Scopus® AI

# Scopus AI champions the Elsevier Responsible AI principles

Elsevier has a long history of partnering with the academic community to developing standards in new and emerging areas, to enable integration and adoption of new technologies, whilst championing standards and policies that preserve academic integrity. For more than a decade, Elsevier has been using AI and machine learning technologies responsibly in our products. With Scopus AI, our intuitive and intelligent search tool powered by generative AI (GenAI), we adhere to Elsevier's core Responsible AI principles and abide by strict guidelines on how personal data is collected and handled.

## Scopus AI champions the Elsevier Responsible AI principles

The following core principles provide guidance for Elsevier teams working on designing, developing, and deploying machine-driven insights, and Scopus AI carefully considers each principle across its development processes:

1. We consider the real-world impact of our solutions on people.
2. We take action to prevent the creation or reinforcement of unfair bias.
3. We can explain how our solutions work.
4. We create accountability through human oversight.
5. We respect privacy and champion robust data governance.

ELSEVIER

## Principle 1: We consider the real-world impact of our solutions on people

- **We worked closely with the research community** throughout the development process, seeking to better understand their needs, concerns and potential risks.
- We invited thousands of institutions, academics, and other potential users worldwide to test early versions of the solution, and continue to **incorporate their feedback as we refine and iterate Scopus AI**.
- We have a **cross-functional and diverse team** that contributes to our decision-making and assessments of the solution.

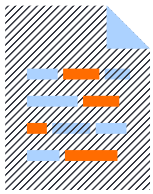

## Principle 2: We take action to prevent the creation or reinforcement of unfair bias

- **Scopus AI grounds its responses in trusted, vetted Scopus content**, aiming to reflect a comprehensive view of the academic literature.
- The **prompt engineering that guides our large language models** is extremely strict, with clear instructions and scope. When Scopus AI makes a claim or assertion about something, a reference is required.
- Scopus AI is one of the first products to pioneer what is rapidly becoming the **gold standard for LLM use** — the retrieval augmented generation (RAG) fusion model. It's an approach that **improves the quality of both vector search retrieval and the generation of LLM summaries**. Together, these factors reduce the risk of hallucinations (or false AI-generated information).

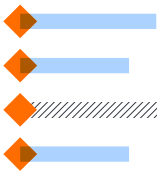

## Principle 3: We can explain how our solutions work

- Scopus AI's responses are based on **trusted, peer-reviewed Scopus content**. Any claims or assumptions **must be backed up by a cited reference**.
- Our prompting aims to inform the user of the confidence of our responses based on the content provided:
  - A **direct response** is made when confidence in the source material is at a high level.
  - An **inferred response** occurs when confidence in the source material is a medium level.
  - **No response** is made when confidence in the source material is low.

## Principle 4: We create accountability through human oversight

- We conduct periodic **Algorithmic Impact Assessments**, in which we acknowledge our risks, weaknesses and strengths.
- We leverage both a **quality and harmful bias evaluation framework** to help us stay informed of what types of queries perform well and which don't for each new release of Scopus AI.
- We also analyze user feedback to identify **poor performing queries, improve the service** and **report to our Sales teams the strengths and limitations of the solution**.

## Principle 5: We respect privacy and champion robust data governance

- As we embed GenAI features in Scopus and other products, we do so in line with Elsevier's Responsible AI Principles and Privacy Principles. To that end, Scopus content selection is **subject to rigorous checks by the Independent Scopus Content Advisory Board**, using both quantitative and qualitative measures.
- For Scopus AI, we are using OpenAI's large language model (LLM) ChatGPT hosted on Microsoft Azure and have an agreement in place that **information passed to this service will not be stored or used for training purposes**. This is an important feature of our implementation, designed to provide privacy and peace of mind to both data publishers and authors.


For more information, visit
elsevier.com/products/scopus/scopus-ai