

Technical Note

The Sustainable Development Goals and Aerospace Engineering: A critical note through Artificial Intelligence

Alejandro Sánchez-Roncero ^a, Òscar Garibo-i-Orts ^{a,f}, J. Alberto Conejero ^a, Hamidreza Eivazi ^b, Fermín Mallor ^b, Emelie Rosenberg ^b, Francesco Fuso-Nerini ^{c,d}, Javier García-Martínez ^e, Ricardo Vinuesa ^{b,d}, Sergio Hoyas ^{a,*}

^a Instituto de Matemática Pura y Aplicada, Universitat Politècnica de València, Camino de Vera 46024, València, Spain

^b ACES – Association of Spanish Scientists in Sweden, Stockholm, Sweden

^c Division of Energy Systems, KTH Royal Institute of Technology, Stockholm, Sweden

^d KTH Climate Action Centre, Stockholm, Sweden

^e Molecular Nanotechnology Lab, Department of Inorganic Chemistry, University of Alicante, Alicante, Spain

^f GRID - Grupo de Investigación en Ciencia de Datos, Valencian International University - VIU, València, Spain

ARTICLE INFO

Keywords:

Sustainability
United Nations
Sustainable Development Goals
Artificial Intelligence
Aerospace Engineering






ABSTRACT

The 2030 Agenda of the United Nations (UN) revolves around the Sustainable Development Goals (SDGs). A critical step towards that objective is identifying whether scientific production aligns with the SDGs' achievement. To assess this, funders and research managers need to manually estimate the impact of their funding agenda on the SDGs, focusing on accuracy, scalability, and objectiveness. With this objective in mind, in this work, we develop ASDG, an easy-to-use Artificial-Intelligence-based model for automatically identifying the potential impact of scientific papers on the UN SDGs. As a demonstrator of ASDG, we analyze the alignment of recent aerospace publications with the SDGs. The Aerospace data set analyzed in this paper consists of approximately 820,000 papers published in English from 2011 to 2020 and indexed in the Scopus database. The most-contributed SDGs are 7 (on clean energy), 9 (on industry), 11 (on sustainable cities), and 13 (on climate action). The establishment of the SDGs by the UN in the middle of the 2010 decade did not significantly affect the data. However, we find clear discrepancies among countries, likely indicative of different priorities. Also, different trends can be seen in the most and least cited papers, with apparent differences in some SDGs. Finally, the number of abstracts the code cannot identify decreases with time, possibly showing the scientific community's awareness of SDG.

1. Introduction


In 2015 all state members of the United Nations (UN) adopted the 2030 Agenda for Sustainable Development. The UN intends to promote peace and prosperity for people and the planet with a vision for the near future. To make that vision a reality, the 2030 Agenda consists of 17 Sustainable Development Goals (SDGs) [1]. They represent the actions that countries from all over the world (both developed and developing) should implement as global cooperation for the future of our planet.

The 17 SDGs, see description in [1], are as follows (those most closely related to Aerospace Engineering have been written in italic font):

-  SDG 1: End poverty in all its forms everywhere.
-  SDG 2: End hunger, achieve food security and improved nutrition and promote sustainable agriculture.
-  SDG 3: Ensure healthy lives and promote well-being for all at all ages.
-  SDG 4: Ensure inclusive and equitable quality education and promote lifelong learning opportunities for all.
-  SDG 5: Achieve gender equality and empower all women and girls.

* Corresponding author.

E-mail address: serhocal@mot.upv.es (S. Hoyas).

-  SDG 6: Ensure availability and sustainable management of water and sanitation for all.
-  SDG 7: *Ensure access to affordable, reliable, sustainable, and modern energy for all.*
-  SDG 8: Promote sustained, inclusive, and sustainable economic growth, full and productive employment, and decent work for all.
-  SDG 9: *Build resilient infrastructure, promote inclusive and sustainable industrialization, and foster innovation.*
-  SDG 10: Reduce inequality within and among countries.
-  SDG 11: *Make cities and human settlements inclusive, safe, resilient and sustainable.*
-  SDG 12: *Ensure sustainable consumption and production patterns.*
-  SDG 13: *Take urgent action to combat climate change and its impacts.*
-  SDG 14: Conserve and sustainably use the oceans, seas, and marine resources for sustainable development.
-  SDG 15: Protect, restore and promote sustainable use of terrestrial ecosystems, sustainably manage forests, combat desertification, and halt and reverse land degradation and halt biodiversity loss.
-  SDG 16: Promote peaceful and inclusive societies for sustainable development, provide access to justice for all and build effective, accountable and inclusive institutions at all levels.
-  SDG 17: Strengthen the means of implementation and revitalize the Global Partnership for Sustainable Development.

