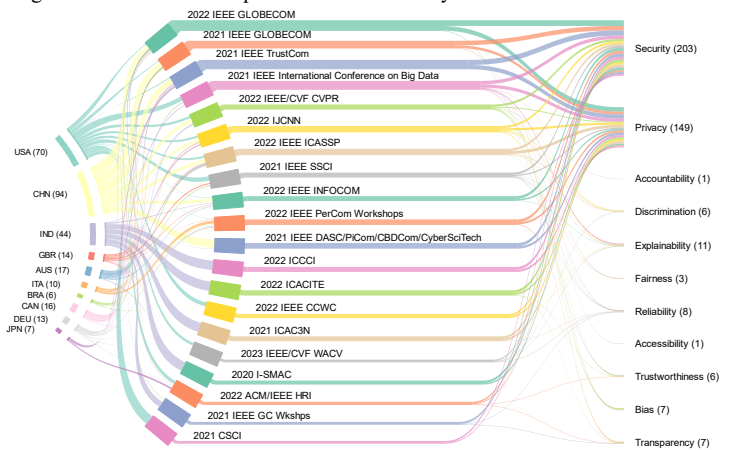


Fig. 4. Portraits of the Responsible AI Community with Conference Proceedings.

B. Topical Hierarchy and Dynamics of Responsible AI

Leveraging OpenAlex’s FoS tags to enable a pre-cleaned hierarchical topical system for detailed topic analyses, we

utilized an article's FoS tags as linked topics and constructed a FoS-based co-occurrence network for the application of the HTT and SEP models: HTT generates a data-driven topical hierarchy, offering insights into the knowledge structure of responsible AI research. SEP identifies predecessor-descendant relationships between research topics, and these evolutionary relationships provide understandings to the historical dynamics of responsible AI research, e.g., the origin of the field, the timeline of involved concepts and technologies, and the potential future directions it might evolve.

Fig. 5 shows the hierarchical topic tree of responsible AI research, delineating AI's knowledge structure into four fundamental components: machine learning, mathematics, data mining, and computer networks. We decided these four initial components (the super parameter k , described in Section III-B-i) based on the seven computer-science-related WoS Categories across mathematic theories, AI, hardware and software, and applications, as well as several round empirical examinations. Note that HTT highlights the discovery of a hierarchical relationship between two research topics within a specific dataset, and thus, this hierarchy does not necessarily indicate a sense of a belongingness between the components.

We specifically annotated the following observations:

- *Machine Learning*: This branch draws a clear hierarchical structure of the machine learning research, consisting of neural networks⁷, deep learning, and their downstream applications, e.g., human-computer interactions, bioinformatics, and recommender systems. Notably, we observed several strong connections with some responsibility issues, e.g., accountability when applying to knowledge management and specific tasks like recommender systems, interpretability for inference-related studies, and trustworthiness when performing human-computer interactions.
- *Data Mining*: This branch is application-focused, highlighting the broad applications of AI and data science techniques in real-world cases, e.g., database management, information retrieval, cyber-physical systems, and governance. Information privacy is a key concern, closely linked with internet privacy, information sensitivity and confidentiality, information leakage, privacy software, and personal identities. In addition, interoperability is involved with database and data modelling.
- *Computer Networks*: This branch groups topics in distributed systems, cloud computing, and telecommunications, which is the key contributor to privacy and security related studies in software engineering and electrical engineering.
- *Mathematics*: This branch highlights the theoretical base of computer science and contains topics such as statistics, parametric statistics, and Gaussian.

