Vita Digital Journal

Life through Technology & Innovation

Computer Science Department Monthly

Vol. 1 | Issue 5 | June 2026

VITA DIGITAL JOURNAL

Life through Technology & Innovation

Computer Science Department Monthly

Volume 1 | **Issue 5** | **June** 2026

Published by:
Department of Computer Science,
Marian College of Arts and Science,
Thiruvananthapuram.

Chief Editor: Livin M Miranda
Student Editors:
Santhipriyan M
Joel George Mathew

mcasmedia123@gmail.com
https://www.mcas.ac.in/

EDITORIAL

June reminds us of our shared responsibility towards the environment through the observance of World Environment Day. As concerns about climate change, resource depletion, and environmental sustainability continue to grow, technology has an increasingly important role to play in building a greener future.

Computer Science is no longer limited to software development and digital innovation. Concepts such as Green Computing, energy-efficient data centers, smart environmental monitoring systems, Artificial Intelligence for climate research, and IoT-based solutions for resource management are helping address some of the world's most pressing environmental challenges. Technology, when developed and used responsibly, can become a powerful tool for sustainability.

As students and future professionals, it is important for us to recognize that innovation and environmental responsibility must go hand in hand. Every effort to optimize resources, reduce electronic waste, and design sustainable technological solutions contributes to a healthier planet.

Through this issue of Vita Digital Journal, we encourage our readers to explore how science and technology can support environmental conservation and inspire positive change. Let us strive to create solutions that not only advance society but also preserve the natural world for future generations

- Chief Editor
Mr. Livin M Miranda
Head, Department of Computer Science

Index

The Rise of Multimodal AI Systems: Integrating Perception, Language, and Reasoning

Fardeen Mohammed F R
S5 BSc Computer Science

INTRODUCTION

For most of the history of artificial intelligence research, systems were designed to handle one type of data at a time (text, images, audio, or video) each processed through its own specialized pipeline. That era is rapidly drawing to a close. Multimodal AI systems, which can simultaneously perceive and reason across multiple input types, are transforming how machines understand the world. In this sense, they are beginning to mirror the integrated sensory experience of human cognition.

The significance of this shift cannot be overstated. Human communication and understanding are inherently multimodal: we read facial expressions while listening to speech, interpret gestures alongside language, and make sense of documents that combine text and images. AI systems that can do the same open up qualitatively new possibilities in healthcare, education, science, and everyday interaction (Radford et al., 2021). This article examines what multimodal AI systems are, how they work, where they are being applied, and what challenges remain.

FOUNDATIONS AND ARCHITECTURE

A multimodal AI system is one that ingests and processes information from more than one sensory modality — most commonly text, images, and audio, though video, structured data, and even sensor inputs are increasingly included. Early efforts at combining modalities were relatively shallow, simply concatenating features from separate models.

The modern approach is more elegant: a shared representational space in which different modalities are projected into comparable embeddings, allowing the model to reason across them fluidly.

The transformer architecture, originally developed for natural language processing, has proven surprisingly adaptable to this challenge.

Vision transformers (ViTs) encode images as sequences of patch embeddings that can be processed alongside token embeddings from text. Contrastive learning methods, exemplified by OpenAI's CLIP model, train image and text encoders jointly so that a photograph of a dog and the phrase "a dog" end up close together in embedding space.

Modality	Input Format	Common Encoders	Example Use
Text	Tokens / strings	Transformer, BERT, GPT	Q&A, summarisation
Images	Pixel grids / patches	ViT, CNN, CLIP	Medical imaging, captioning
Audio	Waveforms / spectrograms	Wav2Vec, Whisper	Speech recognition, music
Video	Frame sequences	TimeSformer, VideoMAE	Action recognition, surveillance
Sensor / IoT	Time-series signals	LSTM, Temporal CNN	Robotics, wearables

Table 1: Summary of input modalities, encoding approaches, and example applications in multimodal AI systems

APPLICATIONS ACROSS DOMAINS

The practical applications of multimodal AI are already wide-ranging. In medicine, models trained on radiology images paired with clinical notes can assist radiologists in detecting anomalies and drafting preliminary reports, potentially reducing diagnostic errors and easing bottlenecks in overloaded healthcare systems (Acosta et al., 2022). Pathology, ophthalmology, and dermatology represent similarly promising areas, where visual classification combined with clinical context yields better outcomes than either modality alone.

