Editorial

We are pleased to present the third issue of the sixteenth volume of the ***Journal of Data Processing,*** featuring the breakthrough research outlined below.

The opening paper, "**Latency and Factual Accuracy in Generative AI: An Empirical Analysis of Model Performance, User Satisfaction, and the Role of Prompt Complexity**," proposes an integrated empirical framework that incorporates efficiency, factual accuracy, and user satisfaction to analyze their combined influence on overall Generative AI (GenAI) performance. The results indicate that prompt length is the dominant positive predictor of response latency, significantly outweighing the impact of decoding hyperparameters such as temperature and Top-P. Among the evaluated models, Claude 3.7 demonstrated the optimal balance, achieving the highest success rate and user satisfaction alongside a comparatively low hallucination rate. This study underscores that concise prompt engineering and task-specific model selection are more critical for optimizing Large Language Model (LLM) deployment than model identity alone, providing actionable insights for the development of reliable, user-centric AI systems.

The second paper, "**Clinical Symptom Profiles and Phenotypic Clusters in Familial vs. Sporadic Neurofibromatosis Type 1: A Comprehensive Exploratory Analysis of the UCI NF1 Dataset**," investigates clinical symptom distributions and identifies latent phenotypic clusters using the UCI NF1 dataset. To characterize phenotypic variability, the authors employ correlation analyses, symptom co-occurrence networks, UpSet plots, and hierarchical clustering. Furthermore, they propose a conceptual, multi-layered clinical decision support framework that integrates machine learning and explainable artificial intelligence (XAI) to translate these exploratory insights into personalized precision medicine. The authors note that future prospective validation incorporating multimodal data is warranted to refine these predictive models and ultimately improve long-term clinical outcomes for NF1 patients.

The final paper, "**A Statistically Validated Large-Scale Mathematical Chain of Thought Dataset for Symbolic Reasoning and Large Language Model Benchmarking**," presents a statistically validated, large-scale mathematical Chain of Thought (CoT) dataset. This corpus comprises approximately 100,000 samples spanning six core mathematical domains (Algebra, Geometry, Trigonometry, Matrices, Derivatives, and Integrals) across three predefined difficulty levels. The dataset preserves complete intermediate reasoning traces and final solutions in standardized LaTeX notation. The authors demonstrate that the corpus exhibits rich symbolic diversity, authentic mathematical syntax, and structured command co-occurrence networks. Ultimately, this statistically validated dataset provides a reliable and robust benchmark for supervised fine-tuning, curriculum learning, process supervision, explainable AI, and the evaluation of next generation mathematical LLMs.

The research published in this issue reflects the ongoing advancements and diverse applications within the field of data science.

Editors