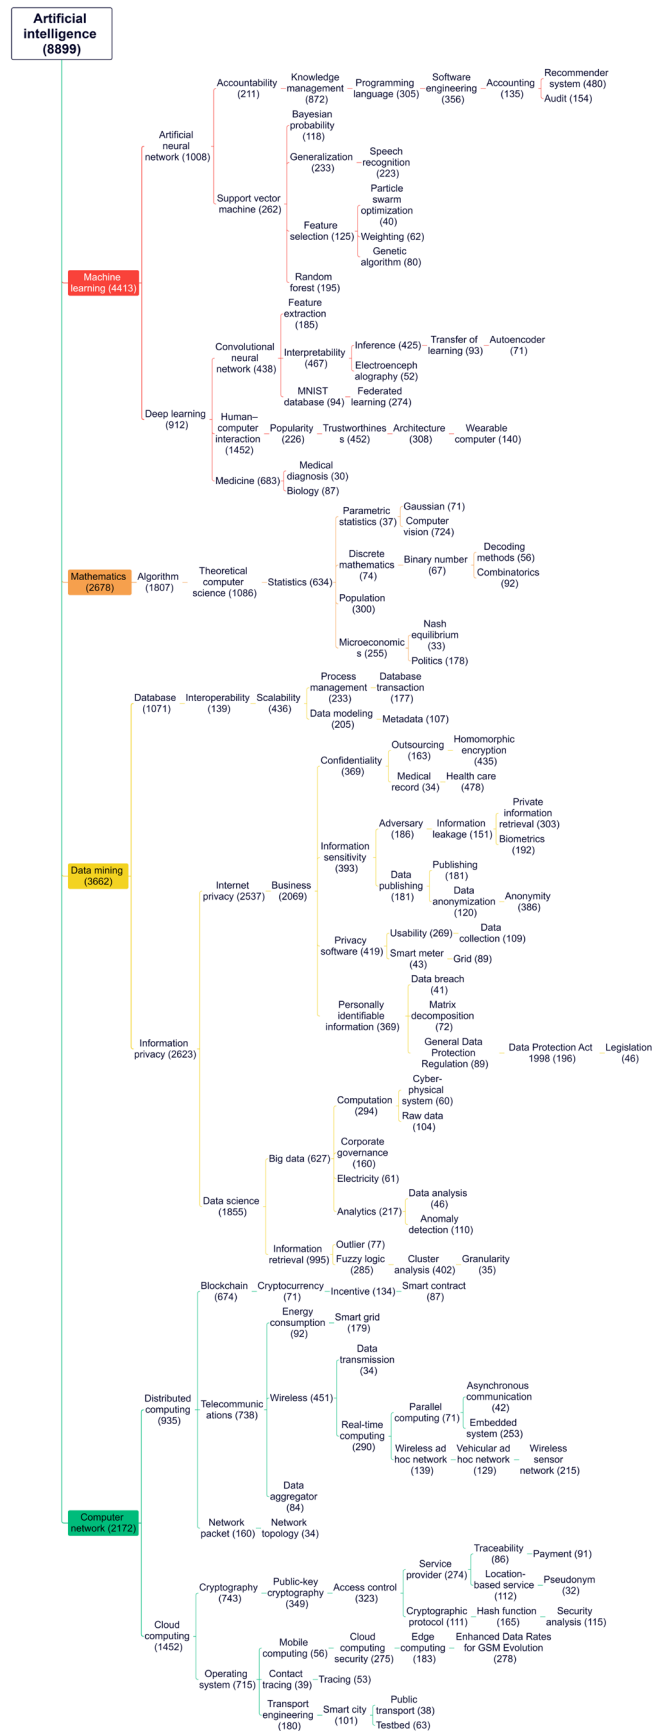


Fig. 5. Hierarchical Topic Tree of Responsible AI Research.

⁷ The topic of support vector machines (SVM) under neural networks can be explained as, after 2015, SVM, together with techniques in its leaf nodes (e.g.,

random forests, Bayesian, and genetic algorithms), frequently co-occurs with neural networks, as comparative baselines, for example.

Fig. 6 illustrates the evolutionary pathways of responsible AI research spanning the years 2015 to 2023. We employed a K-means clustering algorithm on articles published in 2015. Utilizing the elbow method, we identified five topics – artificial intelligence, computer security, internet privacy, encryption, and computer network – as the starting points of the evolutionary pathways. These topics serve as a technical profile of responsible AI before its formal terminology and concepts were established. These initially distinct technologies have gradually coalesced into their respective technological cohorts since 2015, forming the foundational clusters that construct the technological landscape of responsible AI.

We discussed the following observations:

- *Machine Learning and Statistics* (pink cluster): This cluster reveals the evolution of machine learning-based algorithms, techniques, and applications from 2015 to 2023. Key sub-technologies and their evolutionary pathways encompass neural networks, human-computer interaction, robots, and various AI algorithms. Notably, interpretability has emerged as one of the most critical concerns within the context of neural networks.
- *Data and Information Privacy* (blue cluster): This cluster highlights the evolution of privacy-related topics and their interactions, e.g., the mechanisms of privacy

protection in blockchains, privacy safeguards in various data and information sources (e.g., the web, social media, and healthcare records), and some regulations of data protection. Interoperability, trustworthiness, and accessibility are intertwined throughout these pathways.

- *System Security* (green cluster): This cluster progresses from the privacy concerns of data architecture to the security design and implementation of analytical systems, interacting with data science, internet of things, telecommunications, and cloud computing techniques. It emphasizes the development of methods, software, and hardware for privacy and security protection. Some notably techniques include federated learning, adversarial system, private information retrieval, etc.
- *Encryption* (yellow cluster): While relatively small, this cluster draws a clear pathway of cryptography, including both its fundamental techniques and applications.
- *Ethics* (orange cluster): This cluster predominantly addresses AI ethics from the perspective of humanities and social sciences, spanning areas such as business, management science, politics, and law.