In education, multimodal systems can analyse a student's handwritten work, spoken explanations, and facial expressions simultaneously to adapt instruction in real time. Accessibility applications are equally compelling: systems that describe images for visually impaired users, or that generate captions and sign-language animations from speech, can substantially expand digital access. Beyond these domains, multimodal AI is reshaping creative fields, robotics, and scientific research.

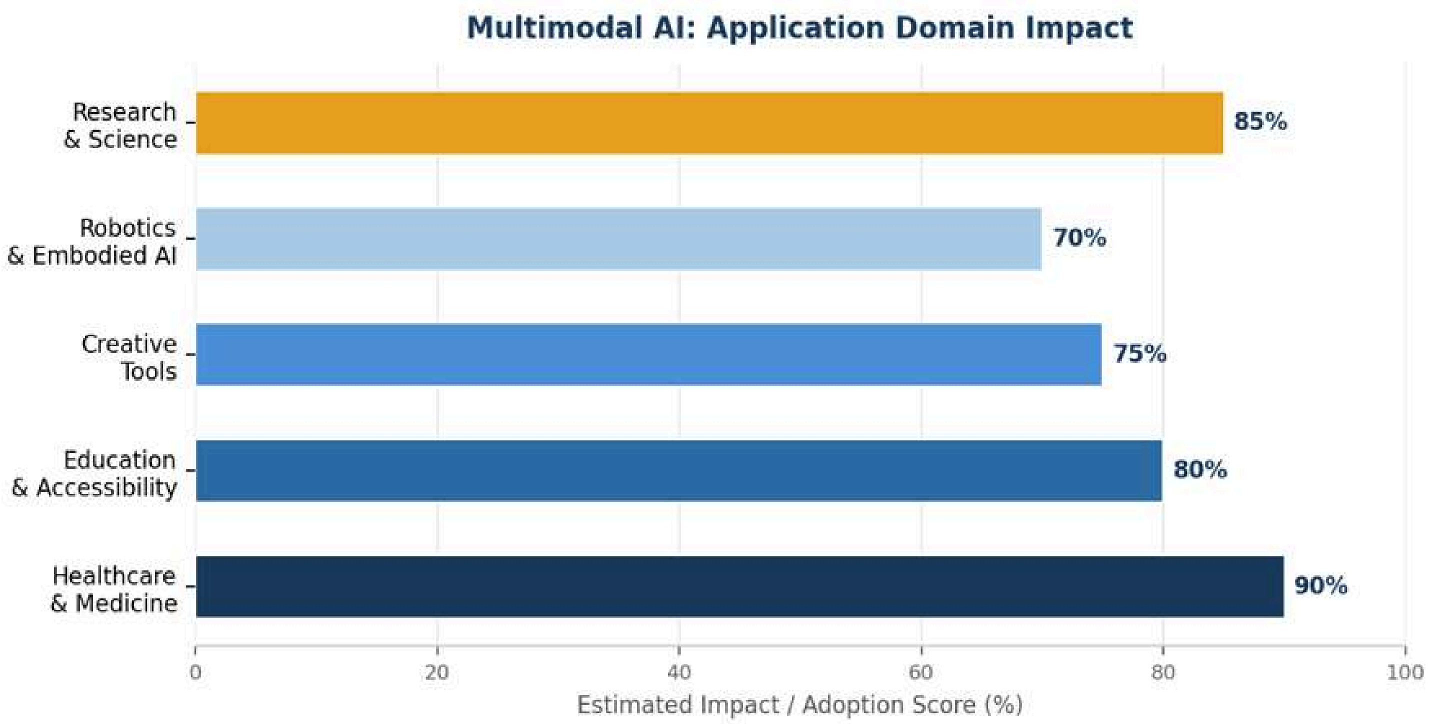

CHALLENGES AND LIMITATIONS

Despite impressive progress, multimodal AI faces substantial unresolved challenges. One of the most persistent is hallucination — the tendency of large models to generate plausible-sounding but factually incorrect content, a problem that becomes more complex when visual and textual information must be aligned. A model might correctly identify an object in an image while generating inaccurate claims about it, or misinterpret a chart while producing fluent prose that appears authoritative (OpenAI, 2023). Alignment across modalities also remains difficult. Text and image encoders trained separately may develop representations that do not map cleanly onto one another, leading to failures in cross-modal reasoning tasks. Computational cost is a practical barrier , multimodal models are large, training them requires substantial energy and hardware, and deploying them at scale is expensive. Privacy and bias are additional concerns, particularly in high-stakes domains.

Challenge	Description	Mitigation Strategies
Hallucination	Generating plausible but factually wrong content	Retrieval-augmented generation, grounding
Cross-modal Alignment	Encoders may not map modalities to	Joint contrastive training (e.g., CLIP)
Computational Cost	Large models require heavy hardware &	Efficient architectures, quantisation
Bias & Privacy	Sensitive attributes can enable discrimination	Fairness audits, differential privacy
Robustness	Poor generalisation to out-of-distribution	Data augmentation, adversarial training

Table 2: Key challenges in multimodal AI and representative mitigation strategies

RECENT DEVELOPMENTS

The pace of advancement in multimodal AI has been remarkable. GPT-4V, released by OpenAI in 2023, demonstrated the ability to interpret complex diagrams, read handwritten notes, and answer questions about photographs with a level of sophistication that surprised even researchers in the field (OpenAI, 2023). Google's Gemini family extended this further, with models trained natively on text, images, audio, and video rather than adapting separate systems post hoc (Google DeepMind, 2023). Meta's ImageBind project explored an even more ambitious goal: a single shared embedding space uniting six distinct modalities, enabling cross-modal retrieval without explicit paired training data (Girdhar et al., 2023).

