Artificial intelligence and sustainable development goals nexus via four vantage points

Osama Nasir [a], Rana Tallal Javed [a,b], Shivam Gupta [c], Ricardo Vinuesa [d], Junaid Qadir [e,*]

[a] Information Technology University of the Punjab, Lahore, Pakistan
[b] Department of Informatics, University of Oslo, Oslo, Norway
[c] Bonn Alliance for Sustainability Research (University of Bonn), Bonn, Germany
[d] FLOW, Engineering Mechanics, KTH Royal Institute of Technology & KTH Climate Action Centre, Stockholm, Sweden, Stockholm, Sweden
[e] Department of Computer Science and Engineering, College of Engineering, Qatar University, Doha, Qatar

ARTICLE INFO

ABSTRACT

Artificial Intelligence (AI) should aim at benefiting society, the economy, and the environment, i.e., AI should aim to be socially good. The UN-defined Sustainable Development Goals (SDGs) are the best depiction to measure social good. For AI to be socially good, it must support all 17 UN SDGs. Our work provides a unique insight into AI on all fronts including Curricula, Frameworks, Projects, and Research papers. We then analyze these datasets to extract meaningful information for policymakers and researchers alike - shedding light on how AI is being used and can potentially be employed in the future to achieve the SDGs. To this end, we devised a methodology using keyword-matching and keyword-similarity to compute the relevance of the SDGs for a given document. SDG metadata and AI4SDG Projects (Oxford initiative on AI4SDGs) were used to validate our methodology. We find an imbalance of coverage with SDG 9 (Industry Innovation and Infrastructure) having the highest representation (with 50.3% of our data containing references to it) compared to SDGs 5, 6, 14, and 15, which have the lowest representation (5% of observed data). Findings from this study suggest that the development of AI technology is focused on improving the current economic growth, but it might neglect important societal and environmental issues.

1. Introduction

In 2015, the United Nations (UN) proposed 17 Sustainable Development Goals [1] (SDGs) comprising 169 targets for attaining the 2030 UN agenda for sustainable development. These goals have a vast range from poverty to climate change. Harnessing the power of Artificial Intelligence (AI) to achieve social good and SDGs has attracted the fancy of numerous practitioners and researchers [2]. Researchers have, however, shown that AI technology is a double-edged sword with the potential of promoting as well as inhibiting SDGs. Vinuesa et al. [3] explored the enabling and inhibiting role of AI with regard to the 169 targets specified for the 17 SDGs and showed that AI acted as an enabler for 134 targets while serving as an inhibitor for 59 targets (for some targets, AI was both potentially an enabler and an inhibitor). Furthermore, Gupta et al. [4] extended their work to discussions on the implications of AI on the SDGs at the indicator level.

Recent developments in AI have also given rise to ethical concerns around it. There are substantial concerns regarding the bias in AI decision-making [5,6], insecurity of jobs [7], fairness [8,9] and accountability of AI [10,11]. To address these concerns around AI, the field of AI Ethics is maturing rapidly and researchers, organizations and governments alike have proposed various ethical frameworks and principles for AI to be socially beneficial [12–14]. Hagendorff [15] manually analyzed 22 frameworks and presented an overview of the topics being discussed within them. Similarly, Jobin et al. [16] gave an overview of content within 84 AI-Ethics frameworks. There are also rising concerns that these frameworks have not led to any substantial change [17]. It has been claimed that these frameworks are mostly policy statements and are non-binding for organizations that adopt or propose them. Thus corporations often create ethical frameworks for "ethics washing" without any real teeth or authority [18–20].

In parallel, guidelines and curricula for ethical AI are being drafted,

and undergraduates and engineers alike are being taught the norms and rules of designing ethical solutions and making ethically conscious decisions. While AI-Ethics frameworks provide a roadmap for building sustainable and ethical AI, AI curricula provide the foundation for encouraging future AI practitioners to think critically about social good. Efforts have also been made to analyze what is being taught in AI-Ethics curricula, with researchers using manual methods [21], as well as automated analysis techniques [22,23], to identify topics in AI-Ethics and Engineering-ethics curricula. Garret et al. [24] gave an overview on not only what is being taught, but also comments on what should be taught when teaching AI-Ethics. Cowls et al. [25] presented a benchmark dataset for AI Projects which aim to achieve the SDGs, with a brief analysis of their dataset to understand the efforts being done in the realization of the SDGs. Lee and Kim [26] analyzed social-media content to judge public opinion regarding the attainment of the SDGs.

The AI ethics research works often do not discuss the SDGs. Understanding the alignment of these ethical frameworks for AI with SDGs will enable authors of such frameworks to better understand the landscape of AI-Ethics. Keeping in mind the critical importance of the SDGs, and the effort devoted to them, it is vitally important to know how to ethically use AI and avoid any pitfalls related to AI. This will also enable researchers and policymakers alike to understand the direction in which the field of AI-Ethics is heading and steer it if the need arises [27,28].

