Editorial

We are pleased to present the final issue of Volume 17 of the ***Journal of Intelligent Computing*****,** which brings together three compelling contributions addressing timely challenges at the intersection of artificial intelligence, reliability, and geospatial analytics.

Large language models (LLMs) frequently produce outputs that are plausible yet factually incorrect, a phenomenon widely known as hallucination. While this problem has attracted considerable scholarly attention, it is often examined through a narrow lens. In **"Multidimensional Hallucination Risk Across AI Models: Frequency, Confidence, Accuracy Mismatch, Detection, and Trust Impact,"** the authors reframe hallucination as a multidimensional risk that spans generation frequency, confidence accuracy mismatch, detectability, and downstream trust impact. Drawing on a comprehensive dataset of 10,000 interactions across six frontier AI models and platforms, they assess hallucination prevalence, epistemic calibration, and human consequences. Their findings reveal that AI hallucination constitutes a systemic, cross platform reliability challenge. Mitigating its impact, they argue, demands integrated strategies that address confidence calibration, proactive detection mechanisms, and user interface designs that encourage appropriate skepticism rather than relying solely on reducing raw factual error rates.

Building on this theme of LLM trustworthiness, the second paper, **"Prompt Clarity, Hallucination Risk, and Trustworthy LLM Performance: Mediation Analysis, Cross-Lingual/Domain Robustness, and Routing Gains,"** investigates how prompt clarity shapes response quality and reliability, with particular attention to the mediating role of hallucination risk. The authors construct a synthetic benchmark of 15,000 prompt response pairs spanning five languages, multiple domains, four difficulty levels, and five model families, and apply causal style mediation analysis, fairness audits, and routing policy simulations. Their results demonstrate that hallucination risk serves as the dominant mediator, accounting for approximately 66% to 71% of the total effect of prompt clarity on response quality a striking finding with practical implications for prompt engineering and model deployment.

Shifting from language models to geospatial intelligence, the third paper, **"Machine Learning Classification of Construction Site Annotations from Geospatial Metadata: Near Perfect Binary Discrimination and the Limits of Construction-Type Recognition,"** evaluates the predictive power of geospatial metadata including polygon geometry, temporal duration, and annotation attributes using Random Forest, Extra Trees, and Logistic Regression models applied to a large scale GeoJSON dataset of construction sites. The authors conduct both binary (site presence) and multiclass (construction type) classification tasks, uncovering a stark contrast: while Random Forest achieves near-perfect binary accuracy (99.9%), multiclass construction type recognition remains highly ambiguous (macro-F1 = 0.610). They conclude that metadata-only models are insufficient for fine grained semantic typing and issue a critical methodological warning for the GeoAI community: rigorous spatial cross-validation and annotation leakage audits are essential to prevent spatially and procedurally inflated performance metrics.

Together, the papers in this issue reflect the vibrant and rapidly evolving research landscape in intelligent computing, tackling questions of model reliability, interpretability, and methodological rigor that will remain central to the field for years to come.

Editors