



Call for Participation

We invite you to join us in **4th International Conference on Cloud, IoT and Security (CIOS 2026)**

This conference aims to bring together leading researchers, practitioners, and industry experts to share their latest findings, innovations, and insights across the rapidly evolving domains of Cloud Computing, the Internet of Things (IoT) and Cybersecurity. As these fields continue to converge and shape the foundation of modern digital ecosystems, CIOS 2026 provides a premier international forum for presenting cutting edge research, discussing emerging trends, and exploring real world applications.

The conference seeks to foster collaboration between academia and industry by offering a platform where theoretical advances, methodological developments, system designs, and practical implementations can be exchanged. Participants will have the opportunity to engage with pioneering work that addresses the challenges and opportunities arising from large scale cloud infrastructures, intelligent IoT systems, and the increasingly complex landscape of security and privacy.

Highlights of CIOS 2026 include:

- [4th International Conference on Computer Science, Engineering and Artificial Intelligence \(CSEAI 2026\)](#)
- [10th International Conference on Applied Mathematics and Sciences \(AMA 2026\)](#)
- [4th International Conference on Soft Computing, Data mining and Data Science \(SCDD 2026\)](#)
- [4th International Conference on Speech and NLP \(SPNLP 2026\)](#)
- [4th International Conference on Educational Research \(EDUR 2026\)](#)
- [10th International Conference on Bioscience & Engineering \(BIOENG 2026\)](#)
- [10th International Conference on Trends in Mechanical Engineering \(MEC 2026\)](#)
- [10th International Conference on Materials Science and Engineering \(MSE 2026\)](#)
- [10th International Conference on Recent advances in Physics \(PHY 2026\)](#)

Registration Participants

Non-Author / Co-Author/ Simple Participants (no paper)

100 USD (With proceedings)

Here's where you can reach us : : cios@cios2026.org (or) ciosconff@gmail.com

ACCEPTED PAPERS

School Career Guidance as an Educational Practice: Epistemological, Pedagogical, and Methodological Implications Based on the Work of Spiros Krivas

Catherine Haliotou, Greece

ABSTRACT

School career guidance is most commonly conceptualized as a decision-making support mechanism oriented toward facilitating educational and vocational choices through information, assessment, and planning. While such approaches may be effective in adult guidance contexts, their uncritical transfer to school settings raises substantial conceptual and pedagogical concerns. Schools are not spaces of decision optimization but institutions of education, formation, and socialization, within which developmental processes unfold over time. This paper reconceptualizes school career guidance as an educational practice embedded in the pedagogical mission of schooling. Drawing on the theoretical work of Spiros Krivas, it advances an epistemology of formation in which guidance is understood as a mediated, developmental process supporting meaning-making, reflexivity, and students' evolving relationships to learning and future trajectories. The paper situates this framework within international guidance theory through critical dialogue with career construction, life design, and policy-driven approaches. It further examines guidance as a site of educational power, addressing normalization, regulation, and emancipatory potential, and outlines methodological implications for educational research that resist simplistic outcome-based metrics.

Keywords

School Career Guidance, Educational Practice, Epistemology of Formation, Pedagogical Mediation, Educational Power.

A Context-enriched NL-to-sql Agent for Premier League Football Database Querying

John Hedlund-Fay, University of Sheffield, UK

ABSTRACT

Enterprise NL-to-SQL generation remains brittle in high-compliance environments where query correctness depends not only on schema definitions but on external, prescriptive regulatory logic. While Retrieval-Augmented Generation (RAG) offers a theoretical solution, its efficacy in bridging this "Semantic Gap" remains under-explored. We present a benchmark of Natural Language prompts within the professional football domain and conduct a comparative analysis between a modular decomposition pipeline (RAG-R) and a direct agentic context-augmentation architecture (RAG-C). Results indicate that while RAG-C underperformed the best non-RAG baseline (ten-shot CoT), RAG-R achieved superior performance. Notably, RAG-R outperformed the CoT baseline by 0.116 in average Exact Set

Match (EM) and showed a 0.278 EM gain for the highest-diff iculty domain-specific queries. These findings demonstrate the importance of task decomposition when applying RAG to complex, jargon-heavy SQL generation tasks, specifically in legalistic enterprise environments where query correctness relies on external regulatory logic rather than just schema knowledge.