The two main questions this article wants to contribute are: is the Aerospace Engineering scientific community focused on fulfilling the SDGs? What are the most relevant SDGs in this community? To the best of our knowledge, there is no published work about this relationship. In this work, our answer is given using Artificial Intelligence (AI) tools.

It is important to note that the recent paradigm change introduced by the fast digitalization of business, academics, daily life, and even policy-making is profound. A recent study by Vinuesa et al. [2] in-depth examined how AI affects the accomplishment of the UN's 2030 Agenda. Although they discovered that 79% of the aims would be positively affected by AI, they also noted that the growth of AI could hinder or even have a detrimental impact on the achievement of 35% of these targets. The SDGs are all interconnected, and while there are numerous synergies, it is vital to recognize and properly document any trade-offs to reach the full potential of AI's ability to contribute to creating a sustainable future. Furthermore, Gupta et al. [3] extended their work to discussions on the implications of AI on the SDGs at the indicator level. In this regard, it is crucial to emphasize that implementing clear and understandable strategies requires employing AI-based technologies to achieve the SDGs. According to Vinuesa and Sirmacek [4], deploying interpretable AI would produce an algorithmic usage that focuses on accountability and transparency.

With this in mind, a preliminary version of our code ASDG (Automatic Classification of Impact to Sustainable Development Goals) can be found in [5]. We believe that a promising way to achieve significant progress in the SDG Agenda is by using AI-based methods to inform policy decisions to maximize the synergies and minimize the trade-offs. With this goal in mind, we created ASDG. This AI-based framework constitutes a step in this direction by enabling the automatic classification of hundreds of thousands of scientific papers by their impact on each SDG.

A critical example of the importance of understanding the situation of a certain funder about the SDGs is climate change. The SDGs must be accomplished while we are amid a climate emergency, as confirmed in the last Intergovernmental Panel on Climate Change [6]. This is particularly important in the case of Aerospace Engineering [7,8]. To summarize the importance of aerodynamics, for example, about a quarter of today's energy is spent moving fluids along pipes or vehicles through air or water. Turbulence dissipates 25% of this energy, which is responsible for up to 5% of the CO₂ dumped by humanity every year [9]. Considering that 340 billion liters of fuel were used in 2017 for air transportation worldwide (as reported by IATA [10]), there is considerable potential for energy savings and fuel consumption reduction. Before the coronavirus Disease-19 (COVID-19) pandemic, this quantity grew yearly at an unsustainable 3% rate.

To summarize, ASDG can be seen as a contribution to SDG17. However, we will not consider this SDG as it is the most difficult to identify. Certainly, few papers coming from the aerospace world are devoted to this issue. Excepting this SDG, ASDG is ready to analyze any field of science and technology. This work presents the first application of ASDG: Aerospace Engineering. A summary of ASDG is given in the next section, together with a description of the database. The results are explained in the third section. Conclusions and future work are described in the last section.

2. Methods

The code employed for this article, ASDG, can identify the connection between a paper and an SDG through its abstract. It uses four different models: Non-Negative Matrix Factorization (NMF) [11], Distributed Representations of Topics (Top2Vec) [12], Latent Dirichlet Allocation (LDA) [13], and BERTopic [14]. Due to their inherently different nature, the information that each model extracts from a text is different. In other words, their functionalities are complementary. To take advantage of this fact, ASDG introduces a voting mechanism. Similar ideas have been used very recently for studying the social network Twitter [15]. In the voting stage, ASDG takes the scores of each model for each text as inputs. Using this information, ASDG decides which identified SDGs have enough confidence to assume that the text relates to them.

