



Cite this: DOI: 10.1039/d5ta05512a

Data integrity in materials science in the era of AI: balancing accelerated discovery with responsible science and innovation

Nik Reeves-McLaren ^a and Sarah Moth-Lund Christensen ^b

Artificial intelligence promises to revolutionise materials discovery through accelerated prediction and optimisation, yet this transformation brings critical data integrity challenges that threaten the scientific record. Recent studies demonstrate that experts cannot reliably distinguish AI-generated microscopy images from authentic experimental data, while widespread errors plague 20–30% of materials characterisation analyses. Generative AI tools can now produce code for data manipulation at pace, creating plausible-looking results that violate fundamental physical principles yet evade traditional peer review. These risks are compounded by inherent biases in training datasets that systematically over-represent equilibrium-phase oxide systems, and by the “black box” opacity of AI models that challenges scientific accountability and epistemic agency. We propose a multifaceted framework for enhanced research integrity encompassing materials-specific ethical governance, professional standards for AI disclosure and data validation, and modular integrity checklists with technique-specific validation protocols. Critical enablers include mandatory deposition of structured raw instrument files, AI-powered fraud detection systems, and cultivation of critical AI literacy through interdisciplinary education. Without immediate action to address these challenges, the materials science community risks perpetuating errors and biases that will fundamentally undermine AI's transformative potential.

Received 8th July 2025
Accepted 31st October 2025

DOI: 10.1039/d5ta05512a

rsc.li/materials-a

1 Introduction

Consider the scene: an early-career researcher, comparing their experimental data to a published standard, instructs a generative artificial intelligence to ‘improve’ their results to better match the reference dataset. The AI complies, subtly altering the data points. The researcher subsequently asks their supervisor whether this action is an acceptable method of data processing – or maybe they don't. This situation illustrates a serious emerging threat to the integrity of the scientific record, one that extends far beyond issues of academic writing and into the manipulation of primary experimental data itself.

The severity of this threat was recently demonstrated in nanomaterials research, where a survey of 250 scientists found that experts could not reliably distinguish AI-generated microscopy images from authentic experimental data.¹ These AI-generated images were created in under one hour using publicly available tools, requiring no specialised technical knowledge. The traditional peer review process, reliant on visual inspection by experts, is no longer sufficient to detect sophisticated image fraud.

These challenges appear at a time when artificial intelligence (AI) is set to reshape materials science, promising rapid discovery of advanced materials by predicting material properties, optimising compositions, and exploring vast chemical design spaces. Emergent examples include the development of alloys with superior mechanical properties,² the generation of numerous MOF candidates, and advances in battery material discovery.^{3–5} Developments such as Google DeepMind's GNOME (graph networks for materials exploration) have demonstrated the potential for large-scale materials discovery, identifying 2.2 million stable crystal structures and representing an order-of-magnitude expansion in known stable materials.⁶ The reliability of such AI models, however, depends entirely on the integrity of their training data.⁷

High-quality, relevant, and representative data is essential for accurate and effective generalisation. The principle “garbage in, garbage out” is key: if training data are limited or flawed, AI models will be inaccurate.⁸ Intense debate followed a 2023 publication on an automated lab for rapid synthesis and characterisation of ‘new’ inorganic materials, with critiques of the work focussed around issues of metadata and what constitutes a novel discovery, the quality of automated analyses, and the ability to model the complexities of real materials, such as disorder.^{9,10}

Despite the clear promise of AI, widespread errors and inconsistencies in data, along with fraudulently manipulated or

^aSchool of Chemical, Materials and Biological Engineering, University of Sheffield, Mappin Street, Sheffield, S1 3JD, UK. E-mail: nik.reeves-mclaren@sheffield.ac.uk

^bCentre for Machine Intelligence, University of Sheffield, Mappin Street, Sheffield S1 3JD, UK



fabricated data threaten research validity in materials characterisation. Generative AI (GenAI) is readily capable of, for example, producing code to manipulate data, and then cover one's tracks, without asking any challenging ethical questions of researchers increasingly under 'publish or perish' pressures. Manipulating data to report a new room temperature superconductor would be discredited within hours of publication, but smaller iterative materials developments are more likely to sneak past peer review.

There is an urgent need for research on, and new approaches to, data integrity. Given the broad uptake of generative AI in materials science, and across all disciplines in engineering and the physical sciences, small errors, biases in foundational training data, and outright unethical conduct risk widespread research misdirection.

2 Outlining the data integrity challenges

The challenges of AI-driven or assisted research can be classified into several areas, each impacting the validity and trustworthiness of scientific findings.

2.1. Widespread errors and underused verification methods

Studies show that a large proportion, from 20% to 30%, of data analyses across various common materials characterisation techniques, contain basic inaccuracies.⁷ A recent study used an AI tool to examine over 3000 papers in Organic Letters and found only 40% of chemical research papers had error-free mass measurements.¹¹ The study also found cases where miscalculated values seemed validated by experimental measurements, casting doubts on researcher understanding, as well as raising concerns about potential data fabrication.

