

EDITORIAL

Trustworthy Intelligence: Reliability, Transparency, and Responsibility in Soft Computing

Iqbal H. Jebril^{1,*} and Tarek Kanan²

¹ Editor-in-Chief, *International Journal of Advances in Soft Computing and its Applications (IJASCA)*; Faculty of Sciences and Information Technology, Al-Zaytoonah University of Jordan, Amman 11733, Jordan.

² Department of Data Science and Artificial Intelligence, Faculty of Sciences and Information Technology, Al-Zaytoonah University of Jordan, Amman 11733, Jordan.

* Corresponding author: i.jebril@zuj.edu.jo. Contributing author: tarek.kanan@zuj.edu.jo.

Received: 1 September 2026 **Accepted:** 6 October 2026 **Published:** 18 October 2026

DOI: 10.15849/ijasca.239 **License:** Open access under CC BY 4.0. Authors retain copyright.

Publisher: Al-Zaytoonah University of Jordan (ZUJ), Amman, Jordan.

Cite this article: I. H. Jebril and T. G. Kanan, "Trustworthy Intelligence: Reliability, Transparency, and Responsibility in Soft Computing," *Int. J. Advance Soft Compu. Appl.*, vol. 18, no. 3, pp. i–vii, 2026, doi: <https://doi.org/10.15849/ijasca.239>.

Abstract

Advances in artificial intelligence and soft computing are frequently assessed through improvements in predictive accuracy, yet numerical performance alone provides an incomplete basis for judging scientific quality and practical value. This editorial examines the broader responsibilities involved in developing, evaluating, and publishing intelligent systems. Drawing on established literature and contributions featured in this issue of the *International Journal of Advances in Soft Computing and its Applications*, it discusses methodological rigor, reproducibility, generalization, explainability, fairness, and computational efficiency as complementary dimensions of trustworthy research. Attention is given to the justification of hybrid architecture, the importance of fair experimental comparisons, and the need to communicate uncertainty and application limitations. The editorial also addresses generative AI assistance in research and manuscript preparation, emphasizing transparent disclosure and continuing human accountability. Similarity and AI-writing reports are considered supporting instruments whose outputs require contextual interpretation rather than automatic judgments about originality or misconduct. Finally, the editorial considers the role of journals in strengthening research practice through constructive peer review, consistent standards, and interdisciplinary dialogue, while introducing a continuing series on soft computing applications. It argues that sustainable scientific progress depends on contributions that deepen understanding and demonstrate reliable, proportionate benefits under clearly stated conditions.

Keywords: Soft Computing; Trustworthy Artificial Intelligence; Research Reproducibility; Explainable AI; Algorithmic Fairness; Research Integrity; Generative AI; Editorial Responsibility.

1 Introduction

When a paper on artificial intelligence reports a higher accuracy score, that is usually the first thing we notice — and it should not be the last. Where did the improvement come from? Would it hold up if the data or the operating conditions changed? Could another researcher reproduce it? And would the people affected by the system's decisions understand what those decisions rest on, or where they might go wrong?

Questions like these sit at the heart of soft computing research today. Intelligent methods now reach into healthcare, education, cybersecurity, engineering, and many other fields. That reach gives us good reason to be ambitious, but it also puts real weight on the people who build, evaluate, and publish these systems.

Soft computing was never meant to pretend uncertainty away. Zadeh's early account of fuzzy logic, neural networks, and soft computing [1] treated incomplete information and approximation as part of the problem, not an inconvenience to be smoothed over. More than three decades later, our models are

larger and our optimization methods more sophisticated, but the need to reason carefully about uncertainty has not gone anywhere.

Writing this as Editor-in-Chief of IJASCA, I keep returning to a simple question: what do we value in a scientific contribution? The papers in this issue, Volume 18, Number 3, range across dataset selection, intrusion detection, explainable learning, and fairness in recruitment — different problems that keep circling back to the same one: what evidence makes an intelligent system worth trusting?