The validation of ASDG was carried out in a previous publication [5]. The model's training (based on 510 manually-curated text files related to each SDG) was described in that work. Briefly, after downloading all papers referenced in [2], for a total of 186 works, we manually selected papers with at least an Abstract and Body differentiated, extracting the sections in 40% of them. A Deep Neural Network [16] was used to extract the remaining 60% automatically. This tool is based on images instead of converting the pdf file to text. We validated this tool with the extracted pdf files and checked out every abstract. As the authors of [2] classified all these papers based on an expert consensus, we labeled these papers to classify all these papers correctly, obtaining an 81% agreement.

The methods mentioned above are briefly described next.

2.1. NMF

Non-negative Matrix Factorization model (NMF) [11]. This method can reduce the space dimension of the problem, extracting essential features. We consider 16 topics, as SDG 17 is currently not considered. All training and validation texts have been preprocessed. This includes:

- Words lemmatization + stop words
- Removing numeric and non-ASCII characters.
- Words frequency and documents frequency were set to 1. This configuration means that no words are excluded.
- Bigrams were allowed.

All training texts are automatically identified with the appropriate SDG, using this information to associate each topic with one SDG. The score corresponding to each topic for each text file is queried after the model has been trained. The named SDGs are multiplied by that score, then recorded in a subject association map (nTopics x 17). The values for each topic are normalized (values/sum (values)), and those topics with scores of less than 0.1 are discarded. The final result is a matrix, where each row represents the likelihood that each SDG will be associated with a particular and single topic.

2.2. Top2Vec

A Top2Vec model [12] was trained using the embedding model “all-MiniLM-L6-v2”. This embedding was pre-trained on a larger corpus, which works better when the training corpus is small. A light preprocessing is required to remove non-ASCII characters. In this case, no document segmentation is defined. The extraction of topics was unsupervised. Since the association of the training texts with the SDGs was known beforehand, we queried the associated texts and their scores for each topic, creating an association matrix as it was done with the NMF model.

2.3. LDA

A latent-Dirichlet-allocation model [13] was also trained with the following configuration:

- Number of topics: 16.
- Passes: 400. Iterations: 1000. Chunk size: 2000
- Bigrams are allowed
- Minimum word count: 1, Maximum word frequency: 0.7

The training and validation texts were preprocessed similarly to the NMF case. In this case, the model assumes that the documents follow a Dirichlet distribution over topics and topics over words. Thus, it inherently allows having more than one topic in each document. The association matrix was calculated as with the other models. Only the UN training texts were used. Note that this method has been successfully employed to automatically classify the AI curricula of a wide range of universities based on their respective contents [17].

2.4. Bertopic

BERTopic is a topic modeling technique very similar to Top2Vec since both are unsupervised clustering-based techniques [14]. BERTopic extracts coherent topic representation via implementing a class-based variation of the term frequency-inverse document frequency (TF-IDF). The steps it follows are:

- Generating the document embeddings with a pretrained transformer-based language model. The embedded words which are semantically similar will be placed close to each other in semantic space. In this way, document-level information is extracted from the corpora.
- The document embeddings are dimensionally reduced. This is because as data increases dimensionality, the distance to the closest point tends to approach the distance to the farthest point. As a result, in high dimensional space, spatial locality becomes ill-defined, and distance measures differ little [14].
- A density-based method cluster is created. This technique assumes that words near the cluster’s centroid are most representative of that cluster. However, in practice, a cluster will not always lie within a sphere around a cluster centroid which might conduce to the extraction of misleading topics.
- Topics vectors are extracted from the cluster. A class-based version of TF-IDF is used to overcome the limitation of the centroid-based

perspective. This has the advantage of separating the clustering technique from the topic generation, allowing more flexibility.

2.5. Voting

A combination of the previously described model is used to take advantage of their respective strengths, as the models complement each other. After a careful study, one document is linked to an SDG if:

- Any model’s score on an SDG is greater than 0.4 (maximum 0.5), or
- The model’s score on an SDG is greater than 0.1 for LDA and Bertopic.

Using this voting system, we successfully classified 81% of the papers based only on the information in the abstract.

