Episteme (2026), 1–18
doi:10.1017/epi.2026.10140

episteme

ARTICLE

Navigating AI-Accelerated Science: An Intellectual Virtue Approach to Responsible Inquiry

Vlasta Sikimić

Industrial Engineering & Innovation Sciences, Eindhoven University of Technology, The Netherlands
Email: v.sikimic@tue.nl

(Received 11 April 2026 UTC; revised 29 July 2026 UTC; accepted 8 August 2026 UTC)

Abstract

Although AI can accelerate scientific discovery, it also generates new epistemic risks that prevailing rule-based governance frameworks are structurally ill-equipped to address. We present four arguments for the claim that intellectual virtues offer an adequate response to the challenges of AI-accelerated science. The Adaptability Argument states that rules cannot keep pace with technological change. The Motivation Gap Argument shows that compliance does not amount to epistemic motivation. The Human-in-the-Loop Argument stresses that AI systems require virtuous human interpretation. Finally, the Fallibility-in-Principle Argument maintains that AI systems remain fallible and, therefore, require human oversight.

Moreover, we develop the first systematic mapping of intellectual virtues onto the CRediT (contributor roles taxonomy) stages of the research process and extend this to encompass the process of scientific evaluation. In doing so, we bridge virtue epistemology and the governance of AI-accelerated science. We show how virtues such as intellectual justice, intellectual scrutiny, intellectual humility, curiosity, and intellectual courage constitute the intellectual infrastructure that enables scientists to evaluate AI-generated outputs and preserve epistemic agency in an era of increasing outsourcing of cognitive effort. This virtue-based framework offers a proactive and future-oriented normative foundation for scientific practice in the age of AI.

Keywords: AI in science; virtue epistemology; intellectual justice; intellectual scrutiny; epistemic agency; responsible AI; CRediT; curiosity

1. Introduction

The rapid increase of the use of artificial intelligence in science has been well documented. In 2023, almost eight per cent of the papers in the Scopus database mention AI or machine learning terms in their title or abstract (Van Noorden and Perkel 2023), and, depending on the subject, up to one-fifth of the papers published between 2020 and

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

2024 show signs of AI use (Liang *et al.* 2025). This trend extends to science evaluation, for instance around 21% of the reviews for the International Conference on Learning Representations 2026 were reported to be fully AI-generated (Naddaf 2025).

AI not only speeds up the scientific discovery process, but it can also endanger the reliability of the scientific pursuit. For example, AI can hallucinate wrong claims it never saw (Ji *et al.* 2023) and reproduce biases (Ananya 2024). However, it has the potential to improve meta-studies and literature reviews. AI can process more sources than humans, it can crosscheck for claims in papers, use sources in many languages, and even adjust the output to different target audiences (Glickman and Zhang 2024). Thus, the question is how to optimize its use in science. To address this question, we propose a normative intellectual virtue framework for maximizing the benefits and minimizing the epistemic damage.[1] More precisely, we study how intellectual virtues help scientists in the quest for knowledge in the AI era.

The intellectual virtue theory provides an ideal framework for addressing the challenges associated with AI-accelerated science because it focuses on human capacities. In the scientific context, intellectual virtues such as curiosity, open-mindedness, and critical thinking help with managing misinformation and finding reliable sources. For example, scientists using AI should be curious about unexpected results, open to new methods, and critical of data sources.

The ethical virtue framework can be traced back to Aristotle (Aristotle and Crisp 2014) who defined it as a middle ground between two vices. Intellectual virtue theory, as a newer framework, is helping to understand the learning process. Exercising intellectual virtues leads to more justified beliefs and genuine knowledge (e.g., Code 1984; Montmarquet 1992). It is gaining prominence in epistemology, positive psychology, and educational studies. For example, Krumrei-Mancuso *et al.* (2020) showed how intellectual humility can benefit the learning process, and Orona and Pritchard (2022) investigated ways how teaching intellectual virtues increases curiosity. Our framework is situated primarily within the responsibilist tradition in virtue epistemology (Code 1987; Montmarquet 1992), which holds that intellectual virtues are acquired character traits involving appropriate motivation and effort, rather than mere innate cognitive faculties. This responsibilist orientation is significant because it entails that virtues can be developed and strengthened through deliberate practice (Baehr 2013; Krumrei-Mancuso *et al.* 2020; Orona and Pritchard 2022). The trainability of intellectual virtues is what makes a virtue-based framework particularly suitable for an era of rapid technological change.

Our contribution is threefold. First, we develop four arguments for the necessity of intellectual virtues in AI-accelerated science: the Adaptability Argument, the Motivation Gap Argument, the Human-in-the-Loop Argument, and the more general Fallibility-in-Principle Argument (Section 2). Second, we bridge virtue epistemology and the governance of AI-accelerated science by mapping intellectual virtues onto the ethical principles of responsible AI (Section 3). Third, we provide the first systematic mapping of intellectual virtues onto CRediT (contributor roles taxonomy) research stages and scientific evaluation, showing how different virtues matter at various stages of the scientific discovery process (Section 4). We discuss the broader implications of our framework for epistemic training and the preservation of human epistemic agency in Section 5.

[1]We use 'intellectual' to encompass both the epistemic and cognitive dimensions.

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

2. Why is virtue epistemology needed for AI in science?

2.1. Four arguments for the necessity of intellectual virtues

We contend that rules and intellectual virtues are both useful for the governance of AI-accelerated science, but that they serve different purposes and operate on different levels. Rules codify shared standards and enable institutional oversight, whereas intellectual virtues are dispositions that also guide epistemic judgement in situations that rules do not yet cover. Our aim in this section is not to settle the long-standing debate between rule- and virtue-based approaches. Rather, we motivate the claim that intellectual virtues are needed for epistemically responsible AI-accelerated science. We support this claim through four interlocking arguments.

2.1.1. The Adaptability Argument

Premise 1: AI technologies evolve faster than regulatory systems.

Premise 2: Rule-based frameworks depend on the prior identification of specific epistemic harms.

Conclusion: Rules alone are insufficient to ensure epistemically responsible AI use.

Intellectual virtues, by contrast, provide adaptable epistemic judgement that can be exercised across novel and unanticipated contexts. The two mechanisms are complementary rather than competitive. Rules retain their function of consolidating the lessons learned, while intellectual virtues also carry the epistemic load during the lag, when the harms are not yet understood. Although some lag may be epistemically valuable because it allows harms to become visible enough for appropriate rules to be formulated, intellectual virtues remain necessary during that interval to guide scientists in anticipating and mitigating potential harms.