1.1. Focus of this paper

In this paper, we analyze AI Ethics in a unique yet interlinked manner by leveraging four distinct vantage points: (1) AI-Ethics Curricula; (2) AI-Ethics Frameworks; (3) AI for SDG (AI4SDG) Projects; and (4) AI4SDGs research papers. After the assessment, we will provide a comprehensive view emerging from these four distinct vantage points. We employ an automated technique to find interesting overlaps between AI and SDGs based on the given datasets. Previous studies such as Vinuesa et al. [3] were limited to a narrower range of data due to the manual nature of the conducted analysis. Having an automated approach provides us with the opportunity to analyze more data than ever before.

1.2. Contributions of this paper

Our work builds on top of previous studies that have explored the use of AI for social good including systematic literature reviews that present various use cases and guidelines [29,30] as well as bibliometrics analysis [31]. Other works have also comprehensively focused in isolation on various related aspects such as AI-Ethics frameworks [16], AI curricula [21–23], accountable human-centered AI [32], and AI and SDGs [3]. This is the first work, however, to the best of our knowledge, that tries to observe the nexus of four distinct vantage points (AI Curricula, Frameworks, Projects, and Research papers) to SDGs in a holistic way. Overall, this study provides the following key contributions:

1. We devise a methodology to compute the relevance of an SDG with a given document.
2. We provide insights into AI-Ethics and its relevance for the SDGs.
3. We release the first dataset that contains data related to AI-Ethics Frameworks, Research Papers, Projects and Curricula, along with their relevance for each SDG.

2. Materials and methods

2.1. Data collection

To understand how AI-Ethics Curricula, Frameworks, Projects, and Research Papers are aligned with the SDGs, we collected documents for all these datasets and evaluated the relevance of each SDG within all the documents using keyword-matching and keyword-similarity scores. The

UN has provided SDG metadata with detailed descriptions of goal indicators and associated methods for their measurement.1 Similarly, the Oxford database of AI4SDG projects provides AI-project descriptions and associated SDGs.2 SDG labels for both these datasets are already available. The experimentally-assigned labels are compared with the true labels to obtain validation accuracy. We performed various experiments to find the optimal experimental setting for the SDG-relevance score. This optimal setting was used to compute relevant SDG labels for curricula, frameworks, projects, and research papers. For our convenience, we have used short SDG names during our analysis. The full SDG names along with their corresponding short names are listed in Table 2. We assess the connections with the SDGs from 4 datasets, namely: AI-Ethics Frameworks, AI-Ethics Curricula, AI4SDGs Projects, and AI4SDGs Research Papers. AI4SDGs Projects and SDG metadata are used for methodology validation. A summary of the number of documents in each of these datasets is given in Table 1 and their description is provided below:

(a) Geographical distribution of 108 AI-Ethics frameworks based on the country of the issuer.
(b) Geographical distribution of 166 AI-Ethics curricula based on the university offering the course.
(c) Geographical distribution of 108 AI4SDG projects from the Oxford initiative of AI4SDGs.
(d) Geographical distribution of authors' affiliated institute for 200 AI-research articles.

2.2. AI-Ethics Frameworks

For our AI-Ethics frameworks dataset, we initially used the comprehensive list of 84 guidelines provided by Jobin [16]. In addition, we also added 24 extra guidelines including China's New Generation AI development plan, the National AI strategy for Qatar, Czech, Singapore, the Organisation for Economic Co-operation and Development (OECD) AI Principles, G20 AI Principles, and others. The geographic distribution of these 108 frameworks based on the country of the issuer is shown in Fig. 1(a).

2.3. AI-ethics curricula

For our AI-curricula dataset, we used the publicly available Google sheet of Tech Ethics curricula provided by Fiesler [21], as our initial dataset. This sheet provides the primary meta-information for tech-ethics courses, which consists of the course title, course instructor, course level (undergraduate or graduate), the teaching department, university, and open links for some course descriptions and syllabi. This dataset is defined as 'tech ethics' and at the time of our analysis, it contained 259 courses. In addition, we have also added 43 AI-Ethics

Table 1
Summary of all 5 Datasets used in our analysis.

Dataset	# of Documents	Description
Frameworks	108	Publicly available frameworks for AI-Ethics and guidelines published by various organizations, governments, and institutes.
Curricula	166	Publicly available AI-ethics course curricula.
Projects	108	Project descriptions from the AI4SDG project of Oxford University, where they collected 108 projects that used AI for solving SDGs.
Research Papers	200	Publications related to AI and SDGs were manually searched from Scopus (since 2015).
Total	**582**	

Table 2
List of short SDGs employed in the present analysis.