2.6. Database and implementation

Regarding the database, we have downloaded 820,000 documents, comprising articles, conference papers, and books from the Scopus database [18]. The search criterion relied on seeking the words “aerospace,” “aeronautics,” “aeronautical,” and “aviation” in all the metadata of the papers. We selected papers from 2011 to 2020, saving the following data:

- Abstract
- Year
- Citations, as of November 2022.
- Country
- Keywords
- Open-access information

For obvious reasons, the language of the document must be English. This procedure may lead to over-represented affiliations in English, and some papers of one of the authors, *i.e.*, [19,20], are not found based on these keywords. However, the casuistic can be extremely long, and it is nearly impossible to add every possible author to the list. Nevertheless, the number of papers studied is high. We firmly believe it represents the state of Aerospace Engineering to SDGs, as we are analyzing a production of more than 80,000 papers yearly. The set was downloaded in packages of around 20,000 documents each, taking special care of not repeating any document. Finally, around one hundred papers were discarded because they did not contain an abstract.

To summarize, and following the flowchart of Fig. 1, for every document, we have performed the following algorithm:

1. Extract the abstract and metadata from a CSV file.
2. Lemmatize and remove any non-ASCII character.
3. Compute the score for every SDG and every method α_x^{sdg} .
4. Evaluate the score for every SDG, following the rules of the box of Fig. 1.
5. Extract the SDG with the maximum score.
6. Save this SDG with the document’s metadata to an output file.

This algorithm was implemented in Python version 3.9. The code is easily parallelizable, as every document can be run independently. We ran it on a typical computer, taking less than 3 hours to classify all the abstracts.

3. Results

To study the results of our analysis, we will use the term frequency, defined as

$$F = \frac{D_{sdg}}{D_{total}}$$

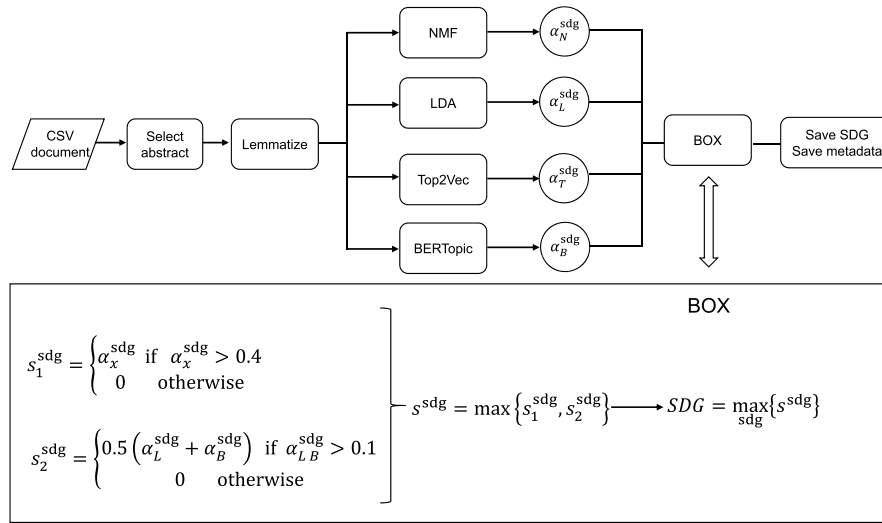


Fig. 1. Flowchart of the ASDG framework, where α_x^{sdg} stands for the score in method x for SDG sdg . This process has been carried out for the 820,000 documents in the database.

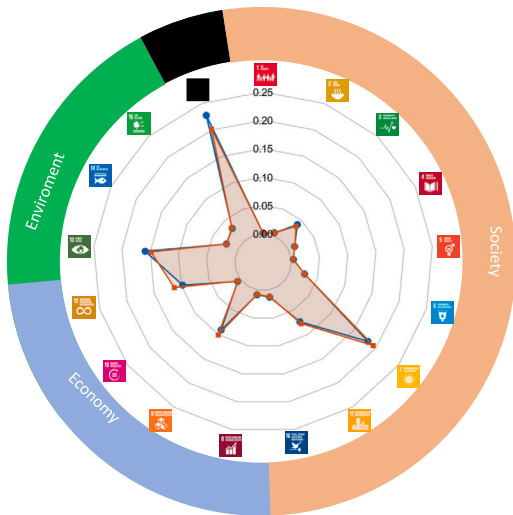


Fig. 2. Distribution of the documents in frequency grouped in two sets, namely 2011–2015 (blue circles) and 2016–2020 (red squares). Frequency is obtained by dividing the papers assigned to a particular SDG over the total papers of that period. Note that the SDGs are grouped into the three categories reported in [2], i.e. Society, Economy, and Environment. The black square indicates the fraction of documents ASDG could not identify.