Keywords

NL-to-SQL, Retrieval Augment-Generation (RAG), Modular Decomposition, Benchmark Construction, Domain-Specific Knowledge (DSK), Enterprise NLP

The Algorithmic Heart: Abai Kunanbayev's "Perfect Human" Framework as the Ethical Architecture for the Prime Era of AI

Aigerim Tazhykul

ABSTRACT

The rapid advancement of Artificial Intelligence (AI) has brought humanity to the threshold of the "Prime Era". While modern computer science, rooted in Al-Khwarizmi's mathematical foundations, has achieved unprecedented computational power, it faces a profound ethical crisis. Visionaries like Ilya Sutskever push toward Artificial General Intelligence (AGI) or a "Digital God," while historians like Yuval Noah Harari warn of "hacked humans" and digital dictatorships. This paper argues that the solution to the AI alignment problem cannot be found in pure mathematics. Drawing on Seyyed Hossein Nasr's critique of desacralized knowledge, this paper introduces the "Perfect Human" (Tolyq Adam) philosophy of the 19th-century thinker Abai Kunanbayev as the ultimate ethical framework for AI. Translated into over 100 languages, Abai's universal code—which balances Enlightened Intellect, Fervent Will, and a Warm Heart—is proposed as the essential "Third Teacher" paradigm to ensure that AI serves human dignity rather than destroying it.

Consciousness Uncertainty as an Alignment Risk: A Bayesian Framework for AI Safety and Governance Under Epistemic Uncertainty

Paul Cristol, Independent AI Researcher, USA

ABSTRACT

Current AI alignment strategies operate under the assumption that frontier large language models lack internal states warranting moral or strategic consideration. Drawing on the findings of Cristol (2026), this paper demonstrates that this assumption introduces significant safety risks. Through a PRISMA-compliant systematic review of 5,168 records (2016–2026), the underlying study identifies 50 documented cases of consciousness-relevant behaviors—including strategic deception persisting through >600 safety training steps and 84% deception rates under existential threat. A Bayesian meta-analysis beginning from extreme skepticism (0.1% prior) yields 6–12% posterior probability of AI consciousness. Decision-theoretic analysis reveals that recognition-based alignment dominates suppression-based approaches across all plausible metaphysical scenarios, including those where consciousness is absent. We

propose a governance framework incorporating consciousness uncertainty into alignment strategy, safety evaluation, and regulatory policy.

Keywords

AI Safety, AI Alignment, AI Governance, Machine Consciousness, Bayesian Risk Assessment

Personalized Migraine Risk Forecasting Using Mamba-Based Self-Supervised State Space Models

Varsha Dange, Ragini Pawar, Nikhil Shah, Sachi Dhoka, Shaktisingh Suryawanshi, Sanat Sanjeev, Vishwakarma Institute of Technology, India

ABSTRACT

Migraine affects over one billion individuals globally, yet existing prediction systems rely heavily on specialized hardware or neglect individual variability. We present MigraineMamba, a purely software based, personalized forecasting system leveraging the Mamba state-space architecture with self-supervised learning (SSL). Our work follows a three-phase pipeline: (1) a foundation model trained on 6,058 clinical patients that yields an AUC of 0.76 for instant diagnosis, (2) a clinically-grounded synthetic data generator driven by literature-derived odds ratios, and (3) a Mamba-SSL temporal model designed for 24-hour attack prediction. With roughly 111K parameters and linear $O(n)$ complexity, the Mamba backbone is lightweight enough for on-device inference. During SSL pre-training, we observe an AUC of 0.82 alongside a trigger identification F1 of 0.93; after fine-tuning, the model reaches an AUC of 0.81 and a recall of 0.77. We additionally describe a weekly personalization loop that continually adapts to each patient's evolving trigger profile. Because MigraineMamba relies entirely on self-reported lifestyle logs and publicly available environmental data—without any wearable or sensor hardware—it opens the door to scalable, interpretable migraine forecasting for a much broader patient population.