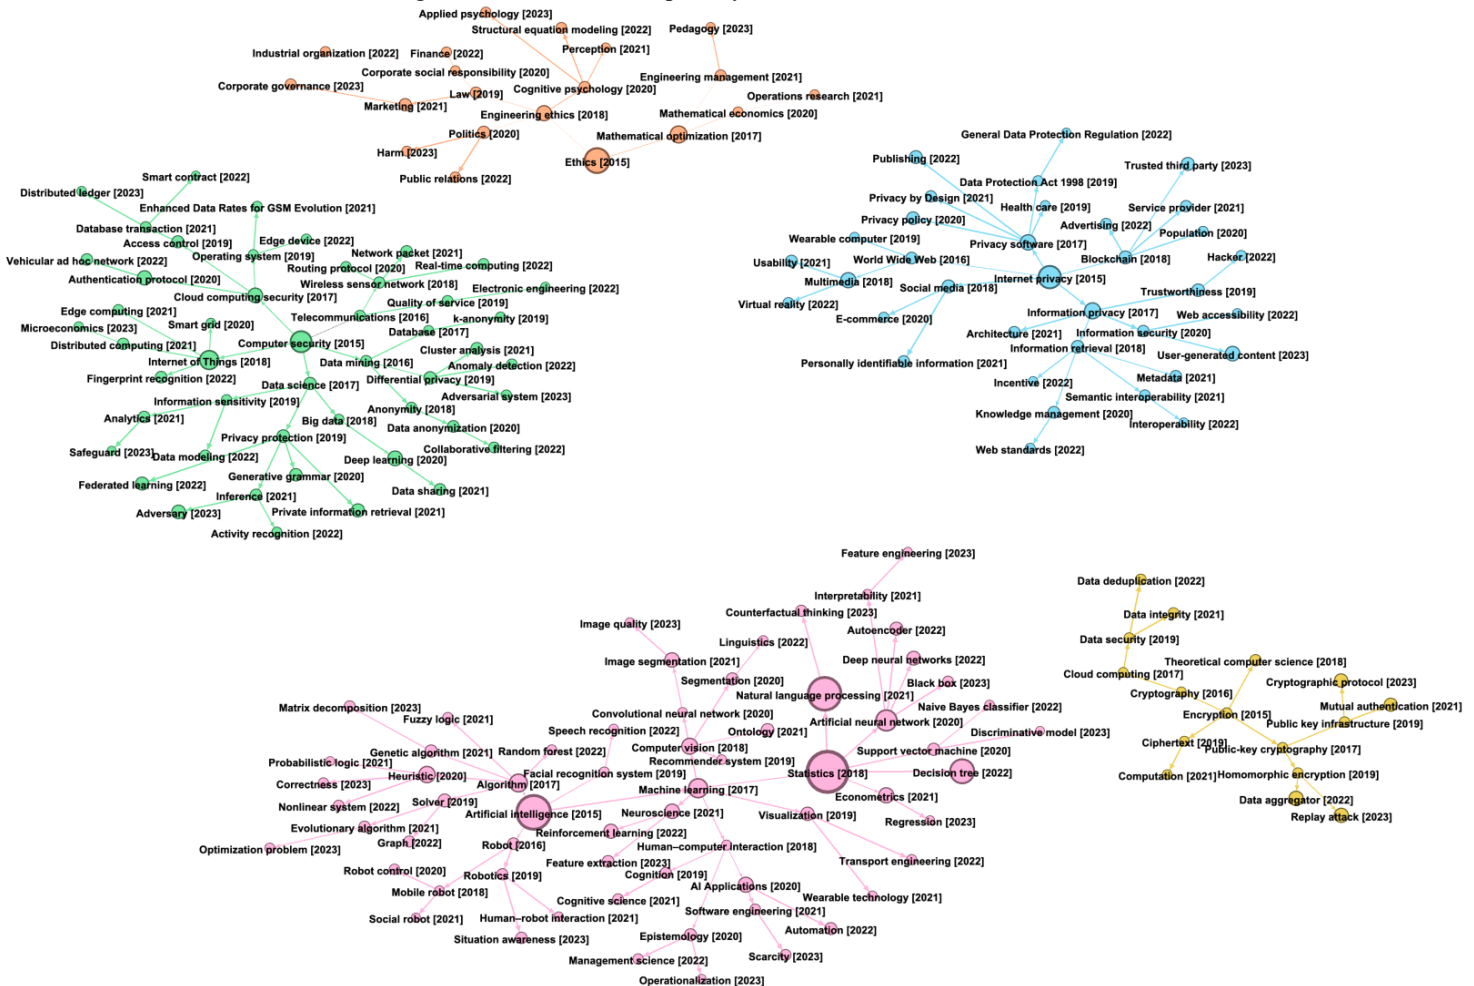


Fig. 6. Scientific Evolutionary Pathways of Responsible AI Research (2015 - 2023).

Notes. The year associated with a node shows the time when the node was born, indicating when the related knowledge initially became influential within the entire dataset.

C. Responsibility Principles vs. Primary AI Techniques

Revisiting the responsibility principles defined in Section II-B and aligning them with the FoS tags, we identified 12 principles for mapping. We added “responsibility”, but excluded “intelligibility” due to its extremely low frequency and “inclusiveness” since it mainly aligns with social science topics when existing computer science literature typically addresses similar concepts under the term “fairness” [50]. Then, we selected 16 primary AI techniques based on their frequency in our dataset. To visualize the co-occurrence between the responsibility principles and primary AI techniques, we utilized the Circos approach [51], as illustrated in Fig. 7.

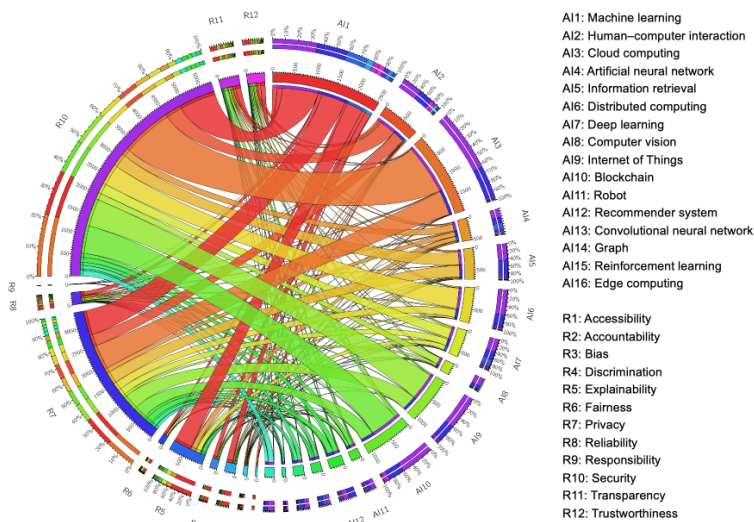


Fig. 7. Interplays between 12 responsibility principles and 16 selected primary AI techniques.