In this equation, D indicates any set of abstracts with a particular restriction. For example, D could be the set of all papers published in 2011. D_{total} is the total number of papers on that set, and D_{sdg} is the subset of papers identified for a particular SDG. In many cases, this parameter will be preferable to the number of documents. In every case, the definition of the particular subset will be absolutely clear.

As the SDGs were launched in 2015, there has been enough time to see their introduction’s consequences. The global picture is shown in Fig. 2. Here we show the frequency of every SDG and the non-classified abstracts. To facilitate the presentation of the data, we have divided the data set into two large groups, before and after the adoption of the SDGs in 2015. The years 2011–2015 are represented by blue circles, and 2016–2020 by red squares. We have also grouped the SDGs following the classification of [2]. Several ideas can be gained from this image. First of all, there is a large number of unidentified papers. There are two causes for this:

Table 1

Frequency expressed as a percentage of selected SDGs in 2011 and 2020. The last row shows the difference between these two rows.

SDG	3	7	9	11	12	13	15
2011	5.30	21.92	12.8	9.85	12.1	20.4	4.00
2020	4.84	24.10	13.1	9.80	14.43	17.39	4.16
Diff	-0.46	2.19	0.30	-0.05	2.29	-2.96	0.16

- This study has been done with abstracts, which makes the identification more difficult.
- We have used a high threshold, avoiding false identifications as much as possible.

Second, some SDGs seem to be more important in Aerospace Engineering, namely:

- Society: SDGs 3 (good health), 7 (clean energy) and 11 (sustainable cities).
- Economy: SDGs 9 (industry) and 12 (responsible consumption).
- Environment: SDGs 13 (climate action) and 15 (life on land).

After studying the global picture, we will focus on these SDGs.

Finally, the variation after the introduction of the SDGs is small. There is a positive point here: ASDG can better identify the abstracts, probably because researchers are more aware of the SDGs. However, there is a negative aspect too: the frequency of papers about climate change is lower after adopting SDGs than before.

This global situation is completed with Fig. 3. The number of papers has almost doubled during the last ten years. This happens for all the SDGs, and the variations are small. Thus, as the number of documents increases steadily by at least 6% yearly, we think the frequency is a better parameter than the number of papers. Fig. 3 (right) also shows that the number of unidentified papers is constantly reducing, reinforcing the idea that researchers are steering their work toward fulfilling the SDGs.

Apart from that, no clear pattern emerges from Fig. 3. The trends in the most popular SDGs seem to be present before the appearance of SDGs. Some further insight can be gained from Table 1. As we can see, SDGs 7 and 12 have attracted more attention. As we mentioned earlier, this has partly been at the cost of SDG 13. This means that climate action is getting reduced as responsible production grows. As these SDGs are tightly coupled, this is probably not as serious as it