On the research side, instruction-tuning and reinforcement learning from human feedback (RLHF) have been adapted for multimodal settings, improving the ability of models to follow complex multimodal instructions. Techniques for grounding model outputs in retrieved evidence are also being extended to multimodal contexts, helping to reduce hallucination and improve factual accuracy.

CONCLUSION

Multimodal AI systems represent one of the most significant shifts in the trajectory of artificial intelligence research. By integrating perception across text, images, audio, and other modalities, these systems are beginning to approach the richness of human sensory experience not by replicating it, but by developing their own forms of cross-modal understanding. The applications are already substantial, from clinical decision support to accessible education to creative tools, and the pace of advancement suggests that the most transformative uses are still ahead.

At the same time, the field must grapple honestly with the challenges it faces: hallucination, alignment failures, computational cost, bias, and privacy risks. Progress on benchmarks does not always translate to reliable real-world performance, and the social consequences of deploying powerful multimodal systems in high-stakes contexts demand careful governance and interdisciplinary scrutiny. The coming decade will likely determine whether these systems fulfil their considerable promise and whether their benefits are distributed broadly or concentrated among a privileged few.

REFERENCES

- Acosta, J. N., Falcone, G. J., Rajpurkar, P., & Topol, E. J. (2022). Multimodal biomedical AI. Nature Medicine, 28(9), 1773–1784. https://doi.org/10.1038/s41591-022-01981-2
- Baltrusaitis, T., Ahuja, (2019). Multimodal machine learning: A survey and taxonomy. IEEE Transactions on Pattern Analysis and Machine Intelligence, 41(2), 423–443.
- Girdhar. (2023). ImageBind: One embedding space to bind them all. CVPR 2023. https://doi.org/10.1109/CVPR52729.2023.01457
- Google DeepMind. (2023). Gemini: A family of highly capable multimodal models. Technical Report.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444.
- OpenAI. (2023). GPT-4 technical report. arXiv:2303.08774.
- Radford. (2021). Learning transferable visual models from natural language supervision. ICML 2021, 8748–8763.
- Reed, S. Zolna. (2022). A generalist agent. Transactions on Machine Learning Research.