Note. The statistics is in Supplementary Table 1.

We derived the following key observations:

- As depicted in Fig. 5 and Fig. 6, the initial studies on responsible AI were primarily centred around privacy and security. Techniques related to cloud computing, machine learning, internet of things, and blockchains contribute more than 50% of related research.
- Machine learning, including its sub-techniques such as neural networks, deep learning, and reinforcement learning, has directly engaged with all principles but demonstrated a tangible connection with explainability (86%), discrimination (83%), bias (74%), and fairness (52%). Clearly, this is consistent with the community’s continuous interests on explainable AI and ethical AI.
- Human-computer interaction occupies the accessibility principle (83%) and is also the second largest contributor to trustworthiness and transparency.
- Some significant matching pairs also include cloud computing with accountability, distributed computing with fairness, and blockchain with accountability and transparency.
- Interestingly, the responsibility principle is linked with

machine learning (31%), robot (31%), and blockchain (15%), indicating some of the AI backbones that initiate the topic of responsible AI.

D. Core Cohort of Responsible AI Research

While our study primarily relied on the dataset of responsible AI-related articles contributed by the AI community since 2015, we maintained a keen interest in exploring the core concept of responsible AI and its developmental progress. Therefore, with a specific focus on the WoS Core Collection to ensure the high standards of coverage, on November 23, 2023, we curated 380 articles directly employing the term “responsible artificial intelligence” or “responsible AI” in their titles, abstracts, and keywords.

Fig. 8 draws a series of portraits for the core cohort, illustrating its publication trend (Fig. 8A), article types (Fig. 8B), and WoS subject categories (Fig. 8C). Intriguingly, the term “responsible artificial intelligence” was firstly mentioned in 2015 [52]⁸, but as evidenced in Fig. 8A, academia did not commence investigations into responsible AI until 2019 and a significant upswing started in 2021. Compared to the large volume of research articles on AI, responsible AI remains a relatively recent and evolving topic.

Certainly, the computer science community is the primary contributor to the core cohort. The category “computer science, artificial intelligence” leads the ranking, comprising 128 articles (33.7%), and the six sub-categories within computer science collectively encompass 275 articles, accounting for 72.4% of the cohort. Remarkably, the inclusion of categories such as law, ethics, social sciences, and information sciences underscores the cross-disciplinary nature of responsible AI research.

Table V presents the top 15 most productive institutions in responsible AI research within the core cohort. The USA takes the lead, followed by the UK, Australia, and the Netherlands. Interestingly, China is conspicuously absent from the list. Given that the core cohort reflects the most recent conceptualization of responsible AI rather than its fundamental technologies and sub-technologies, the dominance of the West, supported by their international-renowned universities, has demonstrated their capabilities in shaping innovative theories and concepts and driving global technological advancements. In addition, Microsoft emerges as one of the most productive institutions, uniquely representing the corporate sector and signalling its ambition in this cutting-edge domain.

TABLE V

TOP 15 MOST PRODUCTIVE INSTITUTIONS ON RESPONSIBLE AI RESEARCH IN THE CORE COHORT

	<i>Affiliations</i>	<i>Country</i>	<i>#Pub.</i>
1	The University of Texas Austin	USA	16
2	CSIRO	Australia	13
3	Umea University	Sweden	11
4	Northeastern University	USA	10

⁸ We conducted a similar search on Google Scholar and no results were retrieved for articles predating the year 2015.

5	University of Oxford	UK	9	11	University of New South Wales	Australia	6
6	Microsoft	USA	8		Cornell University	USA	
	University of Washington	USA			Stanford University	USA	
8	Norwegian University of Science and Technology	Norway	7		University of Sheffield	UK	
	University of Bologna	Italy			Delft University of Technology	Netherlands	
	University of Amsterdam	Netherlands			University of Cambridge	UK	

