Impact Assessment and Project Appraisal

ISSN: 1461-5517 (Print) 1471-5465 (Online) Journal homepage: www.tandfonline.com/journals/tiap20

Natural language processing in impact assessment: a review of applications and concerns

Shiu Fung Hung, Pavitra Kumar & Thomas B Fischer

To cite this article: Shiu Fung Hung, Pavitra Kumar & Thomas B Fischer (2026) Natural language processing in impact assessment: a review of applications and concerns, Impact Assessment and Project Appraisal, 44:1, 68-73, DOI: 10.1080/14615517.2026.2620186

To link to this article: https://doi.org/10.1080/14615517.2026.2620186

Published online: 21 Jan 2026.

IMPACT ASSESSMENT AND PROJECT APPRAISAL
2026, VOL. 44, NO. 1, 68–73
https://doi.org/10.1080/14615517.2026.2620186

Natural language processing in impact assessment: a review of applications and concerns

Shiu Fung Hung[a], Pavitra Kumar[b] and Thomas B Fischer[a,c]

[a]Department of Geography and Planning, University of Liverpool, Liverpool, UK; [b]Department of Computing and Mathematics, Manchester Metropolitan University, Manchester, UK; [c]Research Unit for Environmental Sciences and Management, North West University, Potchefstroom, South Africa

ABSTRACT

The rolling out of Artificial Intelligence (AI) applications is extensively discussed among the impact assessment (IA) community. It is in this context that our letter focuses on natural language processing (NLP) applications. We introduce six categories of NLP applications and examine common concerns over their accuracy, quality, and transparency. We emphasize that it is the users' responsibility to ensure quality and transparency when applying AI. We also suggest that it is time to elevate the discussion on the preparation for further AI use in IA.

ARTICLE HISTORY
Received 19 October 2025
Accepted 16 January 2026

KEYWORDS
Impact assessment; natural language processing; artificial intelligence; IAIA2025

Introduction

The application of artificial intelligence (AI) technologies has received much attention by the impact assessment community (as represented by, e.g. the International Association of Impact Assessment – IAIA) in recent years. Over 1000 delegates attended IAIA's 2025 Conference, focused on the rollout of AI technologies and applications in Impact Assessment (International Association of Impact Assessment 2025; Aung and Fischer, 2025). *Impact Assessment and Project Appraisal* issued a call for papers on the 'Emerging Role of AI in the future of Impact Assessment (IA)', and so far, ten articles have been published in the contribution to the round table discussion. In this letter to the editor, we aim to summarise and discuss the applications and concerns raised in the IA community regarding the rollout of AI applications, particularly the natural language processing (NLP) applications.

AI is an umbrella concept which aims at enabling computers to mimic or simulate human thinking and actions. NLP is one of the major subsets of AI, specialising in the understanding and generation of human language (Khurana et al. 2023). Its application includes the classification of text, labelling, discourse analysis, and creative writing (Zhou et al. 2020; Audichya and Saini 2023; Khurana et al. 2023). NLP-focused tasks are the most common type of usage of popular large language models, such as OpenAI's ChatGPT (Chatterji et al. 2025). At the IAIA25 conference, we saw many demonstrations of the uses of NLP applications shared by multilateral financial institutions, authorities, consultant practitioners, and academics. Meanwhile, at the various themed workshops, participants raised concerns over the rollout of AI interventions in IA practices, including the quality of AI products, transparency, and the potential problems in reviewing AI-processed content. Subsequently, in this letter, we reflect on the discussion and look at the causes for concerns. We then provide our suggestions for preparing the AI rollout and elevate the discussion of 'human-in-the-loop' in the use of AI.