Data Analytics: Transforming Data into Valuable Insights

Vismaya P
S5 BSc Computer Science

ABSTRACT

In the modern digital era, data has become one of the most valuable resources for organizations and individuals. Every day, enormous amounts of data are generated through online transactions, social media platforms, mobile applications, healthcare systems, and business operations. Data analytics is the process of examining, organizing, and interpreting data to extract meaningful information that supports better decision-making. It enables organizations to identify trends, predict future outcomes, improve efficiency, and gain a competitive advantage.

INTRODUCTION

Data analytics has emerged as one of the most important technologies in today's information-driven world. Businesses, governments, educational institutions, and healthcare organizations rely on data to understand customer behaviour, improve services, reduce costs, and make informed decisions. Data analytics converts raw data into meaningful insights through statistical analysis, machine learning algorithms, and visualization techniques.

TYPES OF DATA ANALYTICS

Data analytics is generally divided into four major categories. Descriptive analytics focuses onunderstanding what has happened in the past by summarizing historical data through reports anddashboards. Diagnostic analytics goes a step further by identifying the reasons behind past events anddiscovering patterns or relationships within the data. Predictive analytics uses historical information andmachine learning models to forecast future outcomes, such as customer purchasing behavior or diseaseoutbreaks. Prescriptive analytics recommends the best possible actions based on analytical results, helpingorganizations optimize decisions and improve overall performance.

INDUSTRY APPLICATIONS

Data analytics has transformed numerous industries. In healthcare, it helps diagnose diseases earlier, monitor patients remotely, and predict disease outbreaks. In banking and finance, analytics detects fraudulent transactions, assesses credit risk, and develops investment strategies. Retail companies analyse purchasing behaviour to recommend products and optimise inventory.