That is the question this editorial takes up, alongside a related one: how we use AI responsibly in research and in writing about research. It also marks the start of a continuing series of editorials meant to connect developments in soft computing with the needs of the fields that rely on it.

2 Scientific Quality Beyond Numerical Accuracy

In handling submitted manuscripts, I have noticed that the most useful discussion usually begins right where the performance table ends. Accuracy, precision, recall, and F1-score tell us something about a model, but not, by themselves, why the contribution matters.

A reported improvement might come from a genuinely better method. It might just as easily come from differences in preprocessing, parameter tuning, data partitioning, or simply more resources than the competing approach had. Authors need to show us which of these it is. A comparison only becomes persuasive once readers can see what was held constant, what changed, and why that difference is the one that counts.

The same scrutiny applies to hybrid architecture. There is nothing wrong with pairing a neural network with an evolutionary optimizer or a fuzzy inference system — plenty of good reasons exist for doing so. But the contribution lies in those reasons and in the evidence behind them, not in the combination itself. Authors should say plainly what each component adds and, where it makes sense, test that claim with an ablation study. When the performance gain is small or inconsistent, the added complexity deserves even closer scrutiny, not less.

We would encourage researchers to ask themselves, early on: what limitation are we addressing? Why would this approach overcome it? Are the baselines fair and suitably evaluated? How much do the results depend on experimental choices that could easily have gone another way? Does the benefit justify what it costs computationally? Framing a project around these questions from the start tends to produce a stronger manuscript later.

Two papers in this issue put these concerns into practice. *Genetic-Algorithm Feature Selection for Evaluating Self-Supervised Transformers and Gradient-Boosted Trees in Flow-Based Intrusion Detection* [2] compares models under shared feature sets and data splits. Under those conditions, XGBoost outperformed the self-supervised Transformer on UNSW-NB15, and transfer between datasets stayed unstable — a useful reminder to question both the value of added complexity and the limits of generalization.

Meta-Learning for Dataset Selection: An Adaptive Utility Scoring Framework with Interpretable Meta-Feature Analysis [3] takes on dataset selection through an interpretable utility score built from dataset characteristics. Its component analysis uncovers a trade-off between ranking correlation and how well the method recovers strong candidates among the top three selections. This provides a useful illustration of why the choice of evaluation criterion shapes how we judge a method's contribution in the first place.

For IJASCA, a meaningful advance can take several forms: a theoretical result, a reproducible experimental finding, a reduction in computational cost, or simply a clearer account of when an established method succeeds or fails. A carefully evaluated simple model can be a substantial contribution. So, just as easily, can a study that explains why an apparently promising approach does not generalize.

3 Explainability, Fairness, and Trustworthy Artificial Intelligence

Trustworthiness matters most when a system's output shapes decisions about people. A model can achieve strong aggregate performance while still making unreliable predictions for a particular group, or behaving unpredictably once it meets a new setting. And an explanation can look entirely plausible without actually reflecting how the model reached its output.

Rudin [4] makes a strong case for favoring inherently interpretable models in high-stakes decisions. Kelly and colleagues [5] walk through the validation and implementation hurdles that come up when AI is meant to benefit clinical practice. Ghassemi and colleagues [6] warn against placing too much confidence in today's explainability methods for healthcare. Read together, these three remind us to ask what an explanation actually establishes, and what it leaves untested.

So researchers need to be honest about what an explanation is for and who it is meant to serve. Does it help a specialist spot where the model went wrong? Does it hold up under small, relevant changes to the input? Does it communicate uncertainty in a way that is useful? A visualization or a ranking of influential variables can support that work, but its presence alone proves nothing about reliability.

Fairness calls for the same kind of care. Authors should name the groups, outcomes, and evaluation criteria that matter for their application, and be upfront about the limits of their data. A good overall score can hide differences that matter a great deal to the people the decision affects.