Overview of NLP applications in IA

While NLP has a long history of development (Khurana et al. 2023), the latest applications are closely connected to the development and release of large language models (LLMs). LLMs, such as OpenAI's ChatGPT, Anthropic's Claude, and Google's Gemini, are models that pre-trained with a large amount of data. They are built upon the latest transformer architecture, enabling them to understand and generate natural language to perform a wide range of tasks (Stryker 2025). LLMs are currently the most popular type of AI tools, with ChatGPT, one of the most popular models, reporting 700 million users in July 2025 (Chatterji et al. 2025). Most, if not all, NLP applications demonstrated at the IAIA25 conference were performed by LLMs. Subsequently, we focus only on the use of NLP in IA.[1] We neither differentiate between models used by the presenters, nor do we evaluate the performance of these models.

CONTACT Thomas B Fischer Fischer@liverpool.ac.uk Environmental Assessment and Management Research Centre, Department of Geography and Planning, School of Environmental Sciences, University of Liverpool,Roxby Building, L69 7ZT Liverpool, UK

Summarising the presentations at the IAIA25 conference, articles submitted to IAPA and webinars organised by IAIA, six categories of NLP applications in IA were identified, as follows.

AI-powered search and retrieval-augmented generation (RAG)

Many AI-powered library and Q&A chatbot tools have been introduced around the world. IA specialised examples include myESRA platform (Almasi Simsic 2024), the Danish EA Hub (Kørnøv et al. 2025), and the Canadian Open Science and Data Platform (Habbane 2025). The Dutch IPLO AI presented by van Ravesteijn et al. (2025) was not specific to IA but works in a similar fashion. In comparison to general AI assistant tools (e.g. the basic version of ChatGPT), which search the internet and utilise open information to generate answers for users, RAG searches, and retrieves information from a more specified library and generates answers for users using the retrieved information. It allows AI to provide more accurate responses to the specific domains it is set up for (Almasi Simsic 2024; Kørnøv et al. 2025; van Ravesteijn et al. 2025).

Document automation

Document automation refers to the use of the natural language generation capability to generate and draft a document. In professional use, the user should provide clear instructions and data. An example presented was C-ESRS AI Assistant developed by the World Bank, which leverages the World Bank's data and Project documents to generate drafts for review and refinements by their environmental and social specialists (Song 2025).

Document and content classification

Document classification is a kind of AI-powered document review that involves the task of extracting, analysing, and classifying information within a document. Two types of AI Document Classification were presented. First, there is the use of AI for screening (Cilliers et al. 2025; O'Mahony 2025). In these applications, AI were used to identify and classify if the content in the document would trigger specified screening criteria, for example, if a project profile contains certain activities (Cilliers et al. 2025). The second type of AI-powered Document Review presented was the experimental literature review by Bond et al. (2024), in which AI is used to identify and categorise key issues from published documents.

Machine translation

Machine translation is a straightforward tool that has the AI read the document and generate an alternative version in another language. An example is the Chichewa NLP Project presented by Fagan et al. (2025). Here, an AI translator is used to translate documents and consultation materials into local languages in Africa and translate local comments for practitioners.

Sentiment analysis

Sentiment analysis is a form of text analysis that examines opinions and impressions regarding various topics and subjects (Wankhade et al. 2022). The MELINA model, developed by the International Finance Corporation (IFC), demonstrated the use of AI sentiment analysis for IA. Here, AI identifies Environmental, Social, and Governance (ESG) related keywords and classifies the sentiment as positive, negative, or neutral. Then, it produces an overview of ESG risks presented in the document (Curmally et al. 2022).

Semantic comparison

NLP has the capacity for semantic understanding of documents. There were two demonstrations about using AI to compare and analyse the contents of multiple sets of documents. Gurrieri et al. (2025) showed that AI could be used to analyse the gap among IFC's performance standard, Environmental Health and Safety guidelines and Local laws. Also, Ishii and Kameda (2025) used AI to identify changes made to an environmental plan by comparing the changes in the drafts, public comments submitted, and the response of government.