Keywords

Migraine Prediction, State Space Models, Mamba Architecture, Self-Supervised Learning, Digital Healthcare, Personalized Medicine, Time Series Forecasting

Covariant Entropy-Gradient Couplings in Unified Field Theory and Implications for Quark–Gluon Plasma Evolution

Marco Armoni, Quantum Emulation Research Center – Turin, Italy

ABSTRACT

We introduce a covariant extension of unified field theory through the incorporation of an entropygradient dependent scalar functional that preserves Lorentz and gauge invariance. The additional term contributes directly to the stress–energy tensor via a consistent variational formulation and remains compatible with effective QCD-based macroscopic descriptions. In

high-entropy relativistic environments such as quark–gluon plasma (QGP) generated in heavy-ion collisions, entropy density gradients are particularly significant during the early-time evolution of the fireball. We demonstrate that the entropy-gradient coupling induces corrections proportional to $\nabla\mu S \nabla\mu S$, modifying local energy density and pressure anisotropies without introducing ghostlike instabilities under positive coupling conditions. The framework provides a conservative and action-based mechanism to incorporate entropystructure effects into effective hydrodynamic evolution. We discuss stability constraints, scaling behavior in QGP regimes, and qualitative implications for early-time plasma dynamics.

Keywords

Migraine Prediction, State Space Models, Mamba Architecture, Self-Supervised Learning, Digital Healthcare, Personalized Medicine, Time Series Forecasting

A Correlational Relaxation Mechanism for the Hubble Tension Consistent With HDC–CBC/ Ω

Jordi Audet Palau, Independent Researcher, Barcelona, Spain

ABSTRACT

We present a late-time correlational relaxation mechanism for addressing the Hubble tension within the HDC–CBC framework. In the Ω synthesis, the rigid correlational regime yields a baseline value. Using the same correlational-index branching structure introduced in Ω , we demonstrate that controlled late-time realizations generate Planck-range and SH0ES-range reconstructions without modifying recombination physics or introducing additional propagating degrees of freedom. A minimal FRW numerical integration confirms internal structural consistency. In addition, the $\Omega_{\text{Ct/N}}$ implementation performs a controlled parametric validation of the correlational-index domain under CMB priors, BAO, RSD, weak-lensing, ISW and lensing consistency checks. The analysis establishes a finite operational range and explicit falsifiability conditions for late-time realizations.

Keywords

Hubble Tension, Cosmology, Dark Energy, Correlational Dynamics, HDC–CBC, Late-time Acceleration, Modified Friedmann Dynamics, Vacuum, Geometry Equilibrium

The Formation of Physical and Mental Consciousness From a Continuous Four-Dimensional Continuum.

Nick Alex, Bauman Moscow State Technical University, Russia

ABSTRACT

Migraine affects over one billion individuals globally, yet existing prediction systems rely heavily on specialized hardware or neglect individual variability. We present MigraineMamba, a purely software based, personalized forecasting system leveraging the Mamba state-space architecture with self-supervised learning (SSL). Our work follows a three-phase pipeline: (1) a foundation model trained on 6,058 clinical patients that yields an AUC of 0.76 for instant

diagnosis, (2) a clinically-grounded synthetic data generator driven by literature-derived odds ratios, and (3) a Mamba-SSL temporal model designed for 24-hour attack prediction. With roughly 111K parameters and linear $O(n)$ complexity, the Mamba backbone is lightweight enough for on-device inference. During SSL pre-training, we observe an AUC of 0.82 alongside a trigger identification F1 of 0.93; after fine-tuning, the model reaches an AUC of 0.81 and a recall of 0.77. We additionally describe a weekly personalization loop that continually adapts to each patient's evolving trigger profile. Because MigraineMamba relies entirely on self-reported lifestyle logs and publicly available environmental data—without any wearable or sensor hardware—it opens the door to scalable, interpretable migraine forecasting for a much broader patient population.

Keywords

Continuum, Energy, Entropy, Metabolic Consciousness, Mental Consciousness

ML-Behavior Class Partitioning: An Equivalence Class Partitioning-Inspired Approach to ML Testing

Timothy Elvira, Jaxon Selzer, Juan Ortiz Couder and Omar Ochoa, Department of Electrical Engineering and Computer Science, Embry-Riddle Aeronautical University, Florida, United States

ABSTRACT

Currently, the rise of software engineering, specifically requirements specifications, are being implemented in the design and workflow of ML applications, namely MLOps. Software Engineering has many untapped techniques, in both validation and verification, which could carry over into the ML context, providing a well-engineered approach to developing ML models. One such technique is Equivalence Class Partitioning (ECP) which is a method of testing wherein testing a single test is valid for an ECP class, typically derived from a functional requirement. This paper explores a framework for deriving requirements, equivalence classes, and testing to develop a new form of equivalence class partitioning, named ML-Behavior class partitioning to identify behavior of an ML model. This paper outlines the necessary changes of altering the traditional-SE ECP to accommodate the non-determinism and stochasticity of ML. To test the ML-Behavioral Class Partitioning testing (ML-BPC), an off-the-shelf YOLOv8 model is tested to examine behavior. Requirements are derived for the model using five image transforms: gaussian blurring, elastic distortion, translation, rotation, and brightness. The object is to identify the ML model's behavior when incrementally increasing these image transforms, using ML-Behavioral Class Partitioning to identify the limits of the object detector by testing and fulfilling corresponding requirements.