This issue includes two papers that work through exactly these tensions: *Explainable Multimodal Artificial Intelligence Framework for Real-Time Cognitive Engagement Assessment in Smart Learning Environments* [7] and *FairHR-X: An Explainable and Bias-Aware Artificial Intelligence Framework for Transparent Employee Recruitment and Selection* [8]. Both place prediction, explanation, and the consequences of automated assessment in concrete educational and organizational settings.

Soft computing has real resources to offer here. Fuzzy systems can express approximate relationships, neural networks can learn complex patterns, and evolutionary methods can help with feature selection and optimization. Combining them can pay off when the design is well justified and the resulting behavior is evaluated carefully — including an honest account of where the system should not be relied upon.

4 The Future of Soft Computing: Mathematics, Learning, and Optimization

All the attention deep learning receives should also push us back toward the broader mathematical and computational foundations of the field. Soft computing draws its strength from the interplay of fuzzy reasoning, neural computation, evolutionary methods, probabilistic approaches, and optimization — each bringing its own assumptions and its own way of handling uncertainty and complexity.

There is still substantial work to be done on convergence, stability, how we represent uncertainty, and generalization, and these questions only get harder once several methods are combined. If an evolutionary optimizer improves a neural model, we need to understand how much of that result depends on initialization, how much tuning went into it, and what it cost computationally. Where formal guarantees are not available, authors should at least state clearly how far the empirical evidence actually reaches.

Efficiency deserves a place in the main argument of a paper, not a footnote. For systems that have to run under real constraints on memory, processing power, or response time, computational cost is often what decides whether a method is usable at all. Reporting that cost alongside predictive performance lets readers judge whether a proposed improvement is worth adopting in practice.

We would like to see more work that connects mathematical reasoning with experimental design and real application needs. Progress does not have to mean a new algorithm. It can just as well mean explaining an existing method more carefully, pinning down a failure condition, or simplifying a system without losing the capabilities that matter. That kind of work gives the rest of us something sturdier to build on.

5 AI-Assisted Research and Writing: Encouraging Honest Disclosure

The growing use of generative AI in research raises a related question of trust. Researchers now use these tools to revise language, help with programming, organize material, or check computational procedures. These uses need clearer guidance, especially since many authors are simply unsure what journals expect them to disclose.

Some worry that admitting to AI assistance will count against their manuscript. That concern deserves to be taken seriously by editors. Expectations that are vague, or applied inconsistently from one journal to the next, tend to discourage exactly the openness that responsible publishing depends on.

At IJASCA, we see real value in responsibly used AI tools, and we would rather authors told us how they used them than hid it. Legitimate AI assistance is not, on its own, grounds for rejection. What matters is the nature of the assistance, whether the resulting work is reliable, and whether the authors can stand behind it.

There is an obvious difference between using AI to tidy up a manuscript's language and using it to fabricate results or invent references that do not exist. Assistance with code is no different in principle: accepting a plausible-looking output is not enough. The researcher still has to understand the procedure, verify how it behaves, and judge whether it fits the study. None of that responsibility transfers to the tool.

COPE's position on authorship and AI tools [9] states the underlying principle plainly: AI tools cannot hold authorship responsibilities, and the human authors remain accountable for the manuscript — including the data, the mathematical expressions, the algorithms, the interpretations, and the references, whatever assistance went into producing them.

What disclosure needs to do is give readers enough to judge the tool's role for themselves. Where AI materially shaped a research procedure, that explanation has to be detailed enough for the procedure to be evaluated. For language assistance alone, a brief statement naming the tool and describing the kind of revision involved is usually enough, in line with the journal's requirements.

We would welcome a clear account of responsible AI use in any manuscript we receive. Authors should feel able to give one without it reading as a confession. Our job, as editors, is to build the conditions in which honest reporting is expected, understood, and judged fairly.