In summary, various forms of innovative applications of NLP have been discussed in the field of IA. Currently, applications are in different stages of their development. Some of the AI-powered database and library tools have been released, while many applications for screening and document analysis are at pilot or experimental stages. AI development has been rapid in recent years, and there is an expectation that some of the experimental tools will soon become mature, and that other innovative uses will be developed.

Examining concerns and issues

In the rollout of AI applications in IA, numerous questions and concerns have been raised. Many of the concerns are about using AI in general, rather than specifically for NLP applications. Also, many of the discussions about these concerns are at the surface layer of AI applications without examining the causes of the problems. With AI applications rapidly developing, there is ample need to elevate the discussion. With

a focus on NLP applications, we selected three critical issues and looked more deeply into each of them.

Accuracy and quality of AI products

The accuracy and quality of AI products are one of the key topics in AI development, and concerns have been raised by various authors and commentators (see e.g. Bond et al. 2024; Cilliers et al. 2025; van Ravesteijn et al. 2025), including: Missed or wrong classification (Bond et al. 2024; Cilliers et al. 2025); messed up or made up answers (van Ravesteijn et al. 2025) and wrong interpretations (Bond et al. 2024).

The accuracy and quality of an AI product depend on three major factors: Model robustness, appropriate setup, and users' input. AI models are trained with different datasets and for different purposes. AI models that are trained and calibrated with IA domain-specific data would be essential for more precise and quality AI interventions in IA (Aung and Fischer 2025). If a general-purpose model is to be used, the model performance would also be expected to perform better when fine-tuned on data in a field (see OpenAI 2024). Furthermore, users' prompt inputs play an important role in a model's output quality. AI models would not understand and provide additional information beyond users' input. It is expected to answer specifically based on users' inputs. Hence, better prompt inputs will generate more specific and accurate output in comparison to a general prompt (see example at Cilliers et al. 2025).

Many concerns regarding the quality of the AI products focus on the performance of individual AI models. Meanwhile, the need for model fine-tuning and users' techniques in prompting has remained poorly discussed.

Transparency and AI black box

Transparency is a common concern raised in the application of AI models, and the discussion of AI transparency is usually associated with the 'black box' problem of AI (e.g. Gupta et al. 2025; Jiricka-Pürrer and Stöglehner 2025). However, examining concerns, it becomes clear that the problem is not always caused by the 'black box'.

The 'black box' problem refers to the algorithm system of AI remaining opaque or hidden from human comprehension, which is exacerbated by AI models that utilise deep learning algorithms and multi-layered neural networks (von Eschenbach 2021; Kosinski 2024). Meanwhile, there are two dimensions of transparency concerns of AI applications. First is the transparency between the users' instructions and the AI product. Second is the transparency of AI involvement between the user and the audience.

For the first dimension, it is connected to the ambiguity of the users' instructions and the AI's initial black box in processing the instructions and producing products. As explained in the previous section, AI results are in accordance with users' prompt inputs. AI interprets the users' prompts and creates results using complex parameters it developed through its data training (Kosinski 2024). This is mostly opaque. However, the freedom of creativity in AI's processing is determined by the ambiguity of users' instructions. Unclear instructions provide more room for the AI to interpret, thus creating more unexpected results; clear instructions with high-quality data minimise the AI's creativity and make the product align more closely with users' expectations. As such, the users take a significant portion of responsibility AI's transparency. It would be an accuracy problem if the AI produces results that contradict the users' instructions.

It is important that AI products inherit users' instructions and the process done inside the AI black box. It would not explain which part of the product was created based on users' instructions and which part was generated by the AI's own interpretation. This leads to concerns about transparency between the AI usage and the audience, and questions about how one could review AI-generated content created by others, which we examine in the following section.