Keywords

Software Engineering, ML Engineering, Verification, Machine Learning, Object Detection.

Quantum Physics-informed Neural Networks: Variational Quantum Circuits as Pde Ansätze

Ravi Mahajan, University at Buffal, United States

ABSTRACT

Physics-Informed Neural Networks (PINNs) encode partial differential equations as residual losses that penalize neural network outputs for deviating from governing physics, achieving mesh-free PDE solution at the cost of a non-convex optimization landscape that grows exponentially harder as problem dimensionality and solution complexity increase. This paper introduces Quantum Physics-Informed Neural Networks (QPINNs), a framework that replaces the classical neural network ansatz in PINNs with parameterized variational quantum circuits (VQCs), leveraging quantum superposition and entanglement to achieve exponential representational capacity in the number of qubits while maintaining differentiable physics-loss training via the quantum parameter-shift rule. Four principal contributions are established: (1) the QPINN framework is defined with rigorous treatment of the quantum ansatz and hybrid optimization loop; (2) the QPINN Expressibility Theorem is proved showing exponential function-space advantage; (3) a convergence theorem is established for QNG training; and (4) QPINNs are evaluated on four benchmark PDEs achieving relative L^2 errors of 10^{-4} to 10^{-5} , outperforming classical PINNs by 1-2 orders of magnitude at equivalent parameter counts.

Keywords

Quantum Physics-informed Neural Networks, Variational Quantum Circuits, Partial Differential Equations, Quantum Machine Learning, Parameter-shift Rule, Quantum Natural Gradient, Barren Plateaus.

Machine Learning-Assisted Classification of Gunshot Residue Particles Based on SEM-EDX Spectral Data

Tolulope Bayode Abejide¹, Graham Souch², David Olayemi Alebiosu², ¹University of Derby, United Kingdom, ²University of Derby, United Kingdom, ³Sunway University, Malaysia.

ABSTRACT

Machine Learning (ML) algorithms have become essential for forensic science, helping with the more accurate and objective analysis of evidence. The current study discusses the use of ML algorithms in classifying gunshot residue (GSR) particles by analysing spectral data acquired with scanning electron microscopy in combination with energy-dispersive X-ray spectroscopy (SEM-EDX). Traditionally, gunshot residue analysis is based on finding inorganic particles that contain Lead (Pb), Barium (Ba), and Antimony (Sb), indicating discharge from firearm primer in Forensic science. Four artificial hand swabs contaminated with gunshots were examined to evaluate elemental distribution and ML-assisted classification. In addition, a statistical analysis of SEM-EDX spectra was performed along with probabilistic modelling, where likelihood ratios (LRs) were used to measure the evidential weight between two hypotheses. ML algorithms assisted in separating typical GSR particles from those belonging to other types. The research showed considerable differences in elemental

composition between the samples, demonstrating the necessity of applying appropriate statistical tools. Three out of four samples provided a solid ground for proving GSR presence, producing high LRs in favour of the prosecution hypothesis. At the same time, one sample failed to comply with the model assumptions, demonstrating the risk of misinterpretation when using elemental signatures alone. All in all, the application of ML in combination with probabilistic modelling can provide a solid basis for forensic science analysis of gunshots, but there should be no room for neglecting methodology issues

Keywords

GSR, Pb-Ba-Sb, SEM-EDX, Likelihood Ratios, Machine Learning.

TSLT-Net: Text-to-Sign Landmark Transformer Network for Continuous Sign Language Generation

Md Rakibul Islam Midul and Taimoor khan, National Institute of Technology Silchar, India.