6 Screening Reports and Editorial Judgment

The rise of AI-assisted writing has naturally pushed interest toward automated screening. These tools can surface something worth a closer look, but their output still needs a human to interpret it.

As part of its editorial process, IJASCA runs Turnitin's Similarity Report and AI Writing Report on submissions. Both reports go to the scientific reviewers through the journal's Open Journal Systems (OJS) platform, alongside the manuscript itself and whatever the authors have disclosed.

The two reports do different jobs. The Similarity Report flags text that matches other sources, and that overlap has to be read in context — quotations, references, standard terminology, and descriptions of well-established procedures can all produce matches without amounting to plagiarism. The AI Writing Report flags text the system estimates may have been generated, or modified, by AI. It does not tell us who actually wrote a passage, and it certainly does not establish that misconduct occurred. Turnitin itself acknowledges that the tool can misclassify text and warns against using its AI indicator as the sole basis for any adverse action [10].

So a screening result should raise a specific, answerable question rather than settle anything on its own; writing-provenance flags and research-validity concerns are different questions, and a text-pattern indicator should not, by itself, be read as evidence about whether a method is sound. Unattributed overlap means checking the matched source. A substantive concern about a research procedure should prompt a request for clarification or supporting evidence, independently of any AI-writing indicator.

Neither a similarity percentage nor an AI-writing score can, by itself, decide whether a study is original, valid, or honest.

That judgment still rests with reviewers, who bring the subject knowledge needed to weigh the research question, the methods, the evidence, and the conclusions. Where something remains unclear, they can ask authors to explain how AI was used, clarify an experimental step, or supply what is needed to assess reproducibility — and authors deserve a fair chance to respond to concerns that are clearly stated.

The final decision stays with the Editor-in-Chief or the handling editor assigned to the manuscript, weighing the manuscript itself, the reviewers' assessments, the authors' explanations, and the journal's standards together. Screening technology earns its place only to the extent that it supports this process rather than substituting for it.

7 Journals as Partners in Research Development

A journal shapes its field through the questions it encourages and the standards it holds authors to — and its influence does not stop at the decision on any one manuscript.

That matters especially in soft computing, where useful work often depends on researchers from different disciplines talking to each other. A mathematically elegant method can still miss its application if its assumptions do not fit the problem. An encouraging application result, in turn, can be hard to interpret without a solid account of the method behind it and where that method falls short. Bringing these two sides together takes sustained dialogue among researchers, reviewers, and editors — it does not happen on its own.

Peer review is where a lot of that dialogue happens. A useful review names a weakness precisely, explains why it matters, and points to what evidence or clarification would resolve it. Authors, for their part, strengthen their work by engaging seriously with criticism rather than deflecting it, and by explaining the reasoning behind how they respond.

Reviewers put real time and expertise into this process, and that deserves recognition — as does the care authors take when they revise their work thoughtfully rather than defensively. Editors, meanwhile, are responsible for keeping standards consistent and making sure the back-and-forth stays focused on the science rather than drifting elsewhere.

These are human responsibilities, and they stay central even as the tools around publishing keep changing. In the end, a journal's quality comes down to how carefully its community evaluates and communicates what it knows.

8 A Continuing Editorial Series on Soft Computing Applications

Starting with this issue, IJASCA is launching a continuing series of editorial perspectives on how soft computing and artificial intelligence are actually applied. Each forthcoming issue will carry an editorial built around one application area, working through the questions and opportunities that area raises for further research.

The aim is a focused perspective, not an exhaustive review. Likely topics include healthcare, cybersecurity, intelligent education, sustainable energy, transportation, and industrial optimization, with other areas added as new needs and developments come up.

Within each area, the discussion will tie application challenges back to questions that matter across soft computing: how to reason with incomplete information, how to learn from imbalanced data, how to optimize under uncertainty, how to design hybrid systems that stay efficient, and how to provide decision support that users can evaluate and understand.