Reviewing AI-generated content by others

The question of how to review AI-generated content or documents provided by others was frequently raised by participants during the IAIA25 conference. It not only refers to the documents submitted by IA practitioners, but also documents submitted by the public (see Hung 2025). In this context, we have observed recent cases where campaign groups use AI generators to draft and submit objection letters for planning applications.[2]

Reviewing AI-generated documents involves two parts: identifying the presence of AI generation and assessing the quality of AI-generated content. Identifying AI-generated text can be done by examining the sentence structure, vocabulary and word usage, which usually contains notable characteristic differences from human-written text (Khan et al. 2025). There are also AI-generated text detector tools on the market (e.g. Quillbot and Pangram Labs), despite that the accuracy and reliability of these tools are often found to be inconsistent (Elkhatat et al. 2023; Weber-Wulff et al. 2023).

Reviewing the quality of AI-generated content provided by others remains a critical challenge. There is no known method to reverse-engineer the AI product to find the user instructions provided to the AI model, the data it used or the creation process by the AI. To review the genuineness and quality of an AI-generated work,

it is essential to know users' instructions and data inputs at the prompting stage. By understanding the prompt, the audience can assess the user's intent, the validity of the data it operates on, and the areas of AI creation, including potential 'hallucinations'. Without knowing the prompts, it would rely on the reviewers' professional knowledge about the subject to verify the facts, logic and reasoning of the document.

Addressing concerns

In response to the raised concerns over the rollout of AI, in particular, NLP applications, we will subsequently provide suggestions in three areas: quality data and database investment; communication and disclosure; and human-in-the-loop system.

Quality data and database investment

Data are the core component in AI applications, and the quality of input data directly affects the quality of the output. Even the latest AI models do not negate the need for users to invest in their own high-quality data and fine-tune models for their specific tasks.

The data provided by a user would be the key to AI's performance. Professional users are expected to fine-tune and teach a model the approach to tasks and desired behaviour. It requires quality instructional datasets and historical data that can identify expected outcomes. AI would then follow it. For individual users, clear instructions and quality structured data are essential for producing high-quality output. Up-to-date, localised, and contextualised information for the tasks can only be provided by the end-users. A library of good prompt instructions and example documents also helps create more consistent results, especially when using a non-fine-tuned model. For both, the organisation and the end-user, collecting and building a comprehensive database is essential for consistent and high-quality output.

Communication and disclosure

Transparency of AI applications is associated with the clarity of user instructions and the initial AI black box that processes these instructions. It is ultimately the user's responsibility to ensure that AI products align with their expectations and communicate the AI involvement to the user/audience.

There has been some discussion about the transparency of AI in IA practices (e.g. Bingham et al. 2025). However, with the widespread adoption of AI among all stakeholders, at this point in time, the discussion should be elevated. The level of detail needed about AI's involvement depends on who would be reading it. For the general user/audience, a simple note acknowledging the use of AI and its purpose may be enough. For reviewers and decision-makers, knowing the instructions and data used in the prompts would be necessary to assess the quality of the document. Meanwhile, the private data used in the prompt may not be suitable to be disclosed to an external audience. Therefore, there needs to be communication among the stakeholders and agreements on the level of disclosure on AI's involvement and how the necessary information could be disclosed to the concerned parties. This also includes the documents provided to and received from the public.

Human-in-the-loop system

Human-in-the-loop (HITL) is a core principle in the development and implementation of AI models. It primarily refers to utilising human expertise to improve machine learning through active supervision and feedback (Kumar et al. 2024). The concept also extends to the use of AI. Aung and Fischer (2025) proposed three mechanisms: 'manual review of flagged AI outputs', 'threshold-based escalation for complex or sensitive assessments', and 'shared responsibility protocols between AI systems and domain experts'.

At the practical scale of using AI in IA, human expertise would be necessary in the loop. AI models may not contain the latest domain-specific information; thus, collaborations between AI and IA experts would be necessary for fine-tuning the model when used professionally. The products of AI are sensitive to the quality of instructions and data provided for the task; human experts are necessary to ensure that correct instructions and data are provided to the AI, as well as to verify if the products fit their expectations. The HITL system here would provide a safety net for quality assurance.