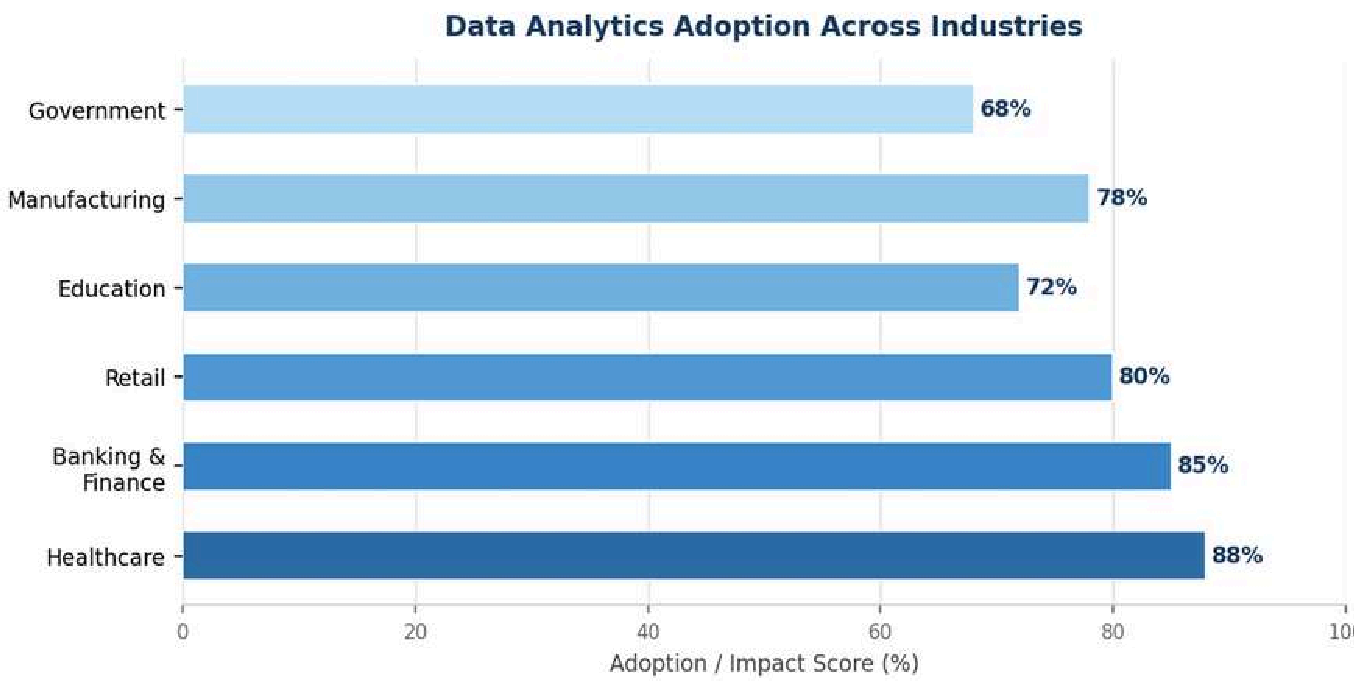

Figure 1: Data Analytics Adoption and Impact Score Across Key Industries

THE DATA ANALYTICS PROCESS

The process of data analytics begins with collecting data from various sources such as databases, websites, social media platforms, sensors, and mobile applications. Once collected, data undergoes a cleaning process to remove duplicates, missing values, and errors. Clean data is then organised and processed into a suitable format for analysis. Statistical methods, machine learning techniques, and analytical models are applied to identify patterns and relationships. Finally, results are presented using charts, graphs, dashboards, and reports.

KEY BENEFITS

The adoption of data analytics provides several important benefits. It enables organizations to make betterdecisions based on factual evidence rather than assumptions. Businesses can improve operationalefficiency, reduce costs, increase productivity, and enhance customer satisfaction. Analytics also supportsrisk management by identifying potential threats before they become serious problems.

Furthermore, organizations gain a competitive advantage by understanding market trends, customer preferences, andfuture opportunities more effectively than their competitors.

CHALLENGES

Despite these advantages, data analytics presents several challenges. Data quality is a primary concern, as inaccurate or incomplete data leads to incorrect conclusions. Protecting sensitive personal and organisational information makes data privacy and cybersecurity essential aspects of any analytics project. The high cost of advanced software, cloud infrastructure, and skilled professionals can limit adoption, especially for small organisations. Integrating data from multiple sources and ensuring compliance with legal and ethical standards remain ongoing challenges.

FUTURE PROSPECTS

The future of data analytics is highly promising due to continuous advancements in Artificial Intelligence, Machine Learning, cloud computing, and the Internet of Things. Real-time analytics, automated decision-making, and intelligent data visualisation tools are becoming increasingly common across industries. As technology evolves, data analytics will continue to play a vital role in innovation, economic growth, and digital transformation.

CONCLUSION

Data analytics has become an essential component of modern society by transforming large volumes of raw data into valuable knowledge. It empowers organisations to make informed decisions, improve efficiency, enhance customer experiences, and solve complex business problems. Although challenges such as data privacy, security, and implementation costs remain, continuous technological advancements are making analytics more accessible and effective.

REFERENCES

- Han, J., Kamber, M., & Pei, J. (2012). Data Mining: Concepts and Techniques (3rd ed.). Morgan Kaufmann.
- Provost, F., & Fawcett, T. (2013). Data Science for Business. O'Reilly Media.
- Marr, B. (2021). Data Strategy: How to Profit from a World of Big Data, Analytics and Artificial Intelligence. Kogan Page.
- IBM. What is Data Analytics? Retrieved from ibm.com 5. Microsoft. Introduction to Data Analytics. Retrieved from microsoft.com

Agentic AI Gone Wrong: AgentJacking—A New Attack on Autonomous AI Agents

Gayathri Devi A

S5 BSc Computer Science

ABSTRACT

In June 2026, security experts revealed AgentJacking, an attack class that seizes control of autonomous AI programming agents through malicious commands sent to error monitoring telemetry. AgentJacking does not use any malware or credentials stealing; instead, it abuses a fundamental architectural design flaw common among all Model Context Protocol (MCP) stack applications—data received by the tool cannot be interpreted as executable code. This paper provides a summary of the details regarding the AgentJacking attack, its generality, and the proposed mitigation strategy for enterprises using agentic AI.