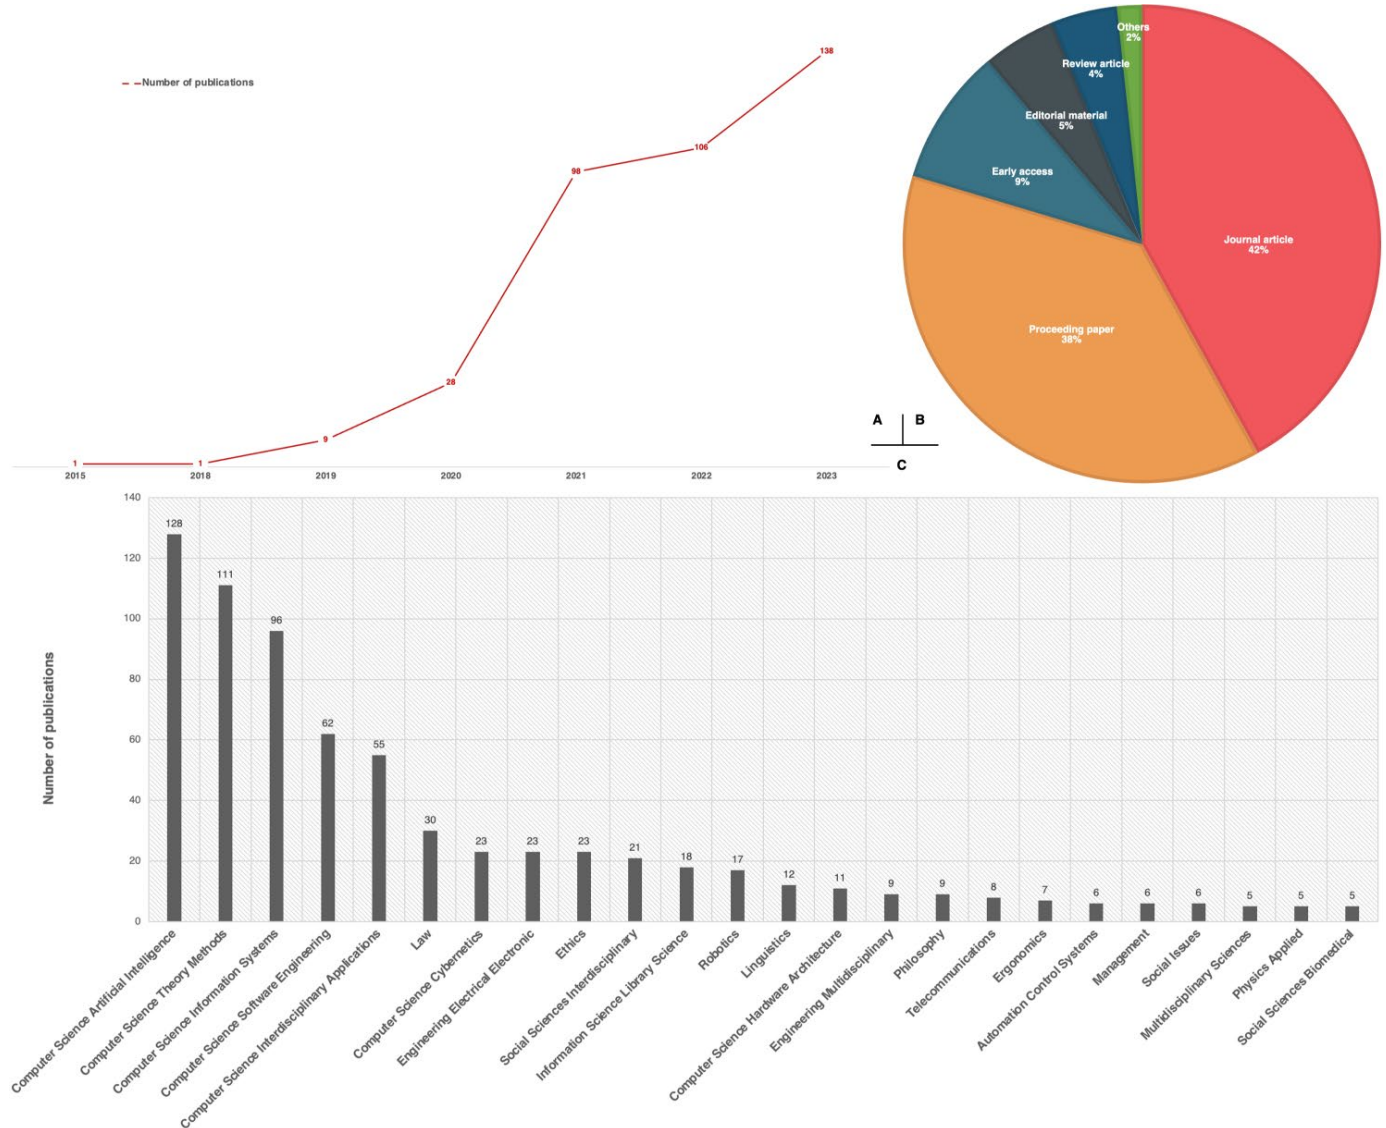


Fig. 8. Portraits of the Core Cohort of Responsible AI Research: A) Publication Trend, B) Distribution of Article Types; and C) Distribution of the WoS Subject Categories (with at least 5 articles).

Notes. 1) An article may have more than one article type, and thus, Fig. 8B exhibit overlaps. 2) While the core cohort spans across 89 WoS subject categories, the 24 subject categories depicted in Fig 7C encapsulate 334 articles – 87.9% of the core cohort.

Aiming to understand the content of the core cohort, we employed the KeyBERT model [53] to automatically extract 1257 terms from the titles and abstracts of the 380 articles, and created a co-term network using VoSViewer [54], see Fig. 9. Slightly different from previous portraits that feature a large amount of AI-related technical terms, Fig. 9 unveils AI’s extensive engagements across multiple disciplines and research areas, e.g., medicine and healthcare (with bioethics and CoVID-19), governance (with industry, economy, law, policy, and

sustainability), etc. This, again, highlights the cross-disciplinary nature of responsible AI, transitioning from fundamental AI techniques and algorithms to a high-level perspective encompassing the design, development, implementation, and regulation of AI in close coordination with academia, industry, and society. Particularly intriguing is the relatively scattered knowledge landscape in Fig. 9, without clearly defined boundaries and communities. This aligns with the theory of complex systems and their transitions, indicating

and thus suffer from accuracy issues. LLMs can augment external information for broad bibliometric entities (e.g., research topics, authors, and organisations), and help achieve better knowledge representations and benefit downstream tasks, including HTT and SEP.