At the larger scale, IA involves a broad range of stakeholders, conducting numerous types of tasks and having decisions made throughout the process. Human experts would be present. However, this does not necessarily mean that the objective of HITL would be met. Without a proper working and organisational environment, the presence of human actors and experts would not be enough to make HITL effective. It requires users and stakeholders to have a critical mindset to review and identify potential errors in the system, crucial information to be relayed and an organisational environment that enables investigation and adjustment to be made following, e.g. warning signals and feedback. In this regard, the HITL system should not only serve merely as a safety net, but also as an active, continuous organisational learning loop to improve the collaboration between AI and human users. In the rollout of AI in IA, the IA community need to be prepared to act and fulfil the human functions in the loop.

Discussion of the future developments of AI

AI is currently developing rapidly. Model updates keep improving AI's performance with regard to various benchmarks (e.g. see OpenAI 2025). There have been IA domain-specific models released or are in development (e.g. MALENA by IFC (https://malena.ifc.org/), myESRA by E&S solution (https://www.essolutions.ai/), and Nyxion by Global Impact Assessment (https://www.nyxion.tech/)). The development of explainable AI (xAI) technology can also help visualise and explain the approach used in the AI thinking process – basically describing how the prediction results are derived, making it easier for users to evaluate and interpret the models (Dwivedi et al. 2023). However, as we emphasised throughout this letter, the quality of users' prompts is always the most important factor in determining the quality of the AI products. Continuous skill training would be needed for the users to keep up with the latest developments in AI.

We only focused on the NLP applications presented by the IA community in this letter. NLP could be used to assist in other qualitative analysis, such as data coding (Wang et al. 2022) and discourse network analysis (Smith et al. 2021). There are many other types of AI models that utilise other subsets like computer vision and robotics. There is also a trend of integrating/collaborating with multiple AI models for complicated tasks (e.g. see Google Cloud 2025). For example, collaboration between NLP and computer vision could facilitate text-to-image and image-to-text applications (Betker et al. 2023). We only covered a small portion of the potential AI interventions in IA. While our discussion involves many topics that would apply to the general use of most types of AI, there needs to be more comprehensive studies to investigate and expand the AI in IA discussions to cover the wide range of possible interventions.

Besides technological advancement, it is also noted that the legislation and policies around the development and use of AI are being updated at this moment (e.g. the implementation of the EU AI Act). Users should be aware of the changes in restrictions, such as the use of copyrighted material and the processing of sensitive data with AI.

Conclusion

Through reviewing the IAIA25 proceedings and IAPA publications, we identified six categories of potential NLP AI interventions demonstrated by the IA community. We acknowledge that there are common concerns over the accuracy, quality, and transparency of the use of AI in IA. However, users' instructions and data input would be the most determining factor for AI output, and it would be the responsibility of the users to ensure quality and transparency of the use of AI, as well as to communicate with the audience. We suggest that attention and elevated discussions are needed on the investment of quality data, communication, and a Human-in-the-loop system in the rollout of AI in IA. As our letter only covers the NLP applications, further studies on the other types of AI applications are needed to facilitate a more comprehensive discussion on the application of AI.

Notes

1. This letter is not a discussion of the use of LLMs. The scale of LLMs has significantly expanded over the last two years to include other AI subsets such as computer vision and sound recognition (e.g. Google's Vertex AI Studio), which are beyond the scope of this letter.
2. For example, the 'Hills off our hills' campaign (https://handsoffourhills.co.uk/site/.), the objection letters can be accessed via the Scottish Energy Consent Unit's website (https://www.energyconsents.scot/.)

Acknowledgments

The production of this paper is self-funded by the authors. There is no conflict of interest to declare.

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Thomas B Fischer http://orcid.org/0000-0003-1436-1221

References

Almasi Simsic I. 2024. Digital impact assessment series: AI & IA: keeping the human-in-the-loop – technology supporting E&S management (online seminar).International Association for Impact Assessment, Fargo, ND.