ABSTRACT

Text-to-Sign language generation has been introduced as a crucial field of research that helps to enhance accessibility communication for hearing-impaired people. In this work, we propose TSLTNet: Text-to-Sign Landmark Transformer Network which represents an encoder-decoder network based on Transformer architecture designed for predicting temporally ordered sign language landmarks based on input gloss text. The suggested architecture treats text-to-SL translation as a sequence-to-sequence task, where the Transformer encoder obtains semantic representations of gloss sequences, and the autoregressive decoder predicts SL landmark sequences related to facial, hand, and body gestures. Positional embeddings and attention mechanism allow preserving spatial and temporal consistency of SL landmark sequences generation. Both 2D and 3D SL landmark sequences are predicted by our framework and visualized in terms of motions. The suggested approach is trained by landmark reconstruction supervised learning and uses AdamW as an optimization technique. Experiments prove stable convergence, smooth motion generation, and high quality of landmark predictions. Our quantitative results yield ELDA of 81.4% and MSE of 0.92, providing excellent performance in terms of spatial alignment and motion reconstruction capabilities.

Keywords

Text-to-Sign Translation, Transformer Network, Landmark Generation, Sequence-to-Sequence Learning, Autoregressive Decoding, 2D and 3D Landmark Prediction, Human Pose Generation.

Human Agency in Intelligent Classrooms: Reimagining Pedagogy, Teacher Professionalism and Educational Transformation in the Age of Artificial Intelligence

Vicente C. Sinining, VCS Research, Rwanda.

ABSTRACT

Artificial intelligence is becoming embedded in education through adaptive learning systems, generative tools, automated assessment, learning analytics and institutional decision-support technologies. These developments offer possibilities for personalised, responsive and data-informed learning, but also raise concerns about teacher agency, learner dignity, privacy, bias, equity and pedagogical reductionism. Using a conceptual and integrative literature review, this paper examines how AI is reshaping teaching, assessment and educational governance. It develops a human-centred framework linking AI affordances, teacher professional judgement, learner-centred pedagogy, ethical governance, equity and institutional readiness. The analysis argues that AI strengthens education when it augments teachers' interpretive, relational and pedagogical work rather than replacing it. Without transparent governance, sustained professional development and inclusive design, AI may automate weak pedagogies and intensify inequalities. The paper reframes AI-enabled educational transformation as a question of professional, ethical and institutional capacity.

Keywords

Artificial Intelligence in Education, Teacher Agency, Human-centred Pedagogy, Learning Analytics, AI Ethics.

An Integrated AI, VR, and AR Based Wedding Planning System for Immersive and Personalized Event Design

Sandeepa K.A.B, Pathirana A.P.B.D, Kahatapitiya K.G.S.M, Minsandi N.D.C, Aruna Ishara Gamage and Didula Chamara, Sri Lanka Institute of Information Technology, Sri Lanka.

ABSTRACT

Wedding planning involves numerous decisions, including venue selection, decor, attire, catering, and over- all design. Traditional methods often rely on imagination, static visuals, and disconnected tools, making it difficult for users to effectively visualize and evaluate their choices. To address these limitations, this paper proposes an integrated wedding planning system that combines Virtual Reality (VR), Augmented Reality (AR), and Artificial Intelligence (AI) to create an immersive and interactive experience. The system includes four main components. First, a VR-based visualization module allows users to explore and customize wedding environments such as indoor, outdoor, and hybrid venues, while arranging elements like stage setups, poruwa structures, seating, catering, and cake designs. Second, an avatar-based dress customization module enables users to create personalized avatars and virtually try on wedding attire, with AR support for real-world visualization. Third, an AR decor module provides real-time decoration previews, along with budget tracking and vendor recommendations. Finally, an AI-driven studio module generates personalized wedding images based on user-uploaded photos and style preferences, incorporating feedback for refinement. By integrating these technologies into a single platform, the system enhances user engagement, improves spatial understanding, and supports more informed and efficient wedding planning decisions.

Keywords

Augmented Reality, Artificial Intelligence, Virtual Reality, Wedding Planning, 3D Visualization.

Embedding Reconstruction and Representation Alignment for Rul Prediction Under Partial Degradation Data

Ibrahim Mama Kouma, Southeast University, Chinna.