There are widespread issues with the underutilisation of well-established physical consistency checks in materials science data analyses, compounded by many instances of poor understanding of statistical measures employed to judge the perceived quality of work. Rietveld refinement is a powerful tool for extracting structural information from powder diffraction data, but misinterpretations of statistical measures are common; one example is the effective 'goodness of fit' of the diffraction pattern calculated from the refined structure to the experimental data, the reduced chi-squared (χ^2). A critical misunderstanding is evident when reports on a refinement quote a χ^2 value less than 1.0, statistically problematic as it implies a fit that is "better than ideal". This can indicate either that the standard uncertainties associated with the observed data are overestimated, or that too many parameters have been introduced, leading to overfitting of the model to noise rather than true physical phenomena. The result is publication of structural parameters that are statistically unreliable or physically meaningless.^{12,13} Furthermore, many publications fail to report or justify crucial details of the refinement model itself, such as the mathematical function used to model the peak profiles and background, the constraints applied to parameters, or the handling of atomic displacement parameters (ADPs).

This frequently leads to the publication of physically nonsensical results, such as negative ADPs, and structural models that are statistically unsound and ultimately irreproducible.¹⁴

A further example: despite proven utility for ensuring consistency in dielectric functions and accurate optical and electronic property measurements, methods such as *F*-sum rules and Kramers–Kronig (K–K) relations are reportedly often overlooked in research on optical materials.⁷ K–K relations are mathematical constraints linking the real and imaginary components of optical constants, derived from fundamental causality requirements. Violation of these relations – or of *F*-sum rules, which constrain integrated absorption based on electron density – indicates either measurement errors, incomplete spectral data, or data manipulation. The failure to apply such validation methods leaves optical property claims vulnerable to fabrication, particularly as GenAI tools could generate superficially plausible spectra that nevertheless violate basic physical constraints.


These types of widespread shortcomings highlight potentially severe issues around the reliability of reported materials and their properties, creating a substantial barrier to the development of high-performance advanced materials. Without improvement in data integrity, handling and reporting, we risk these shortcomings becoming fixtures of AI training and validation data sets – in turn undermining the promise of AI in materials science and leaving us instead with unreliable models, and misdirected research.

2.2. Deliberate data manipulation and synthetic data risks

Research misconduct can be defined to include data fabrication, falsification, or plagiarism committed intentionally, knowingly, or recklessly, representing a significant departure from accepted research practices.^{15,16} The recent reporting of around 800 papers published in crystallography and exotic-chemistry journals originating from "paper mills" highlights one example of such systemic fraud.^{7,17} Other reports show 3.8% of published papers in biomedical research contain inappropriate image duplication.¹⁸

The arrival of GenAI brings new and complex ethical and scientific problems, at a time when research integrity in materials science is already under pressure. Highly realistic synthetic GenAI data and images can easily be misrepresented as experimental. Real data can be altered to better support scientific hypotheses. This capability poses serious risks to research integrity. Traditional methods for detecting fraud, such as identifying non-random digits, are now obsolete due to GenAI's sophistication, leading to an "arms race" between AI tools for detection and new methods designed to avoid them.¹⁵

The growth in use of text produced using GenAI tools such as ChatGPT, Gemini, Claude *etc.* means that many journals now require authors to declare where these have been used. But what about manipulation or fabrication of raw data? There is much less awareness around this risk. One GenAI tool these authors tested was able to yield reusable Python code for data manipulation (to remove secondary phase peaks in diffraction data and fill the resulting void with randomly generated believable



background, or to manipulate long-term battery testing data to remove noise and glitched cycles) in comfortably less than an hour. This is a challenge for the here and now.

Recent work in nanomaterials characterisation provides sobering empirical evidence of these capabilities. Researchers generated convincing fake atomic force microscopy (AFM), scanning transmission electron microscopy (STEM), and transmission electron microscopy (TEM) images in less than one hour using commercially available generative AI tools.¹ When presented to 250 scientists in a blind survey, experts correctly identified real *versus* AI-generated images only 40–51% of the time for most image pairs – performance indistinguishable from random guessing. For four out of six image pairs tested, statistical analysis (chi-squared test, $p > 0.05$) showed no significant difference in scientists' ability to identify authentic *versus* fake images.