Audichya MK, Saini JR. 2023. ChatGPT for creative writing and natural language generation in poetry and prose. Proceedings of 3rd International Conference on Advanced Computing Technologies and Applications, ICACTA 2023; Mumbai, India: Institute of Electrical and Electronics Engineers Inc. https://doi.org/10.1109/ICACTA58201.2023.10392805

Aung TS, Fischer TB. 2025. Impact assessment in the age of artificial intelligence: reflections from IAIA25. Impact Assess Project Appraisal. 43(3):166–170. https://doi.org/10.1080/14615517.2025.2505265

Aung TS, Fischer TB 2025. Impact assessment in the age of artificial intelligence: reflections from IAIA25. Impact Assess Project Appraisal. 43(3):166–170.

Betker J et al. 2023. Improving image generation with better captions. [accessed 2025 Nov 10]. https://cdn.openai.com/papers/dall-e-3.pdf

Bingham C et al. 2025. Principles for use of AI in IA. [accessed 2025 Sep 30]. Fargo. https://iaia.org/ai-and-emerging-technology/2/

Bond A et al. 2024. Using an artificial intelligence chatbot to critically review the scientific literature on the use of artificial intelligence in environmental impact assessment. Impact Assess Project Appraisal. 42(2):189–199. https://doi.org/10.1080/14615517.2024.2320591

Chatterji A et al. 2025. How people use ChatGPT *. [accessed 2025 Sep 22]. San Francisco, US. https://openai.com/index/how-people-are-using-chatgpt/

Cilliers D, Bond A, Retief F, Alberts R, Roos C. 2025. Large language model-assisted EIA screening: a case study using GPT. Impact Assess Project Appraisal. 43(4):267–277. https://doi.org/10.1080/14615517.2025.2523628

Curmally A, Skene F, Sandwidi B. 2022. Digital IA: what might artificial intelligence offer impact assessment? – the IFC's Malena. [accessed 2025 Oct 5]. https://training.iaia.org/digital-ia-what-might-artificial-intelligence-offer-impact-assessment-the-ifcs-malena/

Dwivedi R et al. 2023. Explainable AI (XAI): core ideas, techniques, and solutions. ACM Comput Surv. 55(9):1–33. https://doi.org/10.1145/3561048

Elkhatat AM, Elsaid K, Almeer S. 2023. Evaluating the efficacy of AI content detection tools in differentiating between human and AI-generated text. Int J Educ Integr. 19(1). https://doi.org/10.1007/s40979-023-00140-5

Fagan B, Ayanoglu C, Ozel O. 2025. Breaking barriers: AI generative tools for inclusive stakeholder engagement in Sub-Saharan Africa's education sector. [accessed 2025 Sep 29]. https://iaia.org/iaia25-conference-proceedings/

Google Cloud. 2025. Innovate faster with enterprise-ready AI, enhanced by Gemini models. [accessed 2025 Apr 10]. https://cloud.google.com/vertex-ai

Gupta R, Moser RE, Singh G. 2025. A review of AI techniques and their application to stakeholder engagement. [accessed 2025 Sep 25]. https://iaia.org/iaia25-conference-proceedings/

Gurrieri S, Vaccari I, Desogus E, Napoli E, Schirinzi E. 2025. AI tool for enhancing the process of gap analysis in ESIA and ESDD studies. [accessed 2025 Sep 25]. https://iaia.org/iaia25-conference-proceedings/

Habbane M. 2025. The open science and data platform with generative AI for impact assessment. [accessed 2025 Sep 29]. https://iaia.org/iaia25-conference-proceedings/

Hung SF. 2025. Publicly available AI tools will change public participation, are we prepared for it? [accessed 2025 Sep 25]. https://iaia.org/iaia25-conference-proceedings/

International Association of Impact Assessment. 2025. Proceedings - IAIA25. [accessed 2025 Sep 19]. https://2025.iaia.org/proceedings/