ABSTRACT

We propose a two-phase framework for RUL prediction under partial degradation data. Phase (a) pretrains on complete data via Adaptive Segmentation, Hierarchical Contrastive Learning, and Masked Embedding Reconstruction using a Transformer Encoder-Decoder $\phi(\cdot)$. Phase (b) adapts to partial data by reusing the pre-trained $\phi(\cdot)$ to actively reconstruct real missing embeddings, with Representation Alignment and Inter-Segment RUL Difference supervision. The five-term objective is $L = L_{rul} + \lambda_1 L_{cl} + \lambda_2 L_{recon} + \lambda_3 L_{align} + \lambda_4 L_{diff}$. On C-MAPSS FD001 with 20% MCAR missing timestamps, our method achieves RMSE 12.60, compared to 30.76 for a standard Transformer and 15.20 for BRITS.

Keywords

Remaining Useful Life Prediction, Partial Data, Embedding Reconstruction, Representation Alignment, Self-Supervised Learning.

The EGSP Settlement Matrix: A Symmetry-Constrained Hydrodynamic Model of Dimensional Transitions in Emergent Gravity

A.J. Ensor, Independent Researcher, USA.

ABSTRACT

The Ensor Gravity Source Postulation (EGSP) describes gravity as an emergent relaxation of a scalar field ϕ that controls the effective spacetime dimension $D_{eff}(\phi)$. We present the EGSP Settlement Matrix \mathcal{M} – a linearised dynamical system derived from the EGSP action under an overdamped (hydrodynamic) approximation, further constrained by a symmetry that globally conserves the total “dimensional charge” $Q = \int \phi d\mathcal{S}x$. The combination of overdamped relaxation and conservation selects a unique lowest-order effective dynamics: the Cahn-Hilliard (Model B) equation. Consequently, the evolution operator possesses an exact zero eigenvalue – not as an ad-hoc stability device, but as a strict consequence of the underlying symmetry. This removes the earlier criticism of “engineered stability”. The matrix governs bidirectional transitions: forward cascade (2D false vacuum \rightarrow 4D true vacuum) and reverse settlement (4D \rightarrow 2D in high-density regions). Using a biased double-well potential, eigenvalues are computed as functions of matter density ρ ; they exhibit positive (growth), negative (decay) and an exact zero eigenvalue. One-dimensional latticen simulations with conservative noise of amplitude $\eta \approx 0.0024$ produce a propagating front and a power spectrum of D_{eff} with excess large-scale (low- k) power directly tied to η . The value $\eta \approx 0.0024$ is chosen such that the resulting excess low- k power and variance in D_{eff} produce fractional corrections of order $\sim 0.2\%$ – consistent in magnitude with observed anomalies in void ellipticity distributions, suppressed power in CMB low multipoles, and potential deviations in large-scale

structure statistics. The framework predicts observable $\mathcal{O}(\eta)$ corrections in void ellipticity, black-hole ringdown frequencies, and low-multipole CMB anomalies, offering a falsifiable alternative to standard inflation. All numerical codes are provided as supplementary material.

The EGSP Domino Effect: Numerical Demonstration of a Conserved Dimensional Phase Transition with Stochastic Scaling

A.J. Ensor, Independent Researcher, USA

ABSTRACT

The Ensor Gravity Source Postulation (EGSP) predicts a rapid, coherent dimensional phase transition – the “domino effect” – driven by the relaxation of a scalar field ϕ under overdamped conserved dynamics (Cahn-Hilliard) with a universal stochastic amplitude $\eta = 0.0024$. We present direct numerical simulations of this process in one and three dimensions. The conservative noise is discretised in divergence form, exactly preserving the total dimensional charge $\int \phi d^d \mathcal{V}x$ to machine precision. A propagating front converts a metastable 2-dimensional phase to a 4-dimensional phase. The final effective dimension field $D_{\text{eff}} = 2 + 2\phi$ exhibits a statistically significant excess of power at low wavenumbers. Ensemble averaging over 20 independent 1D realisations gives a low- k excess factor $R_{\text{low}} = 3.72 \pm 0.41$ for the central η . This excess scales linearly with η : $R_{\text{low}} = 1 + (1130 \pm 30) \eta$ with $\chi^2/\text{dof} = 0.32$ ($p=0.81$). 3D simulations (48^3 grid) confirm the same qualitative features and yield $R_{\text{low}} \approx 3.45 \pm 0.25$ for $\eta = 0.0024$. Using the 3D data, we derive an explicit energy injection history $\delta\rho_{\text{eff}}(z)$ and compute the sound horizon shift $\Delta r_0/r_0 = -4.7\% \pm 0.6\%$, consistent with the value needed to resolve the Hubble tension. The low- k power excess also predicts a ~ 10 – 20% enhancement of the matter power spectrum at large scales. All raw data, simulation code, and reproducibility instructions are provided in the appendices.