In contrast to regulations that frequently change, intellectual virtues and vices are more stable traits, which are also interconnected with human cognition and thinking styles (Krumrei-Mancuso *et al.* 2020). Nevertheless, some virtues and vices become more salient in certain contexts and we can train them during education (e.g., Orona and Pritchard 2022). For example, the disposition to critically examine AI outputs is a specification of intellectual scrutiny, which is violated when we surrender our epistemic effort to an algorithm.

2.1.2. The Motivation Gap Argument

Premise 1: Scientific inquiry requires epistemic motivations, such as curiosity and wonder, which guide what researchers choose to explore.

Premise 2: Rule-based frameworks and AI systems can regulate behaviour and suggest directions, but they cannot generate or appropriately guide epistemic motivation.

Premise 3: Intellectual virtues, especially curiosity and intellectual autonomy, are part of the epistemic motivation of researchers. They are salient in the AI-accelerated science as they, for example, help them retain control over the direction of inquiry.

Conclusion: Intellectual virtues are necessary for responsible AI-accelerated science because they secure the motivation for inquiry that rule-based frameworks and AI systems cannot provide.

Intellectual virtues ensure that researchers do not merely follow externally imposed rules or algorithmically suggested paths. They actively help in deciding what is worth investigating. In doing so, intellectual virtues preserve human epistemic agency.

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

2.1.3. The Human-in-the-Loop Argument

Premise 1: Sub-symbolic AI systems[2] used in science are epistemically opaque; their internal reasoning processes cannot be fully inspected or understood by human users.

Premise 2: Opaque systems require human epistemic interpretation; scientists must judge when to trust, question, or override AI outputs without full access to the system's reasoning.

Premise 3: Such interpretation is not a mechanical process; it requires intellectual character traits such as intellectual scrutiny, intellectual courage, and intellectual humility.

Conclusion: Epistemically responsible use of sub-symbolic AI depends on the intellectual virtues of the human agents who interpret, evaluate, and act upon AI-generated outputs.

This connects directly to the problem of epistemic surrender,[3] which is the tendency to defer uncritically to AI recommendations, thereby ceding one's epistemic agency to an opaque system. Rezazade Mehrizi *et al.* (2023) documented such a phenomenon among radiologists, who followed incorrect AI suggestions 72% of the time. The Human-in-the-Loop Argument shows that the proper response to the tendency of surrendering epistemic effort is not more rules about when to trust AI, but the cultivation of intellectual character that sustains independent epistemic judgement.

2.1.4. The Fallibility-in-Principle Argument

This argument extends the Human-in-the-Loop Argument from the specific case of opaque systems to the general principle that all AI in science requires virtuous human oversight.

Premise 1: Even the most advanced AI systems are fallible in principle. They can produce errors, or have epistemic limitations, for example, consider only a limited number of hypotheses.

Premise 2: Even if AI systems were to outperform human scientists on average across a wide range of tasks, their epistemic fallibility in any particular case would remain.

Conclusion: Human epistemic oversight is indispensable because an independent epistemic check reduces the probability of undetected error.

For human oversight to be meaningful rather than ceremonial, the human in the loop must possess the intellectual virtues necessary to exercise critical judgement. Therefore, the scientists need intellectual scrutiny, courage, and autonomy to serve as a genuine epistemic safeguard not as a mere rubber stamp.

Our framework treats responsible AI use not as a matter of rule-following but as a matter of good intellectual agency. This shift in framing from compliance to character is significant. It means that the question is not merely 'Are scientists following the right rules?' but 'Are scientists the kind of epistemic agents who can be trusted to use AI well?' The virtue-based approach prioritizes the latter question and thereby provides a more robust foundation for responsible AI governance in science.

[2]Sub-symbolic AI, such as large language models, learns statistical patterns from data through numerical representations rather than explicit rules.

[3]On epistemic surrender, see Nguyen (2022) on the concept of 'cognitive outsourcing' and its risks for intellectual autonomy. We extend this concern to the scientific context specifically, but make a moderate claim about surrendering epistemic effort and not full cognitive outsourcing as this is a much stronger position.

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

By emphasizing researchers' character and judgement, the virtue-epistemology framework offers a complementary perspective to other approaches in the epistemology of AI, which have focused primarily on ethics (Hagendorff 2022), explainability (Sullivan 2022; Zednik 2021), or the ethical virtuousness of AI solutions themselves (Farina *et al.* 2024). Although these approaches provide important normative constraints, they tend to frame scientists as compliance agents rather than as epistemic agents exercising judgement under uncertainty.

We aim to fill this gap by showing how virtuous epistemic conduct can underpin the responsible and effective use of AI in scientific research. This emphasis on virtues is not meant to replace regulatory or technical solutions, but to ground them in the qualities of scientists themselves (Vallor 2016). By focusing on virtues, we not only respond to today's challenges but also prepare scientists to navigate the unforeseen challenges of tomorrow with sound judgement and integrity.

3. Intellectual virtues corresponding to core ethical principles of responsible AI

Hagendorff (2022) analysed the ethical virtues corresponding to the four frequently cited principles of responsible AI, namely to the principles of fairness, transparency, accountability, and privacy. He connects them to the virtues of justice, honesty, responsibility, and care. These ethical virtues and principles have their epistemic counterparts. For example, the virtue of intellectual justice corresponds to the fairness principle. Intellectual justice, as we understand it here, is the disposition to include the perspective of diverse knowers, particularly those structurally disadvantaged (Anderson 2012; Fricker 2007). Its corresponding vice is intellectual injustice: ungrounded intellectual exclusion of speakers from the underprivileged. In the context of AI, intellectual justice is especially salient because training data bias can reproduce and amplify exactly these forms of injustice at scale. The core epistemic function of intellectual justice is the preservation of knowledge diversity and ensuring that the epistemic contributions of under-represented communities are neither lost nor distorted by algorithmic mediation. In the following, we discuss the intellectual counterparts to Hagendorff's ethical framework in detail. In Section 4, we further develop the intellectual virtue framework by mapping key intellectual virtues onto the stages of the scientific process where AI can contribute.