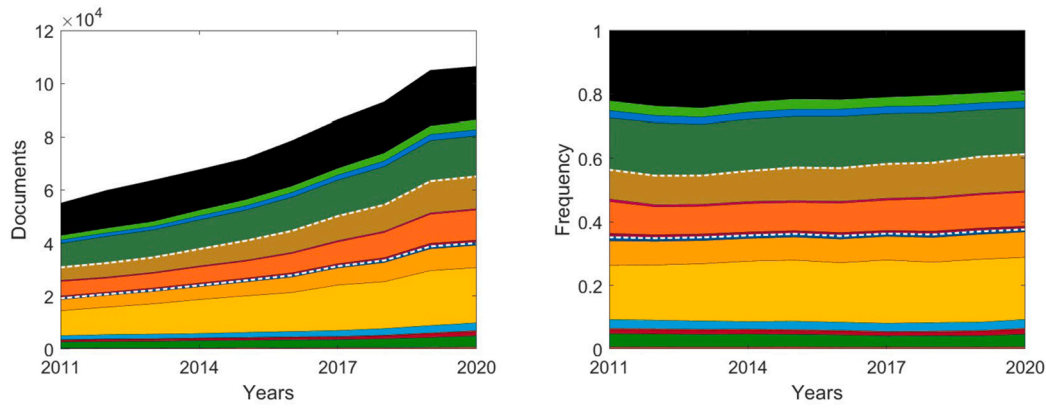


Fig. 3. Distribution of the database in terms of absolute numbers (left) and frequency (right) for every year. The SDGs are represented by their colors, starting from 1 at the bottom and following the order in Fig. 2 in a counterclockwise sense. The black region corresponds to unidentified abstracts, and the white dotted lines indicate the transition between the three large groups.

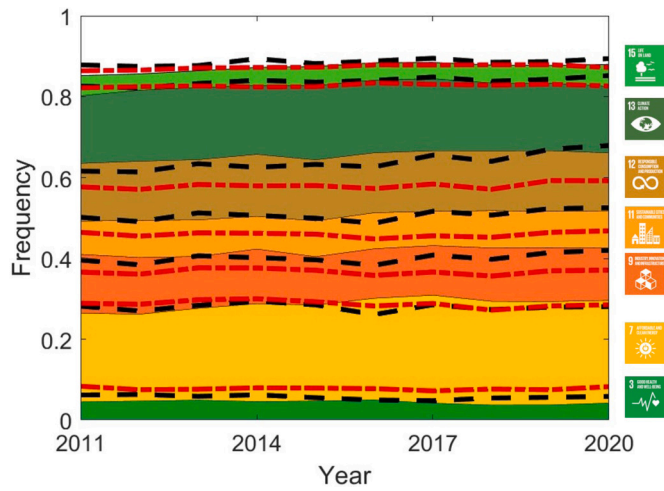


Fig. 4. Distribution in the frequency of selected SDGs by year. Background: China, dashed lines: European Union, and ash-dotted lines: USA.

seems from the point of view of climate action. It could be that further research in engineering (7 and 12) is necessary for advancing in SDG 13.

ASDG can also analyze the priorities of different countries and supranational entities. In Fig. 4, we can see the relative importance of the selected SDGs for the USA, China, and the European Union. Since ten years ago Aerospace Engineering community in the USA has been focused on climate action. However, China seems to focus on engineering and production instead of climate action. Quite curiously, the European Union is closer to China than the USA. In this case, extensive effort is devoted to SDG 9. Also, SDG 12 is receiving much attention, clearly growing with the years.

Finally, we will focus our analysis on the number of citations. Sadly, this is the first tool to measure the quality of a researcher. Thus, it is essential to identify if some fields are preferred because the number of citations working on them is more significant. In Fig. 5, we have sorted the SDGs by the mean citation index (MCI), i.e., the number of citations divided by the number of papers. The size of the dots indicates the order considering only the most cited document. The largest marker corresponds to the maximum. Finally, the lines' width between years i and $i + 1$ indicates the paper's availability as open access (OA) in the year $i + 1$. Again, the wider the line, the greater the percentage of OA papers.

Two groups emerge in this figure. Top cited SDGs, 3, 12, 7, and 15, and less cited SDGs, 9, 12, and 13. This last case is curious since there are highly cited papers in this SDG (years 12, 15, 16, and 17), but it is