Energy materials research offers another good example of how research fields are facing vulnerabilities to AI-assisted data manipulation, in this case due to the complex, multi-parameter nature of electrochemical measurements. Photovoltaic current–voltage characteristics are readily susceptible to algorithmic enhancement, where fill factors could be artificially improved from experimentally observed values of 0.83 to theoretically optimal values approaching 0.89 through subtle modification of series resistance contributions.¹⁹ Such manipulations remain within plausible ranges for peer review assessment whilst significantly inflating reported power conversion efficiencies. Electrocatalyst performance data present similar vulnerabilities through fabrication of Tafel slope values. Experimental studies demonstrate that Pt/C catalysts exhibit Tafel slopes varying from 30 mV dec^{−1} in 0.5 M H₂SO₄ to 120 mV dec^{−1} under fuel cell conditions, with additional dependence upon catalyst loading (63 to 211 mV dec^{−1} across different overpotential ranges for identical materials).²⁰ It is perceivable that GenAI tools could readily generate synthetic data presenting artificially consistent Tafel slopes of 30 mV dec^{−1} across varied conditions, thereby suggesting superior kinetic performance. Electrochemical impedance spectroscopy measurements in battery and fuel cell research are similarly vulnerable, where complex multi-semicircle Nyquist plots can be algorithmically simplified to eliminate inconvenient high-frequency resistances or low-frequency inductive features associated with side reactions or interface instabilities.²¹

If a research group was to fraudulently manipulate data and present the discovery of a room-temperature superconductor based on these types of subtle data hacks, they would likely be found out on the same day. For the more iterative work on less transformative materials that makes up much of the publication record – we propose this is currently much more likely to slip through the net unnoticed.

The Retraction Watch database currently runs to almost 60 000 records, but none yet specifically focus on the use of AI for the purposes discussed here. One tangentially related case claimed to track the productivity of “a thousand material scientists” at a large R&D company, reporting a “44% increase in materials discovery” and an “81% productivity increase for top-decile scientists” due to the introduction of a “machine

learning material generation tool”. The data included were found to be “suspiciously clean and neat: nearly every sub measure of success gave a clear and statistically significant result”. Following an internal, confidential review, the sole author's institution concluded that the paper “should be withdrawn from public discourse” and requested the paper's withdrawal.²²

2.3. Inherent biases and quality shortcomings in AI training datasets

Materials characterisation data can often be noisy, incomplete and inconsistent, which directly impacts the performance of machine learning models. For example, a comprehensive review of over 1300 research papers focusing on X-ray photoelectron spectroscopy (XPS) analyses revealed wide-ranging issues with the quality and reliability of published data, directly exemplifying these challenges. Over 40% of publications contained errors that could have significantly affected the conclusions drawn; 35% neglected to provide details on the spectrometer used, and 85% did not specify the analytical software used.²³

It may seem unintuitive to think of bias in the context of materials science data, but our scientific track record – the AI's training data – fundamentally overrepresent stable, inorganic, equilibrium-phase systems, particularly oxide based materials, especially relative to amorphous, disordered or highly entropic materials.²⁴ AI models trained on such data struggle to generalise or extrapolate to new, unexplored chemistries or processing conditions, often leading to “hallucinations” or unreliable predictions for novel materials. Generative AI models can accidentally learn and then amplify any biases and shortcomings present in their training data.²⁵

For instance, if AI models are predominantly trained on data from materials developed and characterised under specific, well-established (*e.g.*, equilibrium) conditions or for certain applications, they may inadvertently learn to prioritise or ‘hallucinate’ properties that conform to these existing paradigms. This could lead to a biased prediction landscape, where novel materials with unusual or non-equilibrium properties, or those relevant to emerging applications, are systematically overlooked or inaccurately predicted. Such a bias could perpetuate existing research trajectories, effectively ‘stereotyping’ what constitutes a ‘good’ or ‘feasible’ material based on historical data, rather than enabling truly disruptive discoveries.

2.4. Challenges of transparency, explainability, and human oversight in AI systems

A significant hurdle to data integrity in AI-driven materials science is the “black box” problem of AI models, wherein the ‘reasoning’ or the way the systems gets to a specific output is an opaque process for users and programmers alike. As such an AI model may suggest adding a tiny amount of a given element to an alloy will significantly improve its tensile strength, but one might struggle to get an accurate and directly related explanation on *why* it predicts these properties. This issue pertains



especially to AI models, such as deep learning neural networks, where complexity results in opaque decision-making.²⁵ A lack of transparency and explainability severely complicates and hinders understanding of AI outputs, accountability, identification of causes of errors and biases, as well as potentially limiting trust in model predictions.