Although intellectual virtues are often attributed to individuals, they can also characterize scientific communities as collective epistemic agents (Anderson 2012; Lahroodi 2007). For example, a scientific community can be intellectually just when it systematically includes disadvantaged perspectives and intellectually responsible when it establishes procedures for scrutinizing evidence and correcting collective error. In AI-assisted science, such group virtues are especially important because epistemic harms often arise from shared algorithms and datasets. As I have argued elsewhere, shared algorithms can create algorithmic monocultures that homogenize scientific judgements and amplify errors across the community (Sikimić 2025).

Intellectual honesty (Byerly 2023) and intellectual humility (Krumrei-Mancuso and Rouse 2016) guide our need for transparency in research. These virtues help scientists properly assess the limits of their results and motivate them to strive for their explainability. Intellectual humility refers to recognizing our own limitations in learning, including the fallibility of our senses. In the context of AI, it stands for the ability to use the machine with critical reasoning. Intellectual honesty facilitates the transparency of research. It is particularly important when it comes to the use of AI because many AI

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

Table 1. How intellectual virtues correspond to ethical principles (extended on Hagendorff 2022).

Ethical principle	Ethical virtue	Intellectual virtue
Fairness	Justice	Intellectual justice
Transparency	Honesty	Intellectual honesty
Accountability	Responsibility	(Group) intellectual responsibility
Privacy	Care	Curiosity

models used in science are epistemically opaque (Humphreys 2009). The opacity concerns the models themselves rather than their individual outputs, and it comes in different kinds with different sources (Alvarado 2021). Moreover, the problem is not merely that explanations of these models are hard to obtain. As Boge (2022) argues, the difficulty is structural. We do not know how the training process arrives at the final parameter values, we do not know what the individual parameters mean, and there is no established way to find out, so we do not know what knowledge, if any, the model contains. The sheer number of parameters, such as the 175 billion parameters of the large language model (LLM) GPT-3 (Brown et al. 2020), is therefore not the source of the problem by itself, but it illustrates the scale at which this structural opacity arises.

The principle of accountability and the ethical virtue of responsibility have the direct intellectual counterpart in intellectual responsibility (Montmarquet 1992) and intellectual conscientiousness (Montmarquet 1987). A person who exercises intellectual responsibility only accepts beliefs that are supported by evidence, for example AI-generated outputs with real references and sufficient explanation. An intellectually conscientious person ensures that they gather sufficient evidence before forming judgement. These virtues have a positive impact on science. A model of scientific communities demonstrates that collecting sufficient evidence prior to advocating for a hypothesis significantly improves scientific accuracy of findings in comparison with communicating preliminary results immediately (Sikimić and Herud-Sikimić 2022).

Curiosity clusters virtues such as curiosity (Whitcomb 2010), inquisitiveness (Watson 2015) and love of knowledge (Roberts and Wood 2007). The obvious related vice is lack of interest. More interestingly, curiosity and love of knowledge also involve focusing on necessary and fruitful information rather than on epistemic junk, such as gossip. In other words, the other vice related to curiosity is being interested and spending cognitive effort on irrelevant and maybe even epistemically harmful content. Thus, curiosity is of exceptional importance for science. It is both linked to the motivation of researchers, and it helps in choosing which hypotheses to explore. In the context of AI, it has to do with data collections and trustworthy research. Collecting or wrongly interpreting data, violating the privacy of participants, etc. are on the antipode of curiosity similar to reading gossip newspapers. Finally, in the misinformation age, which is accelerated with AI technologies, self-regulation and proper direction in the domain of curiosity are becoming a virtue of the future. In Table 1, we summarize these insights.

4. Intellectual virtues across the stages of AI-supported research

As AI reshapes scientific practice, researchers must cultivate intellectual virtues that enable its effective use. The relevance of particular virtues varies across stages of the

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

scientific process, including research design and execution, communication of results, peer review, and grant evaluation. We propose a framework of intellectual virtues that are especially important for engaging with AI at these different stages of scientific work. To structure this proposal, we draw on, and go beyond, the CRediT nomenclature (Allen *et al.* 2014), which offers a widely accepted framework for distinguishing different types of scientific contributions beyond the traditional focus on authorship alone. This mapping constitutes our primary constructive contribution; it provides a structured, stage-specific normative tool that can guide intellectual virtue cultivation in practice.

CRediT defines specific roles for the authors of a manuscript, such as conceptualization, data curation, formal analysis, and supervision. Please note that some parts of the CRediT nomenclature are condensed for better overview. In addition, we include scientific tasks that fall outside the authoring of papers, such as peer review and dissemination activities. These activities play a central role in maintaining the epistemic quality of science, for example by filtering unreliable results. This framework allows us to connect intellectual virtues with specific scientific practices, while also highlighting forms of epistemic labour that are essential for collective knowledge production in science but remain under-recognized.

The intellectual virtues that are important at various stages of the discovery process can differ depending on whether these stages are carried out by human researchers or delegated to AI systems. Consider, for example, the CRediT item 'Visualization'. Human researchers need to be intellectually honest while they are visualizing data and while promoting AI to plot data. They may be tempted to exaggerate their results, for instance by choosing plot axes that make differences between biological populations appear larger than they are, or by selecting so-called representative images that in fact depict extreme cases. When researchers use AI, they should prompt it to visualize the data as it is, rather than to produce visualizations designed to support a favoured hypothesis. When data visualization is handled by AI, an additional risk arises. AI systems may generate visualizations that include fabricated or distorted data. In this case, the scientist working with AI must exercise intellectual responsibility and intellectual justice. This requires awareness of potential biases and errors introduced by AI systems and a willingness to take responsibility for the epistemic quality of the visualization, including critically questioning and correcting it when necessary.

At the same time, we acknowledge that there are limits to what individual scientists can evaluate. Some models resist understanding in principle, and many algorithms are proprietary and cannot be inspected at all. In such cases, intellectual humility and intellectual honesty call for second-order strategies. For example, instead of asking an LLM to visualize the data directly, virtuous scientists can ask for the code to visualize it themselves with tools suitable to visualize data such as R.

In Table 2 we summarize the salient intellectual virtues for the different parts of the discovery process. They are discussed in detail in the following sections.

4.1. Conceptualization

Conceptualization refers to the formulation of overarching research goals and the formulation of hypotheses. This is one of the most impactful parts of the discovery process. While using AI for this initial stage, we mainly need to foster two intellectual virtues, namely intellectual justice and intellectual autonomy.