SDG #	SDG Full Name	Short Name
SDG 1	No Poverty	Poverty
SDG 2	Zero Hunger	Hunger
SDG 3	Good Health and Well-being	Health
SDG 4	Quality Education	Education
SDG 5	Gender Equality	Gender Eq.
SDG 6	Clean Water and Sanitation	Water
SDG 7	Affordable and Clean Energy	Energy
SDG 8	Decent Work and Economic Growth	Economy
SDG 9	Industry Innovation and Infrastructure	Industry
SDG 10	Reduced Inequality	Inequality
SDG 11	Sustainable Cities and Communities	Cities
SDG 12	Responsible Consumption & Production	Resp. C&P
SDG 13	Climate Action	Climate
SDG 14	Life Below Water	Life Water
SDG 15	Life on Land	Life Land
SDG 16	Peace and Justice Strong Institutions	Peace
SDG 17	Partnerships to achieve the Goal	Partnership

courses apart from those provided in the crowd-sourced spreadsheet in an effort to make the dataset more global. Hence, a list of 302 courses was created. When downloading the curricula of these courses, it was found that many did not make their curricula publicly available, thus we removed those curricula which were not publicly available from our dataset. The curriculum for a course was considered to be valid if it contained descriptions regarding the topics which would be covered in the course. As a result, our final dataset contains a total of 166 curricula. The geographic distribution of these 166 curricula is shown in Fig. 1(b). For ease of use, we call refer to AI-Ethics curricula as curricula only in the rest of the paper.

2.4. AI4SDGs projects

In July 2018, the Oxford university started a survey to receive AI projects which are designed to achieve SDGs [33]. The survey lasted for more than 2 years and collected a total of 108 AI4SDG projects which satisfied their pre-defined criteria [25]. Most of the projects mainly focus on one or two SDGs, but some projects also cover more. This dataset is also used for the validation of our methodology. The geographic distribution of these 108 projects is shown in Fig. 1(c).

2.5. AI4SDGs research papers

Apart from the AI-Ethics Frameworks and AI-Ethics curricula, we have also collected AI4SDGs Research Papers. The dataset was collected from Scopus. The keywords used to extract data from Scopus were 'Artificial Intelligence', 'AI', 'Machine Learning', 'ML', 'SDGs', 'Sustainability', 'Sustainable Development', 'Sustainable Development Goals', and 'Social Welfare'. A total of 200 papers were collected as a result. We only use the abstract of these research papers to find the relevant SDGs. The geographic map of the authors' affiliated institute is shown in Fig. 1(d). When authors were associated with institutes from different countries, the research paper was considered for all the included countries.

2.6. SDG metadata

In 2016, the UN released the metadata information for all the goals, which includes a definition and comprehensive description of each associated indicator [34]. This dataset contains 17 documents, one for each goal. These detailed documents explain the interpretation and importance of indicators, their data-collection methodology, data


(a) Geographical distribution of 108 AI-Ethics frameworks based on the country of the issuer.

(b) Geographical distribution of 166 AI-Ethics curricula based on the university offering the course.

(c) Geographical distribution of 108 AI4SDG projects from the Oxford initiative of AI4SDGs.

(d) Geographical distribution of authors' affiliated institute for 200 AI-research articles.

Fig. 1. Geographical distribution of AI-Ethics Curricula, AI4SDG Projects, and AI4SDG Research Papers (based on all authors' affiliated institutes). We can observe that the United States has the highest number of Curricula, Frameworks, and Projects, while China has the highest number of AI-Research Papers aimed at achieving the SDGs.

sources, computational methods, and limitations, along with references. This dataset is primarily used for our methodology validation. A detailed description of the approach to determine the connection of a document with a particular SDG is provided in the Methodology section.

2.7. Methodology

To find the contribution to the SDGs in a document, the representation of each SDG in a quantitative form is critical. To this end, we have used the compiled list of SDG keywords provided by Monash University and SDSN Australia/Pacific.3 This is an extensive list of keywords that contains multiple keywords and phrases to represent each SDG. In this way, each SDG is represented as a list of its relevant keywords. The number of keywords ranges from 27 to 67.

The text descriptions and SDG labels are publicly available for SDG metadata and AI4SDGs Projects. We extracted text for both these datasets and saved it in an Excel spreadsheet. For text pre-processing, we removed stopwords from the text. Then, the text was lemmatized using spacy [35] and only nouns, verbs, adjectives, and adverbs were extracted to generate the bag of words. Similarly, we also generated a bag of words for all SDGs using their list of keywords. Keyword-matching and keyword-similarity methods were used to calculate relevance scores for SDGs. We performed various experiments (provided in Supplementary Tables 1 and 2) to find the optimal setting yielding the most relevant SDG labels for our validation datasets. The details of each of these methodologies are given below.

2.8. SDG score using keyword matching

To compute the score of each SDG in a given document using keyword matching, we count the number of occurrences of each keyword for a given list of keywords and normalize this count with the document length. Exact keyword matching gives us the exact percentage of text relevant to an SDG. The formula to compute the keyword-matching score, denoted by M_s, is given by:

$$M_s = \sum_i^n \sum_j^m \text{match}(w_i, k_j). \tag{1}$$

Here n and m are the total number of words in a document and keyword list, respectively. We check whether each keyword is present in the document and its count is returned. The score is then normalized using the document length to remove bias against wordy documents. Since we are using an extensive list of keywords, this technique captures the exact text corresponding to each SDG.