To reduce these risks and ensure responsible AI deployment, meaningful human control and oversight are essential. This involves actively monitoring AI behaviour and developing plans to prevent harmful effects on users, with human validation being crucial for high-risk decisions.²⁵ Ethical guidelines for trustworthy AI, such as those from the European Commission, outline different levels of human involvement and oversight of AI system activity, including consideration of societal and ethical impacts, and ultimate decision-making.²⁶

The impact of AI on human “epistemic agency” – the control individuals have over their beliefs, the questions they ask, and the reasons they entertain – is also a critical concern. There is an ongoing discussion about whether AI-based science poses a social epistemological problem, particularly concerning trust in opaque models and the responsibility of scientists for outputs based on AI models.²⁷ Some argue that full transparency is not always needed for trust if systems follow established academic and institutional norms, but this applies to human and institutional actors, not to AI models which cannot be held to norms in the same way. Therefore, how can materials scientists ensure research integrity and fully take responsibility for AI tools, when they cannot foresee, fully understand nor verify how these tools get to a given specific output? For example, if an AI model identifies a novel battery electrolyte composition but cannot explain why certain additives improve ionic conductivity, researchers cannot properly assess safety risks or optimise the formulation further. Similarly, if an AI predicts a ceramic will exhibit ferroelectric properties but provides no mechanistic insight, experimental validation becomes trial-and-error rather than hypothesis-driven science. This requires an evolving understanding of scientific responsibility and epistemic agency in the AI era. The traditional idea of scientific responsibility assumes a human agent's full understanding and control over their research tools.²⁷ AI's opacity directly challenges this, raising basic questions about who is accountable when an AI system makes a flawed decision or generates inaccurate data, and curtailing the scientist's ability to articulate the reasons and evidence supporting AI-generated hypotheses. This implies a profound shift in the epistemology of science, and a new understanding of human agency in research integrity and accountability.^{25–27}

3. Developing frameworks for enhanced research integrity in AI-driven materials science

Dealing with data integrity issues in AI-driven materials science will require a multi-faceted approach, combining governance, professional standards, and human vigilance.

3.1. Ethical foundations and materials-specific governance

With rapid developments in AI, the corresponding AI regulatory and governance landscape is trying to play catchup under the headings of primarily “ethical AI”, “responsible AI”, and “trustworthy AI”. Despite the language flux, core principles are emerging that mirror the technical limitations and concerns for particularly GenAI: transparency, explainability, accountability, fairness, privacy and safety. Key initiatives, such as the National Institute of Standards and Technology's (NIST) research on trustworthy AI characteristics and UNESCO's global standard on the Ethics of Artificial Intelligence, highlight these shared priority principles for AI development.

Despite the concerted efforts in AI governance, the practical challenges for materials science-oriented adoption of these principles and governance frameworks is and will not be straightforward. Even these broad principles need to be translated into specific guidelines and governance mechanisms relevant to the specific challenges of materials science data. Further, the current generalised list of ethical AI principles can and should not be expected to be exhaustive. One must expect that considerable work will be needed to identify potential ethical challenges posed by the use and development of AI specifically in and for materials science research. This includes ethical challenges requiring both technical and non-technical address, for example potentially relating to sustainability, dual use, and how AI might change perceptions and assumptions about materials science research practises.

Frameworks for assessing “AI-ready” data are appearing to deal with some of these problems. The SciHorizon framework, for instance, suggests four main aspects: quality, FAIRness (findable, accessible, interoperable, reusable), explainability, and compliance.⁸ Key parts of the ‘quality’ component include completeness, accuracy, consistency (both internal coherence within a dataset and external alignment to related datasets), and timeliness (prompt publication and continuous updating). ‘Compliance’ stresses the importance of data provenance (clear documentation of data sources, authorship, and licensing), ethics & safety (adherence to scientific ethical standards), and trustworthiness (compliance with national regulations and sustainability of data services). This shifts the focus from sheer data volume to the quality, relevance, and representativeness of the data, but implicit in this is a fundamental re-evaluation of how scientific data is collected, curated, and prepared – with both good practice and AI in mind.

3.2. Professional standards and best practices for materials data and AI

Professional bodies and scientific publishers increasingly set specific standards and best practices for ethical AI use in scientific research and its dissemination through journal articles. For the materials science community, Nature Portfolio journals, including *Nature Materials*, have established clear guidelines for authors and peer reviewers regarding AI tools.²⁸ Their policies establish that accountability for work cannot be effectively applied to AI tools, precluding large language models such as ChatGPT from author attribution on publications. A



nuanced distinction is made for AI-assisted improvements to human-generated texts for readability, style, grammar, spelling, and tone, which does not require declaration. However, in all cases, human accountability for the final text remain paramount.

The Nature Portfolio also largely prohibits the use of generative AI for images due to unresolved legal copyright and research integrity issues. This stance, while understandable from a legal and ethical perspective, may be an implicit hurdle for AI-driven materials discovery workflows reliant on generative models to propose novel material structures or microstructures, where visual representation is key.

Recent calls from the microscopy community emphasise the need to reframe expectations around image quality, recognising that not every nanomaterial or assembly is perfect, and that pristine images may signal manipulation rather than excellence.¹ Excessive demands for polished images create pressure on researchers that can inadvertently incentivise AI usage. Reviewers should not request “better looking” images unless visual improvements would change the authors’ scientific conclusions. The purpose of images is to support conclusions and enable fair judgment, not to serve as polished content for dissemination. Editors must actively dismiss such reviewer comments when they are scientifically unjustified, recognising that this pressure contributes to the integrity crisis.