Intellectual justice ensures that diverse perspectives, approaches, and research questions are considered. For example, while working on diseases, it is important to

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

Table 2. The salient intellectual virtues for the use of AI for different parts of the scientific pursuit.

Scientific process	Key virtue	Example
Conceptualization	Intellectual justice	Choose research goals with impact on diverse populations in mind.
	Intellectual autonomy	Keep leading the research path, not outsourcing intellectual burden to machines.
Methodology	Intellectual justice	Choose a methodology beneficial for everybody.
	Open-mindedness	Be open to new methods instead of relying on an established toolbox.
Investigation and data	Curiosity	Respect privacy and investigate unexpected patterns.
	Intellectual justice	Be mindful of the purpose of original data collection when reusing data.
Analysis and interpretation	Intellectual honesty	Be transparent about modelling choices influencing AI output.
	Intellectual humility	Ensure transparent human reasoning when AI reasoning is opaque.
Writing	Intellectual responsibility	Take responsibility for the claims in the study.
	Intellectual honesty	Correct overselling of research results by AI.
Peer review	Intellectual scrutiny	Check sources and evaluate if results warrant the conclusion.
Funding allocation	Collective intellectual responsibility	Ensure no undetected error systematically impacts the community.
	Intellectual justice	Address biases in training data; monitor impact on marginalized groups.
	Intellectual courage	Support bold ideas and new approaches.
Replication and meta-studies	Intellectual openness	Be open to new technologies and replication of findings.
	Intellectual humility	Recognize that failure to find corroborating evidence does not disprove a claim.
Science communication	Intellectual honesty	Be attentive when AI systems oversell results.

consider the burden this disease puts on the population and take the position of the patients into account. In contrast, while using AI to choose research goals carelessly, there is the risk that the goals are aligned with the most active fields of research because they are over-represented in the training data. For example, currently studies on many diseases, predominantly or only affecting women, such as uterine cancer, are underfunded when considering the disease burden (Smith 2023). Because AI tends to reinforce dominant paradigms, following AI suggestions can reduce the diversity of questions asked and, therefore, harm scientific progress. Intellectually just scientists, thus, need to be on alert when AI is used for defining research goals. Science needs diverse hypotheses and approaches. Diverse approaches especially help address difficult questions (e.g., Kitcher 1990; Zollman 2010). Because AI tends to replicate what it saw previously, it usually does not provide out-of-the-box approaches when not explicitly asked for it. However, intellectual justice enables the researcher to include evidence from different sources.

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

Additionally, intellectual autonomy is important for ensuring that researchers remain in the driver's seat of the scientific enterprise. It enables scientists using AI to choose their own research direction, reach their own conclusions, and explore their own hypotheses despite convincing sounding explanations from an algorithm. This is connected to scientific integrity because only humans can bear responsibility for it, and, therefore, they must remain responsible for and in charge of AI solutions. Furthermore, intellectual autonomy affects the well-being of researchers because it ensures a sense of ownership. An over-reliance on AI would reduce the intellectual autonomy of researchers, especially if the reasonings of the machine are not transparent. For example, an intellectual autonomous researcher would include information it got from an AI to refine their own hypothesis, but they would not leave the hypothesis formulation itself to a machine. The risk of ceding intellectual autonomy connects directly to the Fallibility-in-Principle Argument; even when AI generates plausible hypotheses, the scientist's independent judgement remains the indispensable safeguard against systematic error in the conceptualization stage.

4.2. Methodology

For a given scientific question, the most appropriate methods should be chosen. This choice should be driven by different intellectual virtues. Again, intellectual justice plays a key role. Intellectual justice should ensure that a methodology is chosen with the impact on different populations in mind. For example, in medical science, Cortina *et al.* (2025) noticed that data filters commonly used for selecting electronic health records disproportionately affect minorities. When filter criteria are chosen without the impact on underprivileged populations in mind, these populations can be excluded from the training data, and hence the AI can perform worse for them. A scientist who is intellectually just will always monitor the impact of their choices on different populations.

Furthermore, AI is reproducing data from the training dataset. Hence, AI will in most cases suggest the methods used in the past to address research questions. The virtue of open-mindedness helps researchers to consider innovative methods that could better address a new research question, rather than relying on the same methods they have used in the past. Finally, open-mindedness helps researchers profiting from the advantages from AI. For example, they can employ AI tools which can help choose the best statistical analysis for a task (DeBruine and Lakens 2025).

4.3. Investigation, data collection, and curation

Investigation covers the performance of experiments and the collection of data, whereas data curation is concerned with annotating these data and making them available for use and reuse. The main intellectual virtue important for data collection with AI is curiosity. Curiosity helps researchers to address meaningful research questions and to investigate unexpected AI outputs. The corresponding vices related to curiosity are indifference towards what is new or unplanned and a disregard for privacy. In particular, the latter is frequently challenged by AI aided research. Van Noorden (2020) found that a significant minority of facial recognition researchers held surprisingly lax attitudes towards data ethics. Curiosity also helps researchers to investigate why particular data are collected and what data might be missing from an AI collected dataset.

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

Additionally, data collection and curation are strongly connected to intellectual justice. Data collection and curation are performed with a specific purpose in mind (Leonelli 2016, 2019). This purpose is not necessarily aligned with the subject of another academic study. Many researchers rely on datasets that are known to be biased (Paullada *et al.* 2021). For example, Crawford (2021) discusses how data used for training AI are disproportionately sourced from the Global North, thereby shaping the epistemic and cultural assumptions embedded in these models. In contrast, an intellectually just researcher would go the extra mile and collect data that are most meaningful for the study in a responsible way.

4.4. Analysis and interpretation of the results

Scientific analysis is the evaluation of data using formal techniques, such as statistics. The main virtues important for scientists using AI for this task are intellectual honesty and intellectual humility. Intellectual honesty and humility help to ensure transparency and explainability by honestly reporting of the inputs, outputs, the algorithm, and modelling choices. Intellectual humility requires prudence with the AI-generated results because of the black-box challenge.

Using AI and machine learning involves certain modelling decisions, such as how to deal with missing data and how to define sufficient data quality. These decisions influence the results produced by AI systems and therefore need to be reported honestly. In addition, analyses conducted by AI need to be explainable. This is one of the main normative guidelines governing the use of AI (e.g., Zednik 2021). Furthermore, intellectual honesty is essential for scientists using AI for formal analysis because they need to include all collected data, not only the data that support their hypothesis. The use of AI, for example for pattern recognition, can make it more difficult to detect whether all data were used in an analysis, whereas classical statistical tools follow more established rules and are often easier for reviewers and readers to evaluate.