ACKNOWLEDGMENT

The authors acknowledge the use of ChatGPT for polishing the language of this paper. The authors clarify that all aspects

of this work were designed, developed, implemented, and reported by the authors, without any contributions from AI-generated content.

SUPPLEMENTAL MATERIALS

Supplementary Table 1 and the high-resolution versions of Figs 2, and 5-9 can be retrieved through:

<https://github.com/IntelligentBibliometrics/Responsible-AI-Review>

case of artificial intelligence," *Advanced Engineering Informatics*, p. to appear, 2021.

REFERENCES

- [1] C. Metz. "What Exactly Are the Dangers Posed by A.I.?" *The New York Times*. <https://www.nytimes.com/2023/05/01/technology/ai-problems-danger-chatgpt.html> (accessed September 18, 2023, 2023).
- [2] Y. Yakimova and J. Ojamo. "AI Act: a step closer to the first rules on Artificial Intelligence." *European Parliament News*. <https://www.europarl.europa.eu/news/en/press-room/20230505IPR84904/ai-act-a-step-closer-to-the-first-rules-on-artificial-intelligence> (accessed September 18, 2023, 2023).
- [3] A. Jobin, M. Ienca, and E. Vayena. "The global landscape of AI ethics guidelines," *Nature Machine Intelligence*, vol. 1, no. 9, pp. 389-399, 2019.
- [4] V. Dignum. *Responsible artificial intelligence: how to develop and use AI in a responsible way*. Springer, 2019.
- [5] A. S. Gillis. "Responsible AI." *TechTarget*. <https://www.techtarget.com/searchenterpriseai/definition/responsible-AI> (accessed September 18, 2023, 2023).
- [6] Microsoft, "Microsoft Responsible AI Standard, V2," Microsoft, 2022.
- [7] R. Eitel-Porter, M. Tricarico, and R. Bianchi. "AI ethics & governance." *Accenture*. <https://www.accenture.com/au-en/services/applied-intelligence/ai-ethics-governance> (accessed September 18, 2023, 2023).
- [8] B. Eshete, "Making machine learning trustworthy," *Science*, vol. 373, no. 6556, pp. 743-744, 2021.
- [9] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," *ACM Computing Surveys (CSUR)*, vol. 54, no. 6, pp. 1-35, 2021.
- [10] L. Cheng, K. R. Varshney, and H. Liu, "Socially responsible ai algorithms: Issues, purposes, and challenges," *Journal of Artificial Intelligence Research*, vol. 71, pp. 1137-1181, 2021.
- [11] J. Madhavan, J. A. Kerr, A. R. Corcos, and B. P. Isaacoff, "Toward trustworthy and responsible artificial intelligence policy development," *IEEE Intelligent Systems*, vol. 35, no. 5, pp. 103-108, 2020.
- [12] A. Tzachor, M. Devare, B. King, S. Avin, and S. Ó hÉigeartaigh, "Responsible artificial intelligence in agriculture requires systemic understanding of risks and externalities," *Nature Machine Intelligence*, vol. 4, no. 2, pp. 104-109, 2022.
- [13] J. Wiens *et al.*, "Do no harm: a roadmap for responsible machine learning for health care," *Nature medicine*, vol. 25, no. 9, pp. 1337-1340, 2019.
- [14] Y. Zhang, A. Porter, S. W. Cunningham, D. Chiavetta, and N. Newman, "Parallel or intersecting lines? Intelligent bibliometrics for investigating the involvement of data science in policy analysis," *IEEE Transactions on Engineering Management*, vol. 68, no. 5, pp. 1259-1271, 2020.
- [15] Y. Zhang, G. Zhang, D. Zhu, and J. Lu, "Scientific evolutionary pathways: Identifying and visualizing relationships for scientific topics," *Journal of the Association for Information Science and Technology*, vol. 68, no. 8, pp. 1925-1939, 2017.
- [16] M. Wu and Y. Zhang, "Hierarchical topic tree: A hybrid model comprising network analysis and density peak search," in *18th International Conference on Scientometrics and Informetrics (ISSI)*, 2021: INT SOC SCIENTOMETRICS & INFORMETRICS-ISSI.
- [17] Y. Zhang, M. Wu, G. Zhang, and J. Lu, "Stepping beyond your comfort zone: Diffusion - based network analytics for knowledge trajectory recommendation," *Journal of the Association for Information Science and Technology*, vol. 74, no. 7, pp. 775-790, 2023.
- [18] Y. Zhang, X. Cai, C. V. Fry, M. Wu, and C. S. Wagner, "Topic evolution, disruption and resilience in early COVID-19 research," *Scientometrics*, vol. 126, no. 5, pp. 4225-4253, 2021.
- [19] M. Wu, C. D. Kozanoglu, C. Min, and Y. Zhang, "Unraveling digital transformation enabling capabilities: A data-driven methodology and the
- [20] A. Subramonian, K.-W. Chang, and Y. Sun, "On the discrimination risk of mean aggregation feature imputation in graphs," *Advances in Neural Information Processing Systems*, vol. 35, pp. 32957-32973, 2022.
- [21] G. A. Kaissis, M. R. Makowski, D. Rückert, and R. F. Braren, "Secure, privacy-preserving and federated machine learning in medical imaging," *Nature Machine Intelligence*, vol. 2, no. 6, pp. 305-311, 2020.
- [22] D. Gunning, M. Stefik, J. Choi, T. Miller, S. Stumpf, and G.-Z. Yang, "XAI—Explainable artificial intelligence," *Science robotics*, vol. 4, no. 37, p. eaay7120, 2019.
- [23] H. Yuan, H. Yu, S. Gui, and S. Ji, "Explainability in graph neural networks: A taxonomic survey," *IEEE transactions on pattern analysis and machine intelligence*, vol. 45, no. 5, pp. 5782-5799, 2022.
- [24] S. M. Lundberg *et al.*, "From local explanations to global understanding with explainable AI for trees," *Nature machine intelligence*, vol. 2, no. 1, pp. 56-67, 2020.
- [25] S. Wachter, B. Mittelstadt, and L. Floridi, "Transparent, explainable, and accountable AI for robotics," *Science robotics*, vol. 2, no. 6, p. eaan6080, 2017.
- [26] K. Gade, S. C. Geyik, K. Kenthapadi, V. Mithal, and A. Taly, "Explainable AI in industry," in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2019, pp. 3203-3204.
- [27] M. Brundage *et al.*, "Toward trustworthy AI development: mechanisms for supporting verifiable claims," *arXiv preprint arXiv:2004.07213*, 2020.
- [28] B. Thuraisingham, "Trustworthy machine learning," *IEEE Intelligent Systems*, vol. 37, no. 1, pp. 21-24, 2022.
- [29] Y. Zhang, M. Wu, G. Y. Tian, G. Zhang, and J. Lu, "Ethics and privacy of artificial intelligence: Understandings from bibliometrics," *Knowledge-Based Systems*, vol. 222, p. 106994, 2021.
- [30] B. Mittelstadt, "Principles alone cannot guarantee ethical AI," *Nature machine intelligence*, vol. 1, no. 11, pp. 501-507, 2019.
- [31] V. Dignum, "Responsible Artificial Intelligence—from Principles to Practice," *arXiv preprint arXiv:2205.10785*, 2022.
- [32] Y. Wang, M. Xiong, and H. Olya, "Toward an understanding of responsible artificial intelligence practices," in *Proceedings of the 53rd hawaii international conference on system sciences*, 2020: Hawaii International Conference on System Sciences (HICSS), pp. 4962-4971.
- [33] R. Eitel-Porter, M. Corcoran, and P. Connolly, "Responsible AI: From principles to practice," 2021.
- [34] R. Feng, Y. Yang, Y. Lyu, C. Tan, Y. Sun, and C. Wang, "Learning fair representations via an adversarial framework," *arXiv preprint arXiv:1904.13341*, 2019, doi: <https://doi.org/10.48550/arXiv.1904.13341>.
- [35] B. Wang *et al.*, "Deep uncertainty quantification: A machine learning approach for weather forecasting," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 2087-2095.
- [36] E. Dai and S. Wang, "Towards self-explainable graph neural network," in *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2021, pp. 302-311.
- [37] S. Russell, S. Hauer, R. Altman, and M. Veloso, "Ethics of artificial intelligence," *Nature*, vol. 521, no. 7553, pp. 415-416, 2015.
- [38] V. Dignum, "Responsible artificial intelligence: Designing AI for human values," *ITU Journal: ICT Discoveries*, no. SI No 1, pp. 1-8, 2017.
- [39] Z. Shen, H. Ma, and K. Wang, "A Web-scale system for scientific knowledge exploration," in *Proceedings of ACL 2018, System Demonstrations*, 2018, pp. 87-92.
- [40] Y. Zhang *et al.*, "Does deep learning help topic extraction? A kernel k-means clustering method with word embedding," *Journal of Informetrics*, vol. 12, no. 4, pp. 1099-1117, 2018.