Ishii K, Kameda A. 2025. Leveraging LLMs to evaluate public comments on Japan's environmental plans. [accessed 2025 Sep 25]. https://iaia.org/iaia25-conference-proceedings/

Jiricka-Pürrer A, Stöglehner G. 2025. AI in strategic planning and assessment–a game changer in decision-making or a risk for reproduction of hidden biases? Impact Assess Project Appraisal. 43(4):319–323. https://doi.org/10.1080/14615517.2025.2542081

Khan HU, Naz A, Alarfaj FK, Almusallam N. 2025. Identifying artificial intelligence-generated content using the DistilBERT transformer and NLP techniques. Sci Rep. 15 (1). https://doi.org/10.1038/s41598-025-08208-7

Khurana D, Koli A, Khatter K, Singh S. 2023. Natural language processing: state of the art, current trends and challenges. Multimed Tools Appl. 82(3):3713–3744. https://doi.org/10.1007/s11042-022-13428-4

Kørnøv L, Lyhne I, Rasmus Sveding K. 2025. Laying the foundation for generative AI and multi-agent systems in environmental assessment: building a curated dataset from the Danish EA Hub. Impact Assess Project Appraisal. 43 (4):253–266. https://doi.org/10.1080/14615517.2025.2532919

Kosinski M. 2024. What is black box artificial intelligence (AI)? [accessed 2025 Oct 1]. https://www.ibm.com/think/topics/black-box-ai

Kumar S et al. 2024. Applications, challenges, and future directions of human-in-the-loop learning. IEEE Access. 12:75735–75760. https://doi.org/10.1109/ACCESS.2024.3401547

O'Mahony T. 2025. Sea practice in Ireland - SEA effectiveness, SEA action planning and AI. [accessed 2025 Sep 29]. https://iaia.org/iaia25-conference-proceedings/

OpenAI. 2024. Fine-tuning [Internet]. [accessed 2024 Oct 15]. https://platform.openai.com/docs/guides/fine-tuning/fine-tuning-integrations

OpenAI. 2025. Introducing GPT-5 [internet]. [accessed 2025 Sep 30]. https://openai.com/index/introducing-gpt-5/

Smith TB, Vacca R, Mantegazza L, Capua I. 2021. Natural language processing and network analysis provide novel insights on policy and scientific discourse around sustainable development goals. Sci Rep. 11(1). https://doi.org/10.1038/s41598-021-01801-6

Song HJ. 2025. AI-powered environmental and social impact assessment platforms of the World Bank. [accessed 2025 Sep 29]; https://iaia.org/iaia25-conference-proceedings/

Stryker C. 2025. What are large language models (LLMs)? [accessed 2025 Sep 24]. https://www.ibm.com/think/topics/large-language-models

van Ravesteijn MP, Muspitasari Y, van Eck G. 2025. How can chatbots help to improve government information provision on EA? [accessed 2025 Sep 25]. https://iaia.org/iaia25-conference-proceedings/

von Eschenbach WJ. 2021. Transparency and the black box problem: why we do not trust AI. Philos Technol. 34 (4):1607–1622. https://doi.org/10.1007/s13347-021-00477-0

Wang Y, Tian J, Yazar Y, Ones DS, Landers RN. 2022. Using natural language processing and machine learning to replace human content coders. Psychol Methods. 29 (6):1148–1163. https://doi.org/10.1037/met0000518

Wankhade M, Rao ACS, Kulkarni C. 2022. A survey on sentiment analysis methods, applications, and challenges. Artif Intell Rev. 55(7):5731–5780. https://doi.org/10.1007/s10462-022-10144-1

Weber-Wulff D et al. 2023. Testing of detection tools for AI-generated text. Int J Educ Integr. 19(1). https://doi.org/10.1007/s40979-023-00146-z

Zhou M, Duan N, Liu S, Shum HY. 2020. Progress in neural NLP: modeling, learning, and reasoning. Engineering. 6(3):275–290. https://doi.org/10.1016/j.eng.2019.12.014