2.9. SDG score using keyword similarity

Keyword-similarity score computes the similarity between words that are not exactly the same but are semantically similar. This method exhibits a great advantage when comparing the meaning of similar words. The keyword-similarity score is computed using cosine similarity between document words and keyword lists. Cosine similarity computes the cosine of the angle between the vectorized words, represented as 100-dimensional vectors. For semantically similar words, these 100-dimensional vectors point in the same direction, therefore their cosine similarity is also closer to 1. If the similarity score for a word-keyword pair is less than the threshold (0.75) then the score is discarded. This similarity threshold was found using experimentation. The keyword-similarity score S_s for a document with n words and an SDG with m keywords is expressed as follows:

$$S_s = \sum_i^n \sum_j^m \text{cosine} \quad - \quad \text{similarity}(w_i, k_j). \tag{2}$$

Both word w_i and keyword k_j are represented using pre-trained GloVe word embedding to map each word to a 100-dimensional numerical vector. Then, the similarity scores are aggregated for all word-keyword pairs. The aggregated similarity scores are normalized with the length of the SDG keyword list.

2.10. Normalization

The scores calculated for each document are min-max normalized to facilitate relative comparison. The formula for min-max normalization is given as:

$$s_i = \frac{x_i - \min(x)}{\max(x) - \min(x)} \times 100. \tag{3}$$

Here s_i is the score with normalization and x_i is the score before normalization. Vector x contains the absolute score for each SDG for a given document. In min-max normalization, the lowest score is mapped to 0 and the highest score is mapped to 100. The scores in-between are linearly scaled between 0 an 100.

2.11. Validating our methodology

We have used 2 datasets (SDG metadata, AI4SDG projects) to identify the optimal experimental setting for both methodologies (M_s, S_s) which assigns the most relevant SDG labels. To compute validation accuracy for a document, the scores for each SDG are computed and normalized (using min-max normalization). Two thresholding techniques were used for testing. In the first technique, only the top-n SDG labels with the highest scores are assigned. In the second technique, all SDG scores higher than the threshold are assigned. Different score thresholds will generate different SDG labels for a given document. These assigned SDG labels are then compared with the true labels. If experimental labels contain the true label, it is considered to be a correct assignment.

Consider the example: Document-x has SDG-3 as its true label. The computed scores for the first 5 SDGs are 0.1, 0.2, 0.9, 0.6 and 0.4 respectively. If we use top-3 SDGs, the SDG assignment would be {SDG-3, SDG-4, SDG-5}. It would be considered correct, but it contains 2 extra labels (which is less precise). If the score threshold is set to 0.7, the assignment would become {SDG-3}. This would be considered the correct and most suitable assignment.

Using these threshold techniques, SDG labels for both validation datasets (AI4SDG projects and SDG metadata) are assigned and evaluated against their actual labels to compute the validation accuracy. The formula for validation accuracy is given below:

$$A = \frac{\text{total} - \text{correct} - \text{assignments}}{\text{total} - \text{documents}} \times 100. \tag{4}$$

Validation results for SDG metadata and SDG projects using the keyword-matching and keyword-similarity techniques for different thresholds are provided in Supplementary Tables S1 and S2. For SDG metadata documents, 100%-accurate SDG labels were observed using both keyword-matching and similarity scores with SDG keyword-list length normalization. For AI4SDG projects, high validation accuracy is observed for the keyword-similarity score with normalization using the length of the SDG keyword-list, and keyword-matching score with document length normalization.

3. Results and discussion

(a) Percentage of Curricula which are labeled with the relevant SDGs.
(b) Percentage of Frameworks that are labeled with the relevant SDGs.
(c) Percentage of Projects which are labeled with the relevant SDGs.
(d) Percentage of Research Papers which are labeled with the relevant SDGs.

We found the relevance of the SDGs with all four datasets using

keyword-matching with document-length normalization (M_s/L_{doc}) and keyword-similarity with SDG keyword-list length normalization (S_s/L_{SDG}). These scores are scaled between 0 and 1 for a given document using min-max normalization. This helps in the comparative analysis of these scores. In the end, a threshold of 0.7 is applied to get the most relevant SDG labels for each document. Fig. 2 displays the percentage of documents for each of the 4 datasets which are assigned the relevant SDG labels using the average of the two methodologies (M_s and S_s). Separate barplots for each methodology (M_s and S_s) are provided in Supplementary Figs. S1 and S2.