- [41] Y. Zhang, M. Wu, W. Miao, L. Huang, and J. Lu, "Bi-layer network analytics: A methodology for characterizing emerging general-purpose technologies," *Journal of Informetrics*, vol. 15, no. 4, p. 101202, 2021.
- [42] C. Zhu and K. Motohashi, "Identifying the technology convergence using patent text information: A graph convolutional networks (GCN)-based approach," *Technological Forecasting and Social Change*, vol. 176, p. 121477, 2022.
- [43] M. Wu, Y. Zhang, G. Zhang, and J. Lu, "Exploring the genetic basis of diseases through a heterogeneous bibliometric network: A methodology and case study," *Technological Forecasting and Social Change*, vol. 164, p. 120513, 2021.
- [44] L. Huang, X. Chen, Y. Zhang, Y. Zhu, S. Li, and X. Ni, "Dynamic network analytics for recommending scientific collaborators," *Scientometrics*, vol. 126, no. 11, pp. 8789-8814, 2021.
- [45] L. Fleming, "Recombinant uncertainty in technological search," *Management Science*, vol. 47, no. 1, pp. 117-132, 2001.
- [46] Y. Zhang *et al.*, "Framework of Computational Intelligence-Enhanced Knowledge Base Construction: Methodology and A Case of Gene-Related Cardiovascular Disease," *International Journal of Computational Intelligence Systems*, vol. 13, no. 1, pp. 1109-1119, 2020.
- [47] Y. Zhang, H. Chen, J. Lu, and G. Zhang, "Detecting and predicting the topic change of Knowledge-based Systems: A topic-based bibliometric analysis from 1991 to 2016," *Knowledge-Based Systems*, vol. 133, pp. 255-268, 2017.
- [48] M. Bastian, S. Heymann, and M. Jacomy, "Gephi: An open source software for exploring and manipulating networks," *Proceedings of International AAAI Conference on Web and Social Media*, vol. 8, pp. 361-362, 2009.
- [49] M. Liu *et al.*, "Further divided gender gaps in research productivity and collaboration during the COVID-19 pandemic: Evidence from coronavirus-related literature," *Journal of Informetrics*, vol. 16, no. 2, p. 101295, 2022.
- [50] Y. Dong, J. Kang, H. Tong, and J. Li, "Individual fairness for graph neural networks: A ranking based approach," in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 300-310.
- [51] M. Krzywinski *et al.*, "Circos: An information aesthetic for comparative genomics," *Genome Research*, vol. 19, no. 9, pp. 1639-1645, 2009.
- [52] T. Kamada, "The issues of automated driving vehicle and the expectations for 3D integration technology," in *2015 International 3D Systems Integration Conference (3DIC)*, 2015: IEEE, pp. KN2. 1-KN2. 4.
- [53] M. Grootendorst, "KeyBERT: Minimal keyword extraction with BERT," *Zenodo*, 2020, doi: 10.5281/zenodo.4461265.
- [54] L. Waltman, N. J. van Eck, and E. C. Noyons, "A unified approach to mapping and clustering of bibliometric networks," *Journal of Informetrics*, vol. 4, no. 4, pp. 629-635, 2010.
- [55] D. Rotolo, D. Hicks, and B. R. Martin, "What is an emerging technology?," *Research Policy*, vol. 44, no. 10, pp. 1827-1843, 2015.
- [56] D. Cetindamar, K. Kitto, M. Wu, Y. Zhang, B. Abedin, and S. Knight, "Explicating AI literacy of employees at digital workplaces," *IEEE transactions on engineering management*, 2022.
- [57] Y. Huang, Y. Zhang, M. Wu, A. Porter, and R. Barrangou, "Determination of Factors Driving the Genome Editing Field in the CRISPR Era Using Bibliometrics," *The CRISPR Journal*, vol. 4, no. 5, pp. 728-738, 2021.