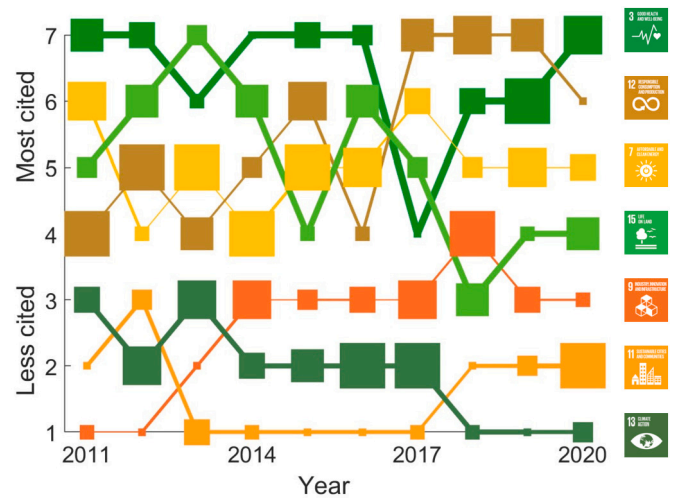


Fig. 5. Sorting of the most important SDGs identified by their colors. The position of every dot indicates the order taking into account the mean citation value of the SDG that year. The size of the marker indicates the order considering only the most cited paper. The width of the line indicates the order considering the percentage of Open-access documents.

almost always the SDG with the lowest MCI and a high percentage of OA papers. This could indicate that many documents in this SDG receive very few (or no) citations, so apart from receiving less attention every year, there are many low-quality papers. On the other hand, SDGs 3, 7, and 12 exhibit a very good MCI. In the case of SDG 7, the number of citations appears to be independent of the OA percentage.

Finally, it is also important to note that SDG 3 results can be biased as health-related publications use to have a very large number of citations.

4. Conclusions

In this work, we have used the tool ASDG [5] to study the alignment of Aerospace Engineering with the Sustainable Development Goals of the United Nations. ASDG uses NMF, LDA, TopVec, and Bertopic methods to identify the SDGs. We identified two main questions in the introduction. First, the introduction of the SDG Agenda did not have a clear impact on the scientific production of the Aerospace Engineering community. This result is quite concerning, as the community's attention is drifting apart from extraordinary and urgent challenges such as climate change, although differences among countries exist. Second, we have identified 7 SDGs to which most of the works on Aerospace Engineering belong.

One possible limitation of this work is that it is based on abstracts, with far better availability than the whole paper. Moreover, at this point, we can only see if there is a connection, but whether this relationship is either positive or negative can not be asserted. So, after this first analysis using ASDG, our objective is to identify the targets of the different SDGs and analyze their positive and negative interactions.

Thus, funders and research organizations can use ASDG to very quickly identify if they are walking on the right path to prosperity. Science maps are possible, even going to smaller organizations like universities. Finally, we recommend that any reader interested in using the code contact us for further collaboration.

CRedit authorship contribution statement

ASR: Software, Validation, Investigation, Visualization.

OGO: Software, Validation, Investigation, Visualization, Writing - Review & Editing.

JAC: Conceptualization, Project definition, Methodology, Investigation, Visualization, Resources, Writing - Review & Editing, Funding acquisition.

HE: Investigation, Writing - Review & Editing.

FM: Writing - Review & Editing.

ER: Investigation, Writing - Review & Editing.

FFN: Investigation, Writing - Review & Editing.

JGM: Writing - Review & Editing.

RV: Conceptualization, Project definition, Methodology, Investigation, Visualization, Resources, Writing - Review & Editing, Funding acquisition.

SH: Software, Conceptualization, Project definition, Methodology, Investigation, Visualization, Resources, Writing - Original draft, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

RV and FFN acknowledge the support of the KTH Climate Action Centre and Digital Futures. SHC is partially funded by project PID2021-128676OB-I00 by Ministerio de Ciencia, innovación y Universidades / FEDER.