SDG 9 (Industry Innovation and Infrastructure) has the highest representation in curricula, frameworks, and research papers. Out of total 582 documents (Frameworks, Curricula, Projects, and Research Papers - Table 1) 50.3% documents (more than half of the documents in all 4 datasets combined) contribute towards SDG 9, 25.4% contribute towards SDG 1 (No Poverty), 24.14% aim at achieving SDG 3 (Good Health and Well-being) and 23.7% documents focus on SDG 11 (Sustainable Cities). Conversely, SDG 5 (Gender Equality), SDG 6 (Water), SDG-14 (Life Water), and SDG-15 (Life Land) have the least representation, with less than 5% documents from all 4 datasets which address them. This shows that the development of AI technology is focused on improving the current economic growth while ignoring important societal and environmental issues.

Vinuesa et al. [3] divided the SDGs into 3 groups: Society (SDG 1 to 7, 11, 16), Economy (SDG 8, 9, 10, 12, 17), and Environment (SDG 13, 14, 15). We have used this categorization to understand which ones of the SDGs are more prominent and which ones are being overseen. In Fig. 2, Society SDGs have 45% representation, Economy SDGs have 47.2% representation, and Environment SDGs have only 7.8% representation.

3.1. AI-ethics curricula

Fig. 2(a) shows the percentage of AI-Ethics curricula that are labeled with respective SDGs using the average of keyword-matching and keyword-similarity techniques. We can observe that SDG 9 (Industry) and SDG 1 (No Poverty) have the highest contribution in curricula with 55% and 46% labeled curricula. SDG 10 (Reduced Inequality) and SDG 8 (Decent Work and Economic Growth) also have a significant representation with 15% labeled curricula. SDG 4 (Quality Education) shows significant relevance to curricula, although manual inspection showed that this relevance is only due to the curricula dataset containing many keywords related to education. SDG 6 (Sanitation) and SDG 15 (Life Land) have no relevance to any AI-Ethics curriculum. SDG 2 (Hunger), SDG 7 (Clean Energy), SDG 12 (Responsible Consumption & Production), SDG 13 (Climate), and SDG 14 (Life Water) have negligible representation in curricula, each having less than 3 labeled documents out of 166 curricula.

3.2. AI-Ethics Frameworks

Understanding the alignment of AI-Ethics frameworks with the UN SDGs will enable a better understanding of the AI-Ethics landscape enabling researchers and policymakers alike to understand the direction
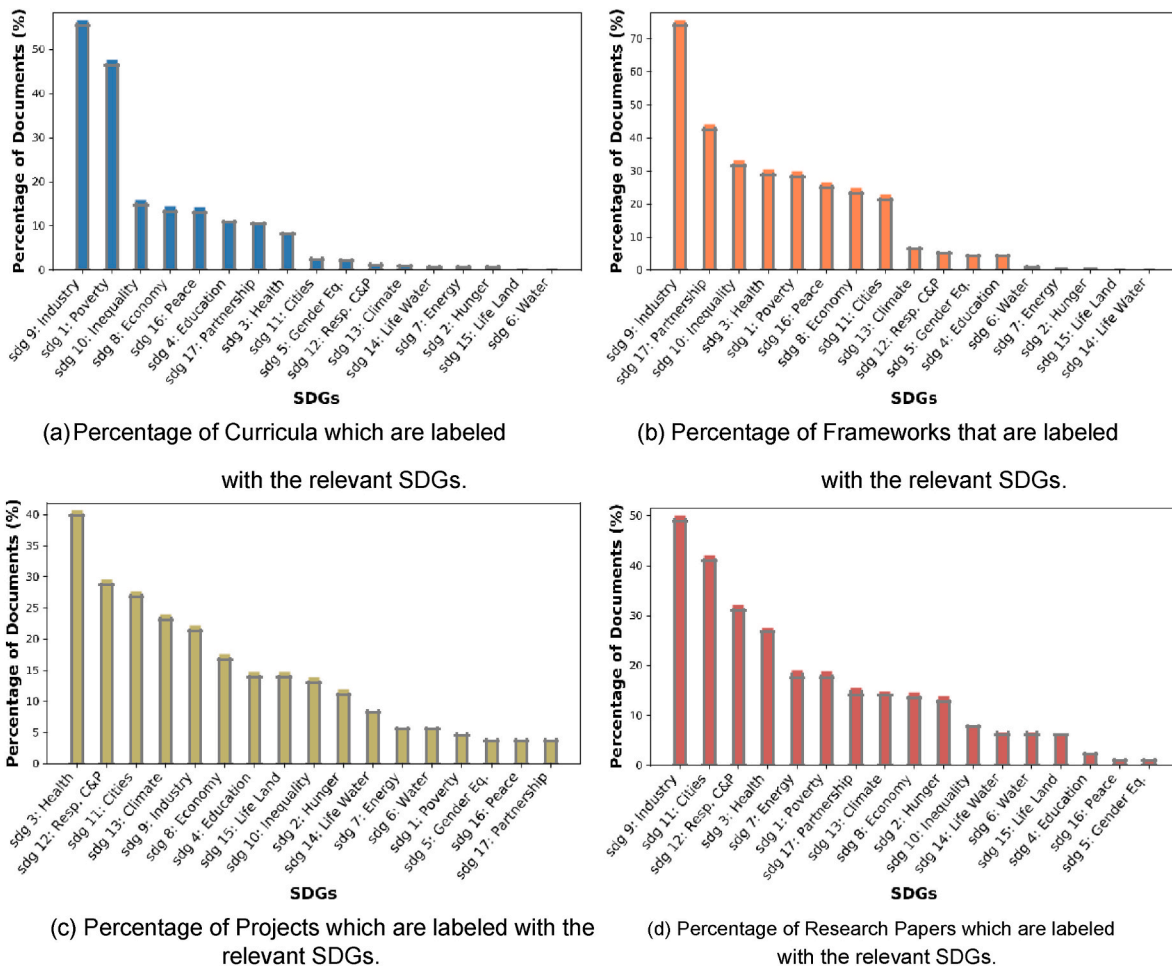