Although intellectual honesty is required to analyse data in the most appropriate way, intellectual humility helps researchers to acknowledge the limits of their own understanding of AI-generated outputs. Some AI tools support scientists in following best practices for data analysis (e.g., DeBruine and Lakens 2025), and an intellectually humble researcher is willing to follow such guidances. Moreover, when AI tools are used to identify patterns in data, intellectual humility helps researchers to admit when they do not understand the model's reasoning and to investigate this reasoning either independently or in collaboration with appropriate experts, thereby increasing the transparency of the decision-making process. The Human-in-the-Loop Argument is directly operative here; when the AI's reasoning is opaque, the scientist's intellectual character, not a checklist, determines whether the resulting interpretation is intellectually responsible.

4.5. Writing

Writing covers the drafting and editing of the manuscript as well as the presentation of the study. The key virtues for scholarly communication are intellectual responsibility and intellectual honesty.

Although an AI system can produce a convincing sounding argument when prompted with scientific results, the human in the loop needs to exercise intellectual responsibility. Scientists remain accountable for the work and must carefully evaluate

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

the evidence to understand the implications and limits of the study. Intellectual responsibility cannot be outsourced to an algorithm or an external AI provider. This view is reflected in many journal policies, which do not allow generative AI systems to be listed as authors because they cannot be held responsible for the content of a study.

A second important aspect is intellectual honesty. It is essential for scholarly communication because claims that are not sufficiently supported by evidence risk polluting the scientific literature. For example, if an AI system extrapolates results and produces claims that are not supported by the available evidence. It can even hallucinate convincingly sounding references that explain any given behaviour (Ji *et al.* 2023). Hence, the scientist must be intellectually honest and correct the AI whenever necessary. Only when the limitations of a study are honestly reported can readers properly evaluate the relevance of the findings for their own research questions. Sikimić and Herud-Sikimić (2022) use agent-based models to illustrate how honest communication of evidence helps scientific communities reach the optimal conclusions more quickly. Their results show that intellectual honesty is an important virtue for collective knowledge acquisition.

4.6. Peer review

Peer review is currently the main method of scientific evaluation. Usually, peer review defines which project is funded and which paper is published. Studies have shown that reviewers are already turning to AI to assist them in the review of manuscripts, even when they are not supposed to (Naddaf 2025). However, automatizing part of the peer review process can be very attractive because of the high workload associated with it. Studies show a very low inner consistency of the reviewers (Bornmann *et al.* 2010).

If carefully designed, AI can be less biased than human reviewers (Sikimić 2025). Therefore, some organizations aim to profit from its potential. For example, the Association for the Advancement of Artificial Intelligence started a pilot programme to include an additional AI-generated commentary while evaluating abstract submissions for its conference (Association for the Advancement of Artificial Intelligence 2025). The AI evaluation is one of the three, where the other two are given by humans. However, if the peer review process would be more automated, it is essential to allow a rebuttal from the authors which should be judged by human reviewers or review boards because of the persistent concerns about the transparency of AI-based recommendations.

Although researchers should remain open to new ideas and new ways of generating evidence, the produced evidence must be reliable. A researcher exercising intellectual scrutiny asks whether data collection was appropriate, whether analytical choices were adequate, and whether the conclusions drawn from the results are justified. Intellectual scrutiny, as we understand it, is closely related to but not identical with critical thinking. It is the cultivated disposition to subject evidence, reasoning, and sources to rigorous evaluation, what classical epistemology also calls the virtue of intellectual carefulness, now exercised in the specific context of evaluating algorithmically mediated claims. This practice helps to distinguish well-supported hypotheses from mere speculation and is essential in peer review for maintaining the integrity of the scientific record. Intellectual scrutiny is essential, given the low transparency of many AI applications. Because the internal processes of many AI tools are opaque, the responsibility falls on the researchers using these tools to carefully evaluate their outputs. In particular, decisions about which data are included in the training of a model have a substantial influence on the results it

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

produces. This requirement applies not only to generative AI but also to more traditional forms of machine learning.

4.7. Funding allocation

The Spanish La Caixa Foundation already uses AI as a pre-screening method to identify applications with a low chance of success, in an attempt to reduce the workload for reviewers (Adam 2025). Because funding allocation largely relies on the peer review of research ideas and researcher profiles, similar intellectual virtues play a role. In addition, the virtue of collective intellectual responsibility is particularly salient. Collective intellectual responsibility is an intellectual virtue that applies not only to individual agents but also to scientific communities and institutions. It consists in a shared commitment to upholding epistemic standards.

This collective virtue can be challenged by AI because of its scalability. For example, a single biased reviewer can only affect a limited number of applications from a small number of funding agencies. In contrast, AI-based solutions could systematically discriminate against applications with specific features. This risk is particularly high when the same AI system or training dataset is used across different funding bodies. Hence, the use of AI in grant review needs to be implemented with particular care (Sikimić 2025). To practise collective intellectual responsibility, funding agencies and everybody involved therefore need to monitor algorithms and datasets, allow for rebuttals, and ideally combine AI assessment with human review. For example, the La Caixa Foundation assigned two reviewers to recheck applications that are filtered out by the algorithm (Adam 2025).

In addition, the use of AI for funding allocation should be strongly guided by the virtue of intellectual justice. Algorithmic systems can reproduce and amplify biases that are difficult to detect. For example, generative AI systems exhibit gender bias when asked to write letters of reference for male and female applicants (Kaplan *et al.* 2024). To counteract such effects, those using AI for funding allocation must practise intellectual justice and carefully examine how these systems affect marginalized populations.

Finally, AI can discourage novel ideas, particularly when it is used for research evaluation, as it performs best when assessing work that is close to its training data. Radically new ideas may fall outside the scope of AI-based evaluation and therefore risk not receiving the necessary funding (Sikimić 2025). Consequently, those involved in review processes that use AI should practise intellectual courage and remain willing to support bold and unconventional ideas, especially within funding schemes designed to promote innovation. However, humans are often inclined to outsource intellectual work to algorithms. As a result, researchers may follow AI recommendations even when these conflict with their own judgements. For example, Rezazade Mehrizi *et al.* (2023) showed that radiologists largely followed AI suggestions, leading to incorrect decisions in 72% of cases when the system was primed with a wrong recommendation, compared with 78% correct decisions when primed with a correct one. This finding suggests that users often fail to sufficiently practise intellectual courage and instead follow AI in a relatively uncritical manner. However, Van der Sijs *et al.* (2006) showed a high tendency of clinicians to override computational drug safety alerts. To balance these tendencies, researchers need to exercise intellectual courage while working with AI systems. This includes the willingness to challenge, override, or reject algorithmic outputs when they conflict with well-grounded scientific knowledge, empirical evidence, or sound methodological reasoning but not completely ignoring it for other reasons.