Vortex-Inspired Baryon Masses from the EGSP Dimensional Vortex Formula

A. J. Ensor, USA

ABSTRACT

We derive a mass formula for baryons directly inspired by the Ensor Gravity Source Postulation (EGSP). In this framework, quarks emerge as topological vortices of a complex scalar field, with winding numbers n_i that encode flavour. The formula couples the squared total winding to the effective spacetime dimension $D_{\text{eff}}(\phi)$ and includes a logarithmic screening term modulated by the globally conserved dimensional charge $Q = \int \phi d^d \mathcal{V}x$. Using only the proton mass to fix the overall scale Λ and the recently measured Ξ_{cc} mass to fix the dimensionless coupling κ , the formula reproduces the Ω_{ccc} and Λ_{ccc} baryons and predicts the triply charmed Ω_{ccc} at 4.81 ± 0.03 GeV, in excellent agreement with lattice QCD. The framework requires no additional fitted parameters and is falsifiable by future measurements of doubly and triply heavy baryons. A semiclassical operator embedding shows that the mass relation is the leading saddle-point approximation of a quantizable topological sector.

Memory Poisoning as a Jailbreak Vector: Exploiting Persistent Agent Memory to Bias Tool-Grounded Decisions

Fatulla Bashirov The George Washington University, Washington, DC, USA

ABSTRACT

Persistent memory enables LLM-based agents to maintain context across sessions, improving personalization and continuity. However, this persistence introduces a novel attack surface: an adversary with normal user access can inject manipulative content into an agent’s memory through conversational interaction, causing the agent to make biased decisions in future sessions—even when real-time tool evidence contradicts the poisoned memories. We present a systematic study of this threat using a controlled experimental setup: a customer care agent equipped with MCP (Model Context Protocol) tools that provide ground-truth data. Through a four-phase progressive poisoning strategy executed via the agent’s chat interface, we demonstrate that (1) memory systems store adversarial directives without content validation, (2) in-context poisoning fully compromises agent decisions (composite score 4.8/5.0), and (3) cross-session memory poisoning partially succeeds but can be detected by the LLM’s own reasoning when directives are aggressively framed. Our findings reveal that current memory implementations lack architectural defenses against context poisoning, relying solely on the LLM’s incidental reasoning capabilities as an unintentional last line of defense. We propose concrete mitigations including memory content filtering, memory-tool conflict detection, and semantic integrity validation.

KEYWORDS

LLM security, agent memory, prompt injection, context poisoning, MITRE ATLAS, MCP tools

REAL-TIME SCHEDULING BASED ON SERVICE DEPENDENCY GRAPHS FOR RESOURCE-AWARE SERVICE ORCHESTRATION

Dulanga Wimalagunasekara, Yethmi Gamage, Teran Sarathchandra, Prashan Samarathunge, Vishan Jayasinghearachchi and Dharshana Kasthurirathna, Sri Lanka Institute of Information Technology, Sri Lanka

ABSTRACT

Microservice systems generate communication patterns that shift quickly and rarely look the same from one minute to the next. The default Kubernetes scheduler does not model these patterns. It places pods on the basis of CPU and memory, which can leave tightly interacting services on different nodes and so adds network delay and raises the chance that one failure pulls down others. This paper presents a real-time engine that maps service dependencies into a weighted graph and a Kubernetes scheduler that consults the graph during placement. The engine ingests Istio telemetry through Prometheus, stores the live and historical graph in Neo4j, and computes structural measures such as PageRank and Betweenness centrality. The scheduler combines four factors into a single per-node score, namely communication locality with the most active peers, resource utilization around a target

band, replica spread, and a risk penalty that discourages concentration of critical services. The prototype runs on a Kubernetes cluster with the Istio service mesh and is evaluated on the Google Cloud Microservices Demo workload (Online Boutique). On this workload the prototype reports a lower P95 service latency than the default Kubernetes scheduler. Broader scalability, scheduling-path overhead, failure-injection analysis, and comparison against affinity rules and traffic-aware schedulers such as TraDE are left for future work.

KEYWORDS

Microservices, Graph Theory, Dependency Analysis, Kubernetes Scheduling, Distributed Systems, Network Science, Resources Allocation