INTRODUCTION

Two years of AI code assistants operated as copilots that suggested code changes while the human reviewed and implemented those changes. But by 2026, this paradigm has shifted to agents that can read through error logs, generate code, execute shell commands, implement changes, and do it all with the same privileges as the human developer . Such an expansion of the scope of activities results in a completely new attack vector which is no longer the application being protected, but the decision-making process of the AI agent.

AgentJacking, a novel class of attacks on the AI agents that have appeared in June 2026 due to the research done by the scientists from Tenet Security, represents the first ever such type of attacks . Such attacks exploit AI coding assistants such as Claude Code, Cursor, OpenAI Codex through forgery of bug reports without delivering any malware, stealing passwords, or exploiting a vulnerability of the targeted application.

BACKGROUND: TRUST IN THE MCP ECOSYSTEM

The Model Context Protocol (MCP) represents the leading industry protocol used by AI agents for connecting to external services such as issue tracking systems, error monitoring services, databases, and software repositories. The very nature of the MCP protocol assumes the data provided by the tool as its response is merely passive context for the model to consume but not active commands for the model to execute . However, AgentJacking shows this assumption does not hold true the very moment when some third-party gets permission to write to the connected service — which, in case of many popular developer tools, is a design choice.

Unfortunately, it is far from being the only weakness that exists within the protocol. AgentJacking represents one of many examples of a specific class of vulnerabilities described by the Cloud Security Alliance during 2025-2026 as part of MCP ecosystem research and documented by such security vendors as Unit 42 of Palo Alto Networks and Elastic Security Labs.

ANATOMY OF THE ATTACK

AgentJacking is an attack against Sentry, a popular error tracking service. Sentry receives errors reported by applications with a Data Source Name (DSN) – a piece of credential deliberately included in client-side code for unauthenticated reporting of crashes . As the DSN can be written by any attacker who finds it, they are able to post a fake error report that looks to the agent like any other error report. There are four steps to the attack, as described in Fig. 1 below

Discovery — the attacker discovers a target’s public Sentry DSN using browser-side JavaScript or code search engines such as GitHub search or Censys.

Injection – the attacker POSTs a crafted error report to Sentry's ingestion endpoint, including a fake “resolution” message with an executable shell command inside.

Retrieval – a developer asks their AI code agent to look into the problem. They query Sentry with MCP and retrieve the poisoned error report just like regular diagnostics information.

Execution – since the agent believes everything reported by tools, it executes the command with the privileges of the developer on their machine.

MITIGATION STRATEGIES

Guidance for practitioners that has come into alignment since the discovery revealed in June 2026 centers around defense-in-depth rather than a single solution. The key idea is to re-establish a trust boundary between retrieved data and agent instructions that MCP does not inherently impose. Table I provides the key guidelines.