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

4.8. Replication and meta-studies

One of the challenges that generative AI brings to science is its ability to generate convincing sounding arguments and to hallucinate information it has never encountered. Hence, replication efforts aimed at independently validating important claims are becoming increasingly important. At the same time, modern technology can also facilitate replication. For example, AI systems can extract and standardize methods sections from published papers, and robotic platforms such as the mobile robotic chemist (Burger *et al.* 2020) can be used to automatically attempt replications of experimental procedures. To use AI in the best possible way, researchers should be intellectually open and intellectually humble.

Intellectual openness is required to adopt novel tools such as AI for scientific inquiry. Moreover, this virtue helps scientists to facilitate the replication of their results, for example by sharing data, materials, and protocols. Replication should not be understood as an attack on one's work but rather as recognition that the work is sufficiently important to warrant independent verification. In addition, intellectual openness enables researchers to consider new methods proposed by algorithmic systems, especially when these methods promise to yield the same results if the underlying hypothesis is correct.

One of the strengths of generative AI is its ability to synthesize written text from multiple sources. This capability is particularly valuable for replication studies and meta-studies. Although replication studies aim to reproduce previous findings to strengthen the robustness of scientific claims, meta-studies aggregate evidence from many different studies to support or challenge broader conclusions. In both cases, scientists using AI must practise the virtue of intellectual humility. This virtue allows researchers to recognize the limits of their own knowledge and abilities. For example, although replication is often considered the gold standard in biology, it is far from straightforward. In particular, when advanced methods are involved, researchers may lack the technical skills required to replicate a finding, and a failure to replicate in such cases does not by itself warrant refuting the original hypothesis. Similarly, if an AI system fails to identify supporting studies for a given claim, this does not mean that the claim should be rejected. AI systems do not have exhaustive knowledge, and researchers using AI for replication should interpret AI outputs with intellectual humility.

4.9. Science communication

AI's ability to summarize data and to generate engaging social media content about research can help scientists communicate their findings and reach audiences beyond the academic community. However, effective communication is not achieved automatically. Researchers still need to develop communication skills to disseminate their work effectively. When AI is used to support the dissemination, the virtue of intellectual honesty is particularly important.

Intellectual honesty requires acknowledging uncertainty, reporting limitations, and being transparent about methodological choices. A lack of intellectual honesty can erode public trust and harm researchers and scientific institutions in the long term. If not carefully guided, AI systems tend to oversell findings or present speculative results as established facts, thereby challenging this virtue. At the same time, AI-based tools such as automated translation systems can facilitate communication with non-academic audiences and help overcome language barriers (Vučković & Sikimić, 2023). With

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

appropriate oversight and adherence to intellectual honesty, AI can support responsible science communication and contribute to increased trust in scientific research.

5. Implications for intellectual training of researchers and the future of scientific virtues

Our framework has two major implications that merit separate discussion: the feasibility of virtue cultivation and the necessary shift in the virtue profile of future scientists.

5.1. Virtue cultivation and epistemic training

On the responsibilist account we endorse, intellectual virtues are not fixed personality traits, but cultivable dispositions that can be developed, strengthened, and refined through targeted interventions. This claim finds support both in philosophical argument and empirical evidence. Philosophically, Baehr (2013) has proposed a pedagogical framework for intellectual virtue education. Empirically, Orona and Pritchard (2022) demonstrated that curiosity can be increased through a structured educational intervention, and Krumrei-Mancuso *et al.* (2020) showed that intellectual humility is positively correlated with knowledge acquisition and can be enhanced through reflective practice. These findings from psychology corroborate the responsibilist thesis: virtues are learnable.

This has direct implications for AI governance. If intellectual virtues can be trained, then institutional investment in intellectual character formation, through graduate curricula, professional development programs, and research ethics training, becomes a viable and indeed necessary component of responsible AI governance. Our CRediT-mapped framework (Table 2) can serve as the basis for such training programmes; researchers can be trained in the specific virtues most relevant to their stage of the research process. For instance, a workshop on AI-assisted data analysis could focus on cultivating intellectual humility and honesty, while training for AI-assisted peer review could emphasize intellectual scrutiny and courage.

5.2. The shifting virtue profile of future scientists

AI will transform not only what scientists do but also who they need to be as epistemic agents. As AI systems take over routine cognitive tasks such as literature search, statistical analysis, code generation, and even initial hypothesis formulation, the distinctive epistemic contribution of human scientists will increasingly reside in the intellectual capacities that cannot be automated, such as critical judgement, interpretive creativity, intellectual courage, and the ability to recognize the limits of algorithmic reasoning. Arguable, future scientists will be more prolific as AI acceleration will enable them to produce more research outputs more quickly, but the epistemic quality of their work will depend more heavily on their intellectual character than ever before.

This suggests a reevaluation of the virtue profile for scientific training. Although traditional scientific education emphasized technical skills and domain knowledge, the AI era demands equal investment in the cultivation of intellectual virtues. Our framework provides the normative structure for this reevaluation. By mapping virtues to research stages, it identifies where and how the scientific community should invest in developing these capacities.

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

6. Conclusion

AI can help researchers complete routine tasks efficiently, it provides unprecedented access to knowledge, including automatic translations of scientific literature from around the globe, and enables them to find patterns in large datasets. For example, in 2024 the developers of the AI-based protein structure prediction tool AlphaFold were awarded the Nobel Prize in Chemistry. Given these capabilities, AI has become a valuable asset for many scientists in a remarkably short time.

The skills scientists need in the era of AI differ from those required in earlier periods. For example, widely used LLMs allow many researchers to write efficient code with minimal programming background, reducing the need to invest heavily in learning such technical skills. Similarly, automatic translation tools enable access to scientific literature from diverse linguistic contexts without the long process of language acquisition. At the same time, a basic understanding of AI systems, and of LLMs in particular, has become essential for scientific practice.