Our hope is that these editorials open up conversations between the researchers developing methods and the specialists who understand the settings where those methods get used. Application problems can expose gaps that deserve real mathematical attention, and theoretical insight can, in turn, point to better ways of handling practical difficulties.

The featured topic in any given issue will not narrow its scope or give any manuscript an advantage in peer review or acceptance. IJASCA continues to welcome submissions across its full aims and scope, and every manuscript is assessed against the same editorial standards regardless of the issue's featured theme.

We invite readers to engage with the questions this series raises, and to bring forward work that tests assumptions, tackles unresolved problems, and opens up useful directions for the field.

9 Recommendations

The preceding discussion points to a short set of practical measures for strengthening the evidence behind intelligent systems and the integrity of scientific communication, to be applied in proportion to the purpose, context, and potential consequences of each study.

Establish the contribution. Identify the limitation being addressed and justify the proposed approach; for hybrid systems, explain each component's role and test it with an ablation study where appropriate.

Design fair, reproducible comparisons. Use consistent data partitions, preprocessing, and evaluation criteria; disclose tuning choices; and share data, code, and configurations where feasible.

Evaluate reliability beyond aggregate scores. Report variability across runs, sensitivity to experimental choices, and uncertainty in reported gains, and state claims in proportion to the conditions actually tested.

Assess explanations and fairness in context. Specify who an explanation serves and what it supports; examine subgroup performance where data permit, and treat insufficient data as a limitation rather than evidence of fairness.

Report efficiency and operational constraints. Disclose computational cost alongside predictive performance, and justify whether the benefit warrants the resources required.

Make AI assistance transparent. Disclose substantive generative-AI use and name the tools, and keep human authors fully responsible for verifying generated text, code, and results.

Use screening reports to inform, not decide. Treat similarity and AI-writing indicators as prompts for specific clarification, never as independent proof of misconduct, and give authors a fair opportunity to respond.

Strengthen review and interdisciplinary exchange. Ground criticism in identifiable weaknesses, recognize well-supported negative results, and connect mathematical and applied perspectives.

10 Looking Ahead

The pace of AI development gives researchers plenty of reasons for optimism, and it also makes careful evaluation more valuable, not less. As new architectures and applications keep appearing, we need to tell apart the advances that genuinely deepen understanding or improve practice from the gains that only hold up under one narrow experimental setup.

For soft computing, that means continuing to ask whether our methods are justified, whether our comparisons are fair, whether our findings reproduce, and whether our conclusions are proportionate to the evidence behind them. It also means making room for results that reveal limitations, clarify uncertainty, or simply show that a simpler approach was good enough all along.

That same standard of responsibility should guide how we use AI while doing research and preparing manuscripts. Technological assistance can genuinely help, and authors should describe it openly while still taking full responsibility for the work they submit.

Our hope is that IJASCA stays a place where innovation gets careful consideration, honest disclosure is welcomed rather than penalized, and constructive criticism helps research mature. The editorial series

launched here is one way of keeping that conversation going across the many communities that make up soft computing.

We thank the authors in Volume 18, Number 3, and the reviewers and editorial board members whose work made this issue possible. Their care, expertise, and willingness to question and improve scientific work are what the journal actually runs on. We look forward to the ideas and discussions our readers bring to future issues.

DECLARATIONS

Funding. This editorial received no external funding.

Conflicts of interest. I. H. Jebril serves as Editor-in-Chief of IJASCA. To avoid a conflict of interest arising from this role, the review and acceptance of this editorial were handled independently by another member of the editorial board. Beyond this, the authors declare no other competing interests.

Ethics approval. Not applicable. This is an editorial essay and does not involve human participants, animals, or primary data collection.

Informed consent. Not applicable.

Data and code availability. Not applicable. This editorial generated no new data or code.

Author contributions. I. H. Jebril and T. G. Kanaan jointly conceived and wrote this editorial. Both authors reviewed and approved the final version.