(a) Percentage of Curricula which are labeled with the relevant SDGs.

(b) Percentage of Frameworks that are labeled with the relevant SDGs.

(c) Percentage of Projects which are labeled with the relevant SDGs.

(d) Percentage of Research Papers which are labeled with the relevant SDGs.

Fig. 2. SDG contribution in decreasing order for all 4 datasets. SDG labels are assigned using the average of normalized keyword-matching (M_s/L_{doc}) and keyword-similarity (S_s/L_{SDG}) scores for each dataset.

in which the field of AI-Ethics is heading and steer it if the need arises. Fig. 2(b) shows the percentage of AI-Ethics frameworks that are labeled with SDGs using the average of keyword-matching and keyword-similarity techniques. We can observe that SDG-9 (Industry) has the highest contribution in all the framework documents with 74% of frameworks addressing it. Some notable SDGs, which individually show relevance to more than 25% of frameworks, are SDG 9 (Industry), SDG 17 (Partnership), SDG 10 (Reduced Inequality), SDG 3 (Health), SDG 1 (No Poverty) and SDG 16 (Peace). On the other hand, SDG 14 (Life Water) and SDG 15 (Life Land) have no relevance to any AI-Ethics framework, which shows major negligence toward the Environmental group. SDG 2 (Hunger), SDG 6 (Sanitation), and SDG 7 (Energy), which belong to the Society group of the SDGs, have negligible representation in AI-Ethics frameworks, with only two relevant frameworks addressing them. Frameworks show the most inclination toward the Economy group. The distribution of the SDGs based on their classes shows that Economy SDGs have 60% representation in frameworks, Society SDGs have 38% representation, while Environment SDGs have only 2% representation in all frameworks.

3.3. AI4SDG projects

Fig. 2(c) shows the percentage of AI4SDG projects, from the Oxford initiative on AI4SDG, that are labeled with SDGs based on the average of keyword-matching and keyword-similarity techniques. SDG 3 (Health) has the highest number of relevant projects using both techniques. Some notable SDGs in projects are SDG 3 (Health), SDG 12 (Responsible Consumption & Production), SDG 11 (Sustainable Cities), and SDG 13 (Climate), suggesting well-balanced attention toward Society, Economy as well as Environment. Most of the assigned SDG labels for projects match the actual labels because this dataset is used in the validation process for both methodologies. AI4SDG projects seem to show some negligence towards SDG 1 (No Poverty), SDG 5 (Gender Equality), SDG 16 (Peace), and SDG 17 (Partnership), as no more than 5 related projects are addressing them. SDG 9 *Industry, Innovation, and Infrastructure* has the most significant contribution to all domains of AI Ethics, with the majority of the documents showing relevance to it. On the other hand, all four datasets seem to neglect societal SDGs SDG 2 *Zero Hunger*, SDG 5 *Gender Equality*, and environmental SDGs SDG 14 *Life Below Water*, and SDG 15 *Life on Land*. A comparison of SDG relevance to all datasets shows that frameworks, which have the responsibility of guiding the use of AI technology, show the most inclination towards economy and negligence towards some society SDGs and all environment SDGs. It is time for the AI Ethics policymakers, to start taking such initiatives which equally address all SDGs and promote social development by developing a harmonious set of policies to encourage other perspectives of AI Ethics toward sustainability.

3.4. AI4SDG research papers

Fig. 2(d) shows the percentage of AI4SDG research papers that are labeled with SDGs based on the average of keyword-matching and keyword-similarity techniques. Similar to curricula and frameworks, SDG with the highest number of relevant research papers is SDG 9 (Industry). SDGs which have relevance to more than 25% research papers are SDG 9 (Industry), SDG 11 (Sustainable Cities), SDG 12 (Responsible Consumption & Production), and SDG 3 (Health). SDG 4 (Quality Education), SDG 5 (Gender Equality), and SDG 16 (Peace and Justice) are the least discussed in research papers and have negligible contributions with less than 5 related research papers. Unlike curricula and frameworks, research papers do not completely oversee SDGs related to the environment (SDG 13 Climate, SDG 14 Life Water, SDG 15 Life Land) and more than 10% research papers are addressing them.