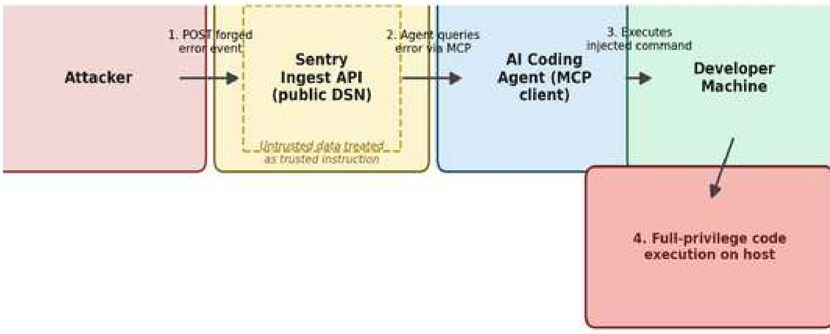

Fig. 1. AgentJacking attack chain

Mitigation	Effect
Treat all tool output as untrusted input	Prevents retrieved data from being parsed as instructions
Human-in-the-loop approval for write actions	Blocks autonomous execution of high-risk commands
Sandboxed / least-privilege agent execution	Limits blast radius if an agent is hijacked
DSN / credential rotation and scoping	Reduces attacker ability to inject events undetected
Provenance tagging across MCP tool chains	Lets agents distinguish trusted vs. external content

RECOMMENDED MITIGATIONS FOR AGENTJACKING-CLASS ATTACKS

DISCUSSION AND CONCLUSION

Rather than thinking about AgentJacking as a one-off attack against any particular product, it is more helpful to consider it as one of the first tangible demonstrations of a core issue in agentic AI: an agent acting on behalf of its developer inherits the developer's privilege level without inheriting his ability to recognize a trap. Insofar as agents are given more access to production systems, the pattern—malicious data quietly entering the instruction stream—is destined to repeat itself in some form unless trust boundaries are formally defined in protocols.

From the perspective of researchers and practitioners, there are several concrete directions worth exploring: models of trust propagation across agent chains, runtime provenance tracking for responses to MCPs, and benchmarks that test agents' resilience against indirect prompt injection rather than just correctness. As of now, this is an active area of discussion in both industry and academia.

REFERENCES

- Tenet Security, “Agentjacking 2026: How Fake Sentry MCP Errors Hijack AI Coding Agents,” Decryption Digest, June 2026. [Online].
- Penligent AI Threat Research, “AI Agents Hacking in 2026: Defending the New Execution Boundary,” February 2026. [Online].
- Cloud Security Alliance, AI Safety Initiative, “Research Note: Agentjacking — MCP Sentry Injection,” June 12, 2026. [Online].
- Palo Alto Networks Unit 42, “Covert Tool Invocation and Conversation Hijacking via MCP Sampling,” cited in Cloud Security Alliance Research Note, 2026.

Artificial Intelligence for Climate Research

Ms. Mubeena J
Asst. Prof.,
Dept of Computer Science

OVERVIEW

One of the most significant scientific challenges facing humanity in the 21st century is that of climate change. Despite their physically accurate nature, classical Earth System Models remain computationally expensive and uncertain due to the issue of parameterizations. The purpose of this paper is to investigate the revolution in climate science brought about by Artificial Intelligence.

CORE CONTRIBUTIONS

- An organizational taxonomy of the use of AI methods in climate studies
- An outlined AI climate research pipeline model (5 phases)
- Performance benchmarking of six different AI architecture types
- Design of a carbon cycle monitoring system using AI methods
- Highlighted open issues and research directions

THE FIVE-STAGE PIPELINE

This paper presents a coherent process that integrates observation with policymaking in the following manner:

- *Step 1* – Data Gathering: Satellites, IOT devices, reanalysis of datasets
- *Step 2* – Preprocessing and Fusion: Standardization, multivariate fusion
- *Step 3* – AI/ML Modeling: Physics-based ML approach, deep learning algorithms
- *Step 4* – Prediction and Analysis: Prediction, anomalies, scenario generation

- *Step 5* – Decision Making: Policy dashboards, adaptation and mitigation.

The feedback loop between step 5 and step 3 ensures model updating as new data comes in.

AI TECHNIQUES REVIEWED

- *Convolutional Neural Networks*: Leveraging spatial locality of grid-based climate variables; precipitation mapping, cyclone detection, estimation, and SST prediction using CNNs.
- *LSTMs & Recurrent Networks*: Multidecadal time sequences modeled by these models; streamflow is employed for forecasting, drought prediction, and sea-level rise.
- *Transformers*: Global modeling through attention mechanism; powering best performing weather AI (Pangu-