Standardised validation protocols are also becoming more prominent. There is a recognised need for new norms, standards, and best practices for conducting research with AI. Data provenance is vital for AI authentication, transparency, and traceability. The detailed requirements for documentation, including researcher responsibilities, workflow, input, output, metadata, origin/access point, and data management, go beyond a traditional citation. This implies that for AI-driven materials science, true reproducibility and trustworthiness depend not just on the final model or results, but on a carefully documented “data trail” from raw source to final output, especially given AI’s potential to hallucinate and/or generate synthetic data.

3.3. Using AI as a tool for materials data quality control

Paradoxically, while AI brings new integrity challenges, it also offers powerful ways to detect misconduct, including data manipulation, image fraud, and fabricated results. AI systems can analyse datasets for statistical inconsistencies and patterns suggesting fabrication or manipulation, using machine learning models to compare experimental results with established scientific principles and assess adherence to expected distributions or statistical norms. Natural language processing tools can check for consistency between text descriptions, figures, and tables, flagging any differences for authors and editors alike.

The potential for automated quality assurance and better peer review processes specifically within materials science is significant. Some publishers and research institutions already use AI tools to scan submitted manuscripts for image integrity problems before peer review. Christmann’s study showed the

power of AI-powered data analysis in uncovering previously unknown systematic errors in chemical publications highlights AI’s capability for automated quality control in chemical and materials data.¹¹ The future must bring better collaboration between AI and human reviewers to improve fraud detection. This will likely involve AI handling issues of scale and initial pattern detection, with human experts then providing critical judgement, contextual understanding, and deal with nuanced cases AI might miss.

The Science family of journals has adopted Proofig, an AI-powered image-analysis tool, to screen for manipulation.²⁹ However, ethical implementation requires that AI-flagged suspicions be reviewed by humans, with outcomes communicated to authors who must have opportunity to respond, in accordance with Committee on Publication Ethics (COPE) guidelines. Such tools should be deployed both during submission and retrospectively to audit previous publications, with the sophistication of anti-fraud measures potentially serving as an indicator of journal quality.

3.4. Cultivating a culture of data responsibility and critical AI literacy in materials science

Ultimately, researchers are responsible for checking the accuracy of data, AI-generated output and ensuring that data provenance is carefully maintained. It is not enough to just use AI tools; researchers must become responsible guardians who understand AI’s abilities, limits, as well as its ethical and scientific implications, actively checking its outputs and ensuring proper disclosure. This necessitates at least two additions or changes to the norms of material science research. First, it requires higher level of general data and AI literacy among researchers. Second, as the state of AI and available AI research tools is far from static, it also requires researchers to continuously reflect upon (1) the potential ethical concerns in their adoption of these tools, and (2) whether their trust, understanding and control of given AI tools and their outputs is adequate for genuine knowledge creation.

Fostering a reflective research culture by bringing AI ethics and sound data analysis skills into scientific education and training is an essential and proactive step. For instance, Freie Universität Berlin’s Department of Biology, Chemistry, and Pharmacy plans to bring AI tools into its curriculum to help students develop strong data analysis skills and critical thinking, preparing them for their future research careers.³⁰ Similarly, Cornell Engineering has started a graduate-level course, “AI for materials”, designed to give the next generation of researchers and engineers the knowledge to drive discovery where AI and materials science meet, highlighting both applications and the challenges involved. The challenges posed by AI can also be seen as an opportune moment to foster new interdisciplinary relations, and to benefit from the strengths of disciplines, who specialise in ethics or meta-science matters. As an example, “Embedded EthiCS” at Harvard is a collaboration between computer scientists and philosophers fuelling both teaching and research on ethical concerns in AI development and adoption.



Encouraging a culture of critical reflection and healthy scepticism towards AI outputs is essential. It is vital to develop a culture within materials science that sees AI models as tools requiring responsible use while instilling strong data analysis skills and critical thinking skills in researchers – at all levels, ensuring we are all equipped with the necessary knowledge and ethical grounding to navigate the complexities of AI-driven science responsibly.

To supplement individual diligence, the community should also consider adopting established structural approaches from other scientific disciplines designed to improve the reliability of research findings. Adversarial collaborations, for instance, unite researchers with conflicting viewpoints to jointly design and conduct a critical experiment, increasing the impartiality of the outcome. The Registered Reports publication format, where methods and analysis protocols are peer-reviewed before

experiments are conducted, mitigates publication bias and questionable research practices. Finally, a 'Red Team' approach, where researchers actively solicit rigorous, structured criticism of a project from designated colleagues prior to submission, can identify weaknesses in argumentation and data interpretation that might otherwise be missed. The adoption of such practices would represent a systemic commitment to research integrity.