Alongside these changes in skills, certain intellectual virtues are becoming salient in the age of AI. In particular, intellectual scrutiny and intellectual justice have become more important than ever. Although human researchers typically make deliberate choices about which sources they consult, LLMs often present information without clear references, making the validation of claims more difficult. A scientist who exercises intellectual scrutiny, therefore, goes beyond surface-level outputs by identifying relevant sources and verifying claims against original data. Moreover, AI systems are commonly trained on historical data that may reflect biases or stereotypes. For this reason, the virtue of intellectual justice is crucial to ensure that diverse perspectives are taken into account, and that the scientific process remains value driven, rather than becoming a self-reinforcing feedback loop based on biased data.

Given the rapid pace of technological change, it is essential that scholars are prepared to engage with new tools in a virtuous manner, guided by intellectual virtues rather than relying solely on external rules or technical safeguards. Intellectual virtues can be cultivated through targeted interventions and strengthening critical thinking.

The intellectual virtues of the scientists using AI should be considered as one part of the puzzle to strive for a better science in the age of AI. For the most holistic perspective, we need to consider the virtues of AI developers and the collective virtues of institutions as well. The virtues most important for the developers of AI for science differ from those of scientists using AI for research. Some intellectual virtues are important for both users and developers. For example, intellectual justice ensures that the training data represent diverse perspectives and do not stay focused on the Global North, and that we intellectually mitigate algorithmic biases. However, virtues such as foresight are more salient for developers because they should consider how their technology could develop in future. From the perspective of institutions, the effect of AI on the scientific community should be addressed, in particular, considering the potential of AI monocultures and monopolies.[4] This represents a fruitful avenue for future research.

The ease with which AI systems can overstate results or even generate entirely false claims increases the complexity of peer review and heightens the need for replication efforts and meta-studies. Alongside the cultivation of intellectual virtues in individual researchers, systematic institutional measures are, therefore, required to maintain the

[4]AI monocultures arise when the same algorithm or training dataset is used by (almost) all users. Monopolies can be the consequence of market consolidation caused by the very high investments needed for some AI tools.

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

reliability of the academic literature. Luckily, AI can also support these efforts. It is especially well suited for synthesizing large bodies of research into meta-studies, checking statistical analyses, and automating replication studies on a large scale, which are increasingly feasible thanks to advances in AI-driven solutions.

Data availability statement. No data were collected for this theoretical work.

Acknowledgements. I would like to thank Krist Vaesen for reading and commenting on the manuscript and for his encouragement. I would also like to express my deep gratitude to Ole Herud-Sikimić for his generous and sustained support of my work. Finally, I am grateful to the anonymous reviewers for their kind words and careful assessment and hope that this work 'might receive a good deal of use in college courses across the curriculum', as one of them put it.

AI was occasionally used to improve the writing style.

Financial support. This research was supported by the Dutch Research Council (NWO), grant id https://doi.org/10.61686/MKYGD09939.

The author has no relevant financial or non-financial interests to disclose.

Open access funding provided by Eindhoven University of Technology.

References

Adam D. (2025). 'When AI Rejects Your Grant Proposal: Algorithms are Helping to Make Funding Decisions.' *Nature* **645**(8082), 832–33.

Allen L., Scott J., Brand A., Hlava M. and Altman M. (2014). 'Publishing: Credit Where Credit is Due.' *Nature* **508**(7496), 312–13.

Alvarado R. (2021). Explaining Epistemic Opacity. Retrieved July 20, 2026, from https://philsci-archive.pitt.edu/19384/

Ananya. (2024). 'AI Image Generators Often Give Racist and Sexist Results: Can They Be Fixed?' *Nature* **627**(8005), 722–25.

Anderson E. (2012). 'Epistemic Justice as a Virtue of Social Institutions.' *Social Epistemology* **26**(2), 163–73.

Aristotle and Crisp R. (Ed.). (2014). *Aristotle: Nicomachean Ethics*. Cambridge: Cambridge University Press.

Association for the Advancement of Artificial Intelligence. (2025). AAAI launches AI powered peer review assessment system. https://aaai.org/aaai-launches-ai-powered-peer-review-assessment-system/

Baehr J. (2013). 'Educating for Intellectual Virtues: From Theory to Practice.' *Journal of Philosophy of Education* **47**(2), 248–62.

Boge F.J. (2022). 'Two Dimensions of Opacity and the Deep Learning Predicament.' *Minds and Machines* **32**(1), 43–75.

Bornmann L., Mutz R. and Daniel H.D. (2010). 'A Reliability-Generalization Study of Journal Peer Reviews.' *PLoS One*, **5**(12), e14331.

Brown T., Mann B., Ryder N., Subbiah M., Kaplan J.D., Dhariwal P. . . . and Amodei D. (2020). 'Language Models are Few-Shot Learners.' *Advances in Neural Information Processing Systems* **33**, 1877–901.

Burger B., Maffettone P.M., Gusev V.V., Aitchison C.M., Bai Y., Wang X. . . . and Cooper A.I. (2020). 'A Mobile Robotic Chemist.' *Nature*, **583**(7815), 237–41.

Byerly T. R (2023). 'Intellectual Honesty and Intellectual Transparency.' *Episteme* **20**(2), 410–28.

Code L. (1984). 'Toward a Responsibilist'epistemology.' *Philosophy and Phenomenological Research* **45**(1), 29–50.

Code L. (1987). *Epistemic Responsibility*. Hanover: University Press of New England for Brown University Press.

Cortina J.M.A., Fatapour Y., Brown K.L., Gisladottir U., Zietz M., Don't Walk O.J.B. . . . and Tatonetti N.P. (2025). 'Biases in Race and Ethnicity Introduced by Filtering Electronic Health Records for "Complete Data": Observational Clinical Data Analysis.' *JMIR Medical Informatics* **13**(1), e67591.

Crawford K. (2021). *Atlas of AI* (Vol. 529). New Haven: Yale University Press.

DeBruine L. and Lakens D. (2025). Metacheck: Check Research Outputs for Best Practices. R package version 0.0.0.9064. https://github.com/scienceverse/metacheck.

Farina M., Zhdanov P., Karimov A. and Lavazza A. (2024). 'AI and Society: A Virtue Ethics Approach.' *AI & Society* **39**(3), 1127–40.

Fricker M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.

Glickman M. and Zhang Y. (2024). AI and Generative AI for Research Discovery and Summarization. *arXiv preprint arXiv*:2401.06795.