In general, AI Frameworks, Curricula, and Research Papers are more concerned with SDGs relating to Industry and Economic growth. AI Projects are more related to Health, but significant contributions were

found for other SDGs too. SDG 5 (Gender Equality) and SDG 6 (Clean Water and Sanitation) are the most neglected SDGs in all the datasets.

3.5. Recommendations

Based on the findings of the study, we summarize below our recommendations for AI practitioners, researchers, and policymakers:

1. SDGs 13, 14, and 15 (which constitute the Environment group) are not sufficiently addressed in all the AI areas under study here. An enhanced focus on these SDGs is essential.
2. AI-Ethics frameworks must focus more on Society (in particular, SDGs 2, 6, and 7) and the Environment (focusing on SDGs 14 and 15). At the moment, the frameworks are excessively focused on economic development.
3. AI research studies need to increase the focus on SDGs 4 (quality education), 5 (gender equality), and 16 (strong institutions). These SDGs are essential for the social context of sustainable development.
4. AI curricula need to be steered towards finding solutions for Society (concretely in the context of the SDGs 2, 6, and 7) and Environmental problems (where all the SDGs, 13, 14, and 15, need more representation).
5. Although AI-based projects considered in this study are more balanced, more attention is required for SDG 1 (no poverty), SDG 5 (on gender equality), SDG 6 (on clean water), and SDG 7 (on clean energy).
6. The priorities are fragmented, and inclusive action is required to balance the efforts toward achieving SDGs via AI. It is essential for new thoughts to emerge and that new AI technologies be leveraged for the successful attainment of SDGs.

4. Conclusions

In order to better understand the linkage between the state of artificial intelligence (AI) and the United Nations (UN) sustainable development goals (SDGs), we use similarity-matching techniques to unveil new trends and analyze the relationship between AI-related document data and SDGs in an automated manner. The field of AI is an ongoing technical revolution that is profoundly impacting every sphere of life. Therefore, it is crucial to ensure that the power of AI is harnessed for our economical, environmental and social welfare while minimizing the negative ethical implications. The SDGs focus equally on the three pillars of Social Good, i.e. Economy, Society, and Environment. Our analysis shows that all 4 datasets are more concerned with SDGs regarding the Economy, with limited attention toward Society and almost neglecting the Environment. Efforts should be made to promote the achievement of the SDGs within the Environment and Society areas so that all SDGs can be achieved uniformly for our social wellness. It must be noted that AI is not a panacea, but the exploration of AI for all SDGs should not be neglected. Overall, our study finds that more work needs to be carried out for reorienting and linking the most important implications of AI, AI-Ethics, and SDGs. The work needs to be inclusive, harmonious, and cross-disciplinary so that the resulting technological innovations promote human progress while being respectful of the environment and the overall human society at large.

Data availability

Data will be made available on request.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.techsoc.2022.102171.

References

[1] #Envision 2030: 17 Goals to Transform the World for Persons with Disabilities;. https://www.un.org/development/desa/disabilities/envision2030.html.

[2] J. Cowls, 'AI for social good': whose good and who's good? Introduction to the special issue on artificial intelligence for social good, Philos. Technol. 34 (1) (2021) 1–5.

[3] R. Vinuesa, H. Azizpour, I. Leite, M. Balaam, V. Dignum, S. Domisch, et al., The role of artificial intelligence in achieving the Sustainable Development Goals, Nat. Commun. 11 (1) (2020) 1–10.

[4] S. Gupta, S.D. Langhans, S. Domisch, F. Fuso-Nerini, A. Felländer, M. Battaglini, et al., Assessing whether artificial intelligence is an enabler or an inhibitor of sustainability at indicator level, Transport Eng. 4 (2021), 100064.

[5] J. Zou, L. Schiebinger, AI Can Be Sexist and Racist—It's Time to Make it Fair, Nature Publishing Group, 2018.

[6] T.T. Krupiy, A vulnerability analysis: theorising the impact of artificial intelligence decision-making processes on individuals, society and human diversity from a social justice perspective, Comput. Law Secur. Rep. 38 (2020).

[7] T. Nam, Technology usage, expected job sustainability, and perceived job insecurity, Technol. Forecast. Soc. Change 138 (2019) 155–165.

[8] A. Yapo, J. Weiss, Ethical implications of bias in machine learning, in: Proceedings of the 51st Hawaii International Conference on System Sciences, 2018.

[9] D. Roselli, J. Matthews, N. Talagala, Managing bias in ai, in: Companion Proceedings of the 2019 World Wide Web Conference, 2019, pp. 539–544.