References

- [1] UN General Assembly (UNGA), Transforming our world: the 2030 Agenda for Sustainable Development, Resolut. A/RES/70/1, 25, pp. 1–35, available at <https://sustainabledevelopment.un.org/index.php?page=view&type=111&nr=8496&menu=35>, 2015.
- [2] Ricardo Vinuesa, Hossein Azizpour, Iolanda Leite, Madeline Balaam, Virginia Dignum, Sami Domisch, Anna Felländer, Simone Daniela Langhans, Max Tegmark, Francesco Fuso Nerini, The role of artificial intelligence in achieving the sustainable development goals, *Nat. Commun.* 11 (2020) 233.
- [3] Shivam Gupta, Sigmone D. Langhans, Sami Domisch, Francesco Fuso-Nerini, Anna Felländer, Manuela Battaglini, Max Tegmark, Ricardo Vinuesa, Assessing whether artificial intelligence is an enabler or an inhibitor of sustainability at indicator level, *Transp. Eng.* 4 (2021) 100064.
- [4] Ricardo Vinuesa, Beril Sirmacek, Interpretable deep-learning models to help achieve the Sustainable Development Goals, *Nat. Mach. Intell.* 3 (2021) 926.
- [5] Alejandro Sánchez-Roncero, Óscar Garibo i Orts, J. Alberto Conejero, Hamidreza Eivazi, Emelie Rosenberg, Francesco Fuso-Nerini, Javier García-Martínez, Sergio Hoyas, Ricardo Vinuesa, ASDG — an AI-based framework for automatic classification of impact on the SDGs, in: *Proceedings of the 15th International Conference on Theory and Practice of Electronic Governance*, Guimarães, Portugal, ACM, 2022.
- [6] Climate change 2021, the physical science basis, sixth assessment report, available at: https://www.ipcc.ch/report/ar6/wg1/downloads/report/IPCC_AR6_WGI_SPM_final.pdf, 2021. (Accessed December 2021).
- [7] Francesco Fuso-Nerini, Julia Tomei, Long Seng To, Aiduan Borrión, Catalina Spataru, Gabriel Anandarajah, Iwona Bisaga, Mairi Black, Priti Parikh, Vanesa Castán Broto, Ben Milligan, Yacob Mulugetta, Mapping synergies and trade-offs between energy and the Sustainable Development Goals, *Nat. Energy* 3 (2018) 10–15.
- [8] Francesco Fuso Nerini, Benjamin Savacool, Nick Hughes, Laura Cozzi, Ellie Cosgrave, Mark Howells, Massimo Tavoni, Julia Tomei, Hisham Zerriffi, Ben Milligan, Mapping synergies and trade-offs between energy and the Sustainable Development Goals, *Nat. Sustain.* 2 (2019) 674–680.
- [9] J. Jiménez, Near-wall turbulence, *Phys. Fluids* 25 (10) (2013) 101302.
- [10] IATA, Annual review, available at: <https://www.iata.org/en/publications/annual-review/>, 2018. (Accessed 27 September 2018).
- [11] Daniel D. Lee, H. Sebastian Seung, Learning the parts of objects by non-negative matrix factorization, *Nature* 401 (6755) (1999) 788–791.
- [12] Dimo Angelov, Top2vec: distributed representations of topics, <https://arxiv.org/abs/2008.09470>, 2020.
- [13] David M. Blei, Andrew Y. Ng, Michael I. Jordan, Latent Dirichlet allocation, *J. Mach. Learn. Res.* 3 (Jan 2003) 993–1022.
- [14] Maarten Grootendorst, Bertopic: neural topic modeling with a class-based TF-IDF procedure, arXiv preprint, arXiv:2203.05794, 2022.
- [15] Roman Egger, Joanne Yu, A topic modeling comparison between LDA, NMF, Top2Vec, and BERTopic to demystify Twitter Posts, *Front. Sociol.* 7 (2022).
- [16] Grobid, <https://github.com/kermitt2/grobid>, 2008–2022.
- [17] Rana Tallal Javed, Osama Nasir, Melania Borit, Loïs Vanhée, Elias Zea, Shivam Gupta, Ricardo Vinuesa, Junaid Qadir, 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.
- [18] Scopus database, available at: <https://www.scopus.com>, 2011–2020.
- [19] Sergio Hoyas, Martin Oberlack, Francisco Alcántara-Ávila, Stefanie V. Kraheberger, Jonathan Laux, Wall turbulence at high friction Reynolds numbers, *Phys. Rev. Fluids* 7 (Jan 2022) 014602, <https://doi.org/10.1103/PhysRevFluids.7.014602>.
- [20] Martin Oberlack, Sergio Hoyas, Stefanie V. Kraheberger, Francisco Alcántara-Ávila, Jonathan Laux, Turbulence statistics of arbitrary moments of wall-bounded shear flows: a symmetry approach, *Phys. Rev. Lett.* 128 (Jan 2022) 024502, <https://doi.org/10.1103/PhysRevLett.128.024502>.