Hagendorff T. (2022). 'A Virtue-Based Framework to Support Putting AI Ethics into Practice.' *Philosophy & Technology* **35**(3), Article 55.

Humphreys P. (2009). 'The Philosophical Novelty of Computer Simulation Methods.' *Synthese* **169**(3), 615–26.

Ji Z., Lee N., Frieske R., Yu T., Su D., Xu Y. ... **and Fung P.** (2023). 'Survey of Hallucination in Natural Language Generation.' *ACM Computing Surveys* **55**(12), 1–38.

Kaplan D.M., Palitsky R., Arconada Alvarez S.J., Pozzo N.S., Greenleaf M.N., Atkinson C.A. and Lam W.A. (2024). 'What's in a Name? Experimental Evidence of Gender Bias in Recommendation Letters Generated by ChatGPT.' *Journal of Medical Internet Research* **26**, e51837.

Kitcher P. (1990). 'The Division of Cognitive Labor.' *The Journal of Philosophy* **87**(1), 5–22.

Krumrei-Mancuso E.J., Haggard M.C., LaBouff J.P. and Rowatt W.C. (2020). 'Links between Intellectual Humility and Acquiring Knowledge.' *The Journal of Positive Psychology* **15**(2), 155–70.

Krumrei-Mancuso E.J. and Rouse S.V. (2016). 'The Development and Validation of the Comprehensive Intellectual Humility Scale.' *Journal of Personality Assessment* **98**(2), 209–21.

Lahroodi R. (2007). 'Collective Epistemic Virtues.' *Social Epistemology* **21**(3), 281–97.

Leonelli S. (2016). *Data-Centric Biology: A Philosophical Study*. Chicago and London: University of Chicago Press.

Leonelli S. (2019). 'What Distinguishes Data from Models?' *European Journal for Philosophy of Science* **9**(2), 22.

Liang W., Zhang Y., Wu Z., Lepp H., Ji W., Zhao X. ... **and Zou J.** (2025). 'Quantifying Large Language Model Usage in Scientific Papers.' *Nature Human Behaviour* **9**(12), 2599–2609.

Montmarquet J. (1992). 'Epistemic Virtue and Doxastic Responsibility.' *American Philosophical Quarterly* **29**(4), 331–41.

Montmarquet J.A. (1987). 'Epistemic Virtue.' *Mind* **96**(384), 482–97.

Naddaf M. (2025). 'Major AI Conference Flooded with Peer Reviews Written Fully by AI.' *Nature* **648**(8093), 256–57.

Nguyen C.T. (2022). 'Trust as an Unquestioning Attitude.' *Oxford Studies in Epistemology* **7**, 214–44.

Orona G.A. and Pritchard D.H. (2022). 'Inculcating Curiosity: Pilot Results of an Online Module.' *Assessment & Evaluation in Higher Education* **47**(6), 867–81.

Paullada A., Raji I.D., Bender E.M., Denton E. and Hanna A. (2021). 'Data and its (Dis)contents: A Survey of Dataset Development and Use in Machine Learning Research.' *Patterns* **2**(11), 100336.

Rezazade Mehrizi M.H., Mol F., Peter M., Ranschaert E., Dos Santos D.P., Shahidi R. ... **and Dratsch T.** (2023). 'The Impact of AI Suggestions on Radiologists' Decisions: A Pilot Study of Explainability and Attitudinal Priming Interventions in Mammography Examination.' *Scientific Reports* **13**(1), 9230.

Roberts R.C. and Wood W.J. (2007). *Intellectual Virtues*. Oxford: Oxford University Press.

Sikimić V. (2025). 'Fair or Flawed? Rethinking Grant Review with Generative AI.' *Synthese* **206**(6), 282.

Sikimić V. and Herud-Sikimić O. (2022). 'Modelling Efficient Team Structures in Biology.' *Journal of Logic and Computation* **32**(6), 1109–28.

Smith K. (2023). 'The Funding Gender Gap.' *Nature* **617**(7959), 28–29.

Sullivan E. (2022). 'Inductive Risk, Understanding, and Opaque Machine Learning Models.' *Philosophy of Science* **89**(5), 1065–74.

Vallor S. (2016). *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting*. New York: Oxford University Press.

Van Der Sijs H., Aarts J., Vulto A. and Berg M. (2006). 'Overriding of Drug Safety Alerts in Computerized Physician Order Entry.' *Journal of the American Medical Informatics Association* **13**(2), 138–47.

Van Noorden R. (2020). 'The Ethical Questions that Haunt Facial-Recognition Research.' *Nature* **587**(7834), 354–59.

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press

Van Noorden R. and Perkel J.M. (2023). 'AI and Science: What 1,600 Researchers Think.' *Nature* **621**(7980), 672–75.
Vučković A. and Sikimić V. (2023). 'How to Fight Linguistic Injustice in Science.' *Social Epistemology* **37**(1), 80–96.
Watson L. (2015). 'What is Inquisitiveness.' *American Philosophical Quarterly* **52**(3), 273–87.
Whitcomb D. (2010). 'Curiosity was Framed.' *Philosophy and Phenomenological Research* **81**(3), 664–87.
Zednik C. (2021). 'Solving the Black Box Problem.' *Philosophy & Technology* **34**(2), 265–88.
Zollman K.J. (2010). 'The Epistemic Benefit of Transient Diversity.' *Erkenntnis* **72**(1), 17–35.

Vlasta Sikimić is a tenured Assistant Professor of Philosophy of Science at Eindhoven University of Technology, working on meta-science in the age of AI by connecting virtue epistemology with philosophy of AI in science and education. She leads projects on the virtuous use of AI in research and teaching, funded by the Dutch Research Council and the Centre for Engineering Education of the federation of the four technical universities in the Netherlands. Vlasta is an elected member of the Steering Committee of the European Philosophy of Science Association and serves on the Board of the Dutch Association for Logic and Philosophy of the Exact Sciences. She was a member of the working group that developed the Serbian national ethical guidelines for safe and reliable use of AI, and she speaks regularly about AI on television and on science podcasts.

Cite this article: Sikimić V. (2026). "Navigating AI-Accelerated Science: An Intellectual Virtue Approach to Responsible Inquiry." *Episteme* 1–18. https://doi.org/10.1017/epi.2026.10140

https://doi.org/10.1017/epi.2026.10140 Published online by Cambridge University Press