Responsible AI Principles

Elsevier provides information-based analytics and decision tools for researchers and health professionals worldwide, helping them advance science and improve healthcare outcomes, for the benefit of society.

For more than a decade, Elsevier has been using AI and machine learning technologies responsibly in our products combined with our unparalleled peer-reviewed content, extensive data sets, and sophisticated analytics to help researchers, clinicians and educators discover, advance and apply trusted knowledge.

Our responsible AI principles

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.

Responsible artificial intelligence principles at Elsevier

Our solutions, both internal and external, enhance human decision-making. This approach is underpinned by our commitment to corporate responsibility, which we define as the way we do business, proactively working to increase our positive impact and prevent negative impact. We are using ever more sophisticated analytics and technology as we deliver higher value solutions. While the confluence of increasingly complex technology, proximity to decision making and potential impact on people creates new challenges, they also create many opportunities to benefit society.

Purpose and scope of the responsible AI principles

Generally, we use the term Artificial Intelligence (AI) to describe machine based systems which infer solutions to set tasks and have a degree of autonomy. The scope of these Principles, however, is broader than AI and includes any machine driven insights resulting from the tools and techniques within the field of data science. These Principles provide high level guidance for anyone at Elsevier working on designing, developing, and deploying machine-driven insights.

They provide a risk-based framework drawing on best practice from within our company and other organizations. Individual business areas own the practical implementation of the principles. Elsevier has robust policies and processes in place that are applicable to AI-enabled solutions. The purpose of the Responsible AI Principles is to complement these. AI is a field that evolves continually, at unprecedented speed and scale. These Principles will iterate over time, based on colleague and customer feedback, as well as industry and legislative trends. This will allow us to be proactive, ensuring our solutions develop in line with our values, and maintaining our stance as a thought leader in the market.

1. We consider the real-world impact of our solutions on people

This level of reflection enables us to create trustworthy solutions in line with our company values.

Recognizing that our solutions may assist our customers in their decision-making, we are mindful of the potential impacts our solutions may have on people. AI is a method to solve a business problem for our customers as well as our own company, implying a set of assumptions and a specific, real-world context. The better that context is understood and the more aware we are of our assumptions, the better the solutions we create, and the higher the value-add for customers.

Source: https://www.elsevier.com/about/policies-and-standards/responsible-ai-principles

We go beyond asking “what are we building, and who is the customer?” We seek to identify the range of people who benefit from our solution and how, and who might be impacted and why. To do so, we define the sphere of influence of the solution. We map stakeholders beyond direct customers, and we think about the domain to which the solution applies — are people’s health, livelihood (including career prospects) or rights affected in any way? These insights enable us to consider the impact of a particular solution.

2. We take action to prevent the creation or reinforcement of unfair bias

This drives high-quality results and averts discrimination.

As a supporter of the United Nations Global Compact, promoting fairness and non-discrimination is at the core of our business philosophy and values. We understand that mathematical accuracy doesn’t guarantee freedom from bias, which is why we act to prevent the creation or reinforcement of unfair bias. When such actions are not taken, bias can be introduced inadvertently via data inputs and/or through machine processing or algorithms. Once introduced, it can be replicated through human decision-making across data science, product management and technology.

That can lead to results that are skewed, and therefore less valuable, output. It also may lead to less favorable outcomes for individuals or groups based on gender, ethnicity, socio-economic status, and other personal attributes. Our actions to prevent the creation or reinforcement of unfair bias include implementation of procedures, extensive review and documentation processes, and use of available automated bias detection tools, among others.

3. We can explain how our solutions work

An appropriate level of transparency creates trustworthiness for users and regulatory bodies.

This principle doesn’t prescribe what models to build and use, and it doesn’t prohibit the use of “closed box” models. The aim is to have an appropriate level of transparency for each application and use-case to ensure that different users can understand and trust the output.

Different contexts and audiences require different explanations. As part of the design process, we consider what elements of the solution will need explaining, to whom they might need to be explained, and how to go about explaining them. We also evaluate a solution’s reliability and are explicit about its intended use.

4. We create accountability through human oversight

This enables ongoing, robust quality assurance of machine outputs and helps pre-empt unintended use.

Our technology assists our customers’ decision-making processes. It is important that humans have ownership and accountability over the development, use, and outcomes of AI systems. We apply human oversight throughout the lifecycle of our solutions. This is core to ensuring the quality and appropriate performance of our solutions.

Once the solution has left our hands, this means the customer takes on the role of ultimate decision-maker. Use of our solutions is controlled by an agreed set of terms and conditions, as well as applicable law. We hold our customers accountable to these requirements. Customer support colleagues play an important role in ensuring the intended use is understood by customers and quality issues are dealt with appropriately by internal teams.

5. We respect privacy and champion robust data governance

This ensures we continue to be recognized as a trusted provider of information solutions.

Source: https://www.elsevier.com/about/policies-and-standards/responsible-ai-principles

Appropriate collection, reproducibility, usage and protection of data are crucial to our long-term success as an information and analytics business. As we maintain and broaden our data assets and discover new ways of generating insights, we recognize that strong data governance is essential. AI systems function more accurately when they are fed large amounts of high quality data, and some data sets are utilized across solutions, for multiple purposes. We need to ensure we have in place robust data management and security policies and procedures.

Some data sets include personal information. We are committed to handling personal information in accordance with all applicable privacy laws and regulations as well as our own Privacy Principles, which require that we always act as responsible stewards of personal information.

Source: https://www.elsevier.com/about/policies-and-standards/responsible-ai-principles