Use of artificial intelligence. Generative AI tools (ChatGPT and Claude) were used during the preparation of this editorial to assist with language editing and revision, to support searches for background information, and in discussing specific points addressed in the text. The authors directed and reviewed all AI-assisted input, verified its accuracy, and take full responsibility for the final ideas, analysis, conclusions, and manuscript content.

Acknowledgements. The authors gratefully acknowledge Al-Zaytoonah University of Jordan for providing an academic environment that supports scientific research, scholarly exchange, and responsible innovation. We extend our appreciation to the editorial board and reviewers of the International Journal of Advances in Soft Computing and its Applications for their dedication to scientific quality and constructive dialogue. We also thank the researchers, colleagues, and contributors whose insights and scholarly work informed the perspectives presented in this editorial.

REFERENCES

- [1] L. A. Zadeh, "Fuzzy logic, neural networks, and soft computing," *Communications of the ACM*, vol. 37, no. 3, pp. 77–84, 1994. <https://doi.org/10.1145/175247.175255>
- [2] M. F. Al-hammouri, B. Vamsi, M. M. AlMi'ani, E. Al Bataineh, and A. Al Bataineh, "Genetic-Algorithm Feature Selection for Evaluating Self-Supervised Transformers and Gradient-Boosted Trees in Flow-Based Intrusion Detection," *International Journal of Advances in Soft Computing and its Applications*, vol. 18, no. 3, pp. 296–316, 2026. <https://doi.org/10.15849/ijasca.166>
- [3] F. AL-Saqqar, A. H. Alkilan, M. I. Nusir, W. AlQassas, and M. El-Bashir, "Meta-Learning for Dataset Selection: An Adaptive Utility Scoring Framework with Interpretable Meta-Feature Analysis," *International Journal of Advances in Soft Computing and its Applications*, vol. 18, no. 3, pp. 87–112, 2026. <https://doi.org/10.15849/ijasca.160>
- [4] C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," *Nature Machine Intelligence*, vol. 1, pp. 206–215, 2019. <https://doi.org/10.1038/s42256-019-0048-x>
- [5] C. J. Kelly, A. Karthikesalingam, M. Suleyman, G. Corrado, and D. King, "Key challenges for delivering clinical impact with artificial intelligence," *BMC Medicine*, vol. 17, article 195, 2019. <https://doi.org/10.1186/s12916-019-1426-2>
- [6] M. Ghassemi, L. Oakden-Rayner, and A. L. Beam, "The false hope of current approaches to explainable artificial intelligence in health care," *The Lancet Digital Health*, vol. 3, no. 11, pp. e745–e750, 2021. [https://doi.org/10.1016/S2589-7500\(21\)00208-9](https://doi.org/10.1016/S2589-7500(21)00208-9)
- [7] A. A. Alhamadi, "Explainable Multimodal Artificial Intelligence Framework for Real-Time Cognitive Engagement Assessment in Smart Learning Environments," *International Journal of Advances in Soft Computing and its Applications*, vol. 18, no. 3, pp. 34–71, 2026. <https://doi.org/10.15849/ijasca.141>
- [8] A. N. Aloqaily, "FairHR-X: An Explainable and Bias-Aware Artificial Intelligence Framework for Transparent Employee Recruitment and Selection," *International Journal of Advances in Soft Computing and its Applications*, vol. 18, no. 3, pp. 140–156, 2026. <https://doi.org/10.15849/ijasca.184>
- [9] COPE Council, "Authorship and AI tools," *Committee on Publication Ethics*, COPE Position Statement, 2024. <https://doi.org/10.24318/cCVRZBms>
- [10] Turnitin, "Using the AI Writing Report," *Turnitin Guides*, accessed October 12, 2026. <https://guides.turnitin.com/hc/en-us/articles/22774058814093-Using-the-AI-Writing-Report>