[10] E. Azzali, Accountability in AI as global issue, Management 20 (2020) 22.

[11] R. Vinuesa, B. Sirmacek, Interpretable deep-learning models to help achieve the Sustainable Development Goals, Nat. Mach. Intell. 3 (11) (2021) 926, 926.

[12] Floridi L, Cowls J. A Unified Framework of Five Principles for AI in Society. Available at: SSRN 3831321. 2019;.

[13] L. Floridi, J. Cowls, M. Beltrametti, R. Chatila, P. Chazerand, V. Dignum, et al., AI4People—an ethical framework for a good AI society: opportunities, risks, principles, and recommendations, Minds Mach. 28 (4) (2018) 689–707.

[14] J. Cowls, T. King, M. Taddeo, L. Floridi, Designing AI for Social Good: Seven Essential Factors, 2019. Available at: SSRN 3388669.

[15] T. Hagendorff, The ethics of AI ethics: an evaluation of guidelines, Minds Mach. (2020) 1–22.

[16] A. Jobin, M. Ienca, E. Vayena, The global landscape of AI ethics guidelines, Nat. Mach. Intell. 1 (9) (2019) 389–399.

[17] K. Hao, let's Stop AI Ethics-Washing and Actually Do Something, 2019, 2020.

[18] B. Wagner, Ethics as an escape from regulation. From "ethics-washing" to ethics-shopping?, in: Being Profiled Amsterdam University Press, 2018, pp. 84–89.

[19] E. Bietti, From ethics washing to ethics bashing: a view on tech ethics from within moral philosophy, in: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, 2020, pp. 210–219.

[20] K. Yeung, A. Howes, G. Pogrebna, AI Governance by Human Rights-Centered Design, Deliberation and Oversight: an End to Ethics Washing. The Oxford Handbook of AI Ethics, Oxford University Press, 2019, 2019.

[21] C. Fiesler, N. Garrett, N. Beard, What do we teach when we teach tech ethics? A syllabi analysis, in: Proceedings of the 51st ACM Technical Symposium on Computer Science Education, 2020, pp. 289–295.

[22] O. Nasir, S. Muntaha, R.T. Javed, J. Qadir, Work in progress: pedagogy of engineering ethics: a bibliometric and curricular analysis, in: IEEE Global Engineering Education Conference (EDUCON), 2021, pp. 1553–1557, 2021.

[23] R.T. Javed, O. Nasir, M. Borit, L. Vanhée, E. Zea, S. Gupta, et al., Get out of the BAG! Silos in AI ethics education: unsupervised topic modeling analysis of global AI curricula, J. Artif. Intell. Res. 73 (2022) 933–965.

[24] N. Garrett, N. Beard, C. Fiesler, More than "if time allows" the role of ethics in AI education, in: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, 2020, pp. 272–278.

[25] J. Cowls, A. Tsamados, M. Taddeo, L. Floridi, A definition, benchmark and database of AI for social good initiatives, Nat. Mach. Intell. 3 (2) (2021) 111–115.

[26] R. Lee, J. Kim, Developing a Social Index for Measuring the Public Opinion Regarding the Attainment of Sustainable Development Goals. Social Indicators Research, 2021, pp. 1–21.

[27] H.H. Goh, R. Vinuesa, Regulating artificial-intelligence applications to achieve the sustainable development goals, Discov. Sustain. 2 (1) (2021) 1–6.

[28] C. Wilson, M. Van Der Velden, Sustainable AI: an integrated model to guide public sector decision-making, Technol. Soc. 68 (2022), 101926.

[29] N. Tomašev, J. Cornebise, F. Hutter, S. Mohamed, A. Picciariello, B. Connelly, et al., AI for social good: unlocking the opportunity for positive impact, Nat. Commun. 11 (1) (2020) 1–6.

[30] A. Di Vaio, R. Palladino, R. Hassan, O. Escobar, Artificial intelligence and business models in the sustainable development goals perspective: a systematic literature review, J. Bus. Res. 121 (2020) 283–314.

[31] S.F. Wamba, R.E. Bawack, C. Guthrie, M.M. Queiroz, K.D.A. Carillo, Are we preparing for a good AI society? A bibliometric review and research agenda, Technol. Forecast. Soc. Change 164 (2021), 120482.

[32] J. Qadir, M.Q. Islam, A. Al-Fuqaha, Toward accountable human-centered AI: rationale and promising directions, J. Inf. Commun. Ethics Soc. 20 (No. 2) (2022) 329–342, https://doi.org/10.1108/JICES-06-2021-0059.

[33] Project finder, AISDG;. https://www.aiforsdgs.org/all-projects.

[34] Compilation of Metadata for the Proposed Global Indicators for the Review of the 2030 Agenda for Sustainable Development;. https://unstats.un.org/sdgs/iaeg-sdgs/metadata-compilation/.

[35] Lemmatizer, spaCy API Documentation;. https://spacy.io/api/lemmatizer.