3.5. Practical implementation: data integrity checklists and validation frameworks

To move from identifying these challenges towards actionable solutions, we first propose a modular checklist for authors, reviewers, and journal editors, inspired by best practices in clinical research such as the INSPECT-SR criteria for systematic

Table 1 Modular data integrity checklist for materials science publications. The core data integrity module is required for all submissions, while technique-specific modules should be applied as relevant to the methods used. This framework provides concrete validation steps for authors, editors, and reviewers while accommodating the diverse methodological landscape of materials research. Additional modules can be developed for other characterisation techniques as community priorities emerge

Module 1: core data integrity (required for all submissions)		
Category	Item	Status
AI Usage & transparency	All AI tools explicitly declared with specific usage scope	☐
	Any AI involvement in data processing clearly documented	☐
	All AI-generated content validated by qualified researchers	☐
	Synthetic/AI-generated data clearly labelled and benchmarked	☐
Data provenance	Raw instrument files organised in structured format (e.g., MAIF) and made publicly available	☐
	Dataset origins and collection conditions documented	☐
	Appropriate uncertainty quantification throughout	☐
	Outlier identification and treatment methods described	☐
	Sample sizes statistically justified for conclusions	☐
Validation	Key results cross-validated where possible	☐
	Known limitations and biases acknowledged	☐
	Sufficient detail provided for independent replication	☐
Technique-specific modules (add as applicable)		
Example module A: X-ray diffraction & Rietveld refinement		
Scan parameters, instrument configuration, and sample preparation fully documented		☐
All significant peaks assigned or unassigned peaks discussed		☐
Peaks appropriately labelled in figures with e.g. phase symbols or Miller indices		☐
Background function choice and number of fitting terms presented and justified; systematic residuals addressed		☐
Overly symmetric peaks or suspiciously smooth backgrounds (or sections thereof) explained		☐
χ^2 values reported and physically reasonable (1.0–3.0; ≈ 1.0 requires justification)		☐
R -factors (R_p , R_{wp} , R_e) present and appropriate for data quality and model complexity		☐
Crystal structure(s) chemically reasonable, refined parameters have estimated standard deviations, and appropriate bond lengths/angles presented and discussed		☐
Diffraction patterns show realistic noise levels and instrument-specific artifacts		☐
Any use of machine learning or AI tools used disclosed and justified – for example in phase analyses, peak fitting or least squares refinement.		☐
Example module B: electrochemistry & battery testing		
Multiple cell testing with error bars and/or reproducibility assessment		☐
Detailed discussion of cell fabrication and testing regimes present		☐
Theoretical capacities and method to calculate (dis)charge 'C-rate' presented		☐
No impossible efficiencies (>100%) or super-theoretical capacities		☐
If appropriate, control experiments included (e.g., electrolyte-only cells)		☐
Environmental conditions (temperature, humidity) documented		☐
Capacity fade rates realistic and mechanistically consistent		☐
Cycling curves show realistic noise levels and expected measurement artifacts		☐
Glitch handling documented: unexplained discontinuities, missing cycles, or data gaps addressed		☐
Overly smooth capacity curves or suspiciously consistent coulombic efficiencies investigated		☐



reviews.³¹ This approach provides a structured framework for data validation specific to common materials science techniques. The core module, Table 1, outlines general principles of data integrity and AI usage declaration applicable to all manuscripts. We have also suggested initial technique-specific modules for X-ray diffraction with Rietveld refinement and for electrochemical battery testing, two areas where data quality is frequently suboptimal. This modular design is extensible, and we invite community contributions to develop and validate further modules.

A critical enabler for this checklist is a policy mandating the deposition of raw experimental data. The distinction between raw and processed data is vital for integrity. It is insufficient to provide only processed data, such as a text file of a diffraction pattern (.xy); journals must require the structured, machine-readable raw data files generated by the instrument itself (e.g., .raw, .xrdml). Raw instrument files contain a rich set of metadata – including calibration parameters, detector settings, and collection times – that are essential for reproducing the analysis and verifying the data's origin. This embedded metadata makes the convincing fabrication of a raw data file substantially more difficult than creating a simple text file of processed numbers. To make verification of these files practical rather than overwhelming, recent proposals from the nanomaterials community suggest adopting standardised data storage structures.¹ The minimal arrangement of instrument files (MAIF) framework proposes that each manuscript has its own folder, with each figure having a subfolder containing primary instrument files specific to that figure, and non-figure data stored in a separate 'additional data' folder.¹ This structured approach – as opposed to idiosyncratic, researcher-specific filing systems – enables efficient checking of key instrument files for legitimacy without overwhelming reviewers or investigators. We recommend that journals require structured raw data files (following MAIF or similar principles) as a publication criterion, published as compressed directories in supplementary information or repositories such as Zenodo, Open Science Framework, or Figshare. This requirement should over time become mandatory rather than merely encouraged.

While the checklist approach we propose addresses the integrity of individual studies, a distinct but related challenge is the quality control of the large, aggregated datasets upon which foundational AI models are built. In fields such as clinical science, where meta-analyses of randomised controlled trials face similar issues with flawed or fraudulent data, researchers have developed formal tools to identify problematic studies before their inclusion in a wider analysis.³¹ A parallel approach is required in materials science to ensure that AI models are not trained on compromised data.

Accordingly, we propose the development of a complementary data-vetting framework specifically for the curation of AI training sets. Such a framework would consist of a series of checks to be applied to any dataset being considered for inclusion in a larger corpus, including: (i) verification of the publication status of the source data, checking for retractions, corrections, or expressions of concern; (ii) screening against public post-publication review platforms for credible criticisms;

and (iii) programmatic checks for statistical anomalies or physically implausible results within the data itself, such as efficiencies exceeding 100% or unrealistic electrochemical parameters.

These proposed actions are not one-and-done fixes, and as shown earlier connects with larger reflections upon how AI interacts with our notion of trustworthy research and knowledge generation, therefore instead requiring continuous involvement and engagement from both the materials science research community and beyond. However, we hope with these initial suggestions to open up for a focused discussion on responsible AI adoption in Materials Sciences, and encourage a wider dialogue on how we can ensure trustworthy and responsible research practices in the AI age.

Conflicts of interest

There are no conflicts to declare.

Data availability

No original data are included in this perspective.

References

- 1 N. Davydiuk, *et al.*, The rising danger of AI-generated images in nanomaterials science and what we can do about it, *Nat. Nanotechnol.*, 2025, **20**, 1174–1177, DOI: [10.1038/s41565-025-02009-9](https://doi.org/10.1038/s41565-025-02009-9).
- 2 F. Wang, *et al.*, Experimentally validated inverse design of FeNiCrCoCu MPEAs and unlocking key insights with explainable AI, *npj Comput. Mater.*, 2025, **11**, 124, DOI: [10.1038/s41524-025-01600-x](https://doi.org/10.1038/s41524-025-01600-x).
- 3 A. Beadle, *How Is AI Accelerating the Discovery of New Materials?*, <https://www.technologynetworks.com/applied-sciences/articles/how-is-ai-accelerating-the-discovery-of-new-materials-394927>, 2025.
- 4 H. Park, *et al.*, A generative artificial intelligence framework based on a molecular diffusion model for the design of metal-organic frameworks for carbon capture, *Commun. Chem.*, 2024, **7**, 21, DOI: [10.1038/s42004-023-01090-2](https://doi.org/10.1038/s42004-023-01090-2).
- 5 A. D. Sendek, *et al.*, Machine Learning-Assisted Discovery of Solid Li-Ion Conducting Materials, *Chem. Mater.*, 2019, **31**, 342–352, DOI: [10.1021/acs.chemmater.8b03272](https://doi.org/10.1021/acs.chemmater.8b03272).
- 6 A. Merchant, *et al.*, Scaling deep learning for materials discovery, *Nature*, 2023, **624**, 80–85, DOI: [10.1038/s41586-023-06735-9](https://doi.org/10.1038/s41586-023-06735-9).
- 7 S. H. Aboutaleb, Ensuring Data Integrity in AI-Driven Materials Science: Why F-Sum Rules and Kramers-Kronig Relations Matter, *Nanoscale Adv. Mater.*, 2025, **2**, 10–15, DOI: [10.22034/nsam.2025.01.02](https://doi.org/10.22034/nsam.2025.01.02).
- 8 C. Qin, *et al.*, SciHorizon: Benchmarking AI-for-Science Readiness from Scientific Data to Large Language Models, *arXiv*, 2025, preprint arXiv:2503.13503, DOI: [10.48550/arXiv.2503.13503](https://doi.org/10.48550/arXiv.2503.13503).



- 9 J. Leeman, *et al.*, Challenges in High-Throughput Inorganic Materials Prediction and Autonomous Synthesis, *Phys. Rev. X*, 2024, **3**, 011002, DOI: [10.1103/PRXEnergy.3.011002](https://doi.org/10.1103/PRXEnergy.3.011002).
- 10 N. J. Szymanski, *et al.*, An autonomous laboratory for the accelerated synthesis of novel materials, *Nature*, 2023, **624**, 86–91, DOI: [10.1038/s41586-023-06734-w](https://doi.org/10.1038/s41586-023-06734-w).
- 11 M. Christmann, What I Learned from Analyzing Accurate Mass Data of 3000 Supporting Information Files, *Org. Lett.*, 2025, **27**, 4–7, DOI: [10.1021/acs.orglett.4c03458](https://doi.org/10.1021/acs.orglett.4c03458).
- 12 A. L. Spek, What makes a crystal structure report valid?, *Inorg. Chim. Acta*, 2018, **470**, 232–237, DOI: [10.1016/j.ica.2017.04.036](https://doi.org/10.1016/j.ica.2017.04.036).
- 13 B. H. Toby, R factors in Rietveld analysis: How good is good enough?, *Powder Diff.*, 2006, **21**, 67–70, DOI: [10.1154/1.2179804](https://doi.org/10.1154/1.2179804).
- 14 L. B. McCusker, R. B. Von Dreele, D. E. Cox, D. Louer and P. Scardi, Rietveld refinement guidelines, *J. Appl. Crystallogr.*, 1999, **32**, 36–50, DOI: [10.1107/S0021889898009856](https://doi.org/10.1107/S0021889898009856).
- 15 D. B. Resnik, M. Hosseini, J. J. H. Kim, G. Epiphaniou and C. Maple, GenAI synthetic data create ethical challenges for scientists. Here's how to address them, *Proc. Natl. Acad. Sci. U. S. A.*, 2025, **122**, e2409182122, DOI: [10.1073/pnas.2409182122](https://doi.org/10.1073/pnas.2409182122).
- 16 J. Gu, *et al.*, AI-enabled image fraud in scientific publications, *Patterns*, 2022, **3**, 100511, DOI: [10.1016/j.patter.2022.100511](https://doi.org/10.1016/j.patter.2022.100511).
- 17 J. A. Byrne, *et al.*, A call for research to address the threat of paper mills, *PLoS Biol.*, 2024, **22**, e3002931, DOI: [10.1371/journal.pbio.3002931](https://doi.org/10.1371/journal.pbio.3002931).
- 18 E. M. Bik, A. Casadevall and C. Fang Ferric, The Prevalence of Inappropriate Image Duplication in Biomedical Research Publications, *mBio*, 2016, **7**(3), e00809.
- 19 M. A. Green, Solar cell fill factors: General graph and empirical expressions, *Solid-State Electron.*, 1981, **24**, 788–789, DOI: [10.1016/0038-1101\(81\)90062-9](https://doi.org/10.1016/0038-1101(81)90062-9).
- 20 T. Shinagawa, A. T. Garcia-Esparza and K. Takanabe, Insight on Tafel slopes from a microkinetic analysis of aqueous electrocatalysis for energy conversion, *Sci. Rep.*, 2015, **5**, 13801, DOI: [10.1038/srep13801](https://doi.org/10.1038/srep13801).
- 21 A. C. Lazanas and M. I. Prodromidis, Electrochemical Impedance Spectroscopy—A Tutorial, *ACS Meas. Sci. Au*, 2023, **3**, 162–193, DOI: [10.1021/acsmeasuresciau.2c00070](https://doi.org/10.1021/acsmeasuresciau.2c00070).
- 22 MIT Economics, Assuring an accurate research record, <https://economics.mit.edu/news/assuring-accurate-research-record>, 2025.
- 23 G. H. Major, *et al.*, Assessment of the frequency and nature of erroneous x-ray photoelectron spectroscopy analyses in the scientific literature, *J. Vac. Sci. Technol., A*, 2020, **38**, 061204, DOI: [10.1116/6.0000685](https://doi.org/10.1116/6.0000685).
- 24 M.-H. Van, P. Verma, C. Zhao and X. Wu, A Survey of AI for Materials Science: Foundation Models, LLM Agents, Datasets, and Tools, *arXiv*, 2025, preprint arXiv:2506.20743, DOI: [10.48550/arXiv.2506.20743](https://doi.org/10.48550/arXiv.2506.20743).
- 25 UK Government, *Government Digital Service*, 2025.
- 26 European Commission, *Directorate-General for Research and Innovation*, Brussels, 2025.
- 27 U. Peters, Science Based on Artificial Intelligence Need not Pose a Social Epistemological Problem, *Social Epistemology Review and Reply Collective*, 2024, **13**, 58–66.
- 28 Springer Nature, *Artificial Intelligence (AI)*, <https://www.nature.com/nmat/editorial-policies/ai>, 2025.
- 29 H. H. Thorp, Genuine images in 2024, *Science*, 2024, **383**, 7, DOI: [10.1126/science.adn7530](https://doi.org/10.1126/science.adn7530).
- 30 Freie Universität Berlin's Department of Biology, Chemistry, Pharmacy, *Freie Universität Berlin's Department of Biology, Chemistry, Pharmacy*, 2025.
- 31 J. Wilkinson, *et al.*, Protocol for the development of a tool (INSPECT-SR) to identify problematic randomised controlled trials in systematic reviews of health interventions, *BMJ Open*, 2024, **14**, e084164, DOI: [10.1136/bmjopen-2024-084164](https://doi.org/10.1136/bmjopen-2024-084164).