Yi Zhang (SM'23) is an Associate Professor at the Australian Artificial Intelligence Institute, University of Technology Sydney (UTS). He received a dual-PhD degree in Management Science and Engineering from Beijing Institute of Technology (2016) and Software Engineering from UTS (2017). His research aligns with bibliometrics and technology management. He has published more than 100 research articles in leading journals and conferences in related fields.

Dr Zhang was the receipt of the 2019 Australian Research Council's Discovery Early Career Researcher Award (DECRA). He serves as an Executive Editor of *Technol Forecast Soc*, a Specialty Chief Editor of *Front Res Metr Anal*,

and an Associate Editor of *IEEE T Eng Manage and Scientometrics*.



Mengjia Wu is a Postdoctoral Research Associate at the Australian Artificial Intelligence Institute, University of Technology Sydney (UTS). He holds a PhD degree in Information Systems from UTS (2023) and has published around 30 papers in refereed journals and conference proceedings. These publications cross multiple disciplines, aligning with his research interests in incorporating network analytics and AI methods, especially natural language processing and graph neural networks, to address information science issues. Dr Wu serves as the Associate Editor of *Front Res Metr Anal*.



Guangquan Zhang is an Associate Professor and Director of the Decision Systems and e-Service Intelligent Research Laboratory at the University of Technology Sydney, Australia. He received the Ph.D. degree in applied mathematics from Curtin University of Technology, Australia, in 2001. His research interests include fuzzy machine learning, fuzzy optimization, and machine learning. He has authored five monographs, five textbooks, and 460 papers including 220 refereed international journal papers.

Dr Zhang has won seven Australian Research Council (ARC) Discovery Projects grants and many other research grants. He was awarded an ARC QEII fellowship in 2005.

He has served as a member of the editorial boards of several international journals, as a guest editor of eight special issues for IEEE transactions and other international journals and co-chaired several international conferences and workshops in fuzzy decision-making and knowledge engineering.



Jie Lu (F'18) is an Australian Laureate Fellow, IFSA Fellow, ACS Fellow, Distinguished Professor, and the Director of Australian Artificial Intelligence Institute at the University of Technology Sydney, Australia. She received a PhD degree from Curtin University in 2000.

Her main research expertise is in transfer learning, concept drift, fuzzy systems, decision support systems and recommender systems. She has published over 500 papers in IEEE Transactions and other leading journals and conferences. She is the recipient of two IEEE Transactions on Fuzzy Systems Outstanding Paper Awards (2019 and 2022), NeurIPS2022 Outstanding Paper Award, Australia's Most Innovative Engineer Award (2019), Australasian Artificial Intelligence Distinguished Research Contribution Award (2022), Australian NSW Premier's Prize on Excellence in Engineering or Information & Communication Technology (2023), and the Officer of the Order of Australia (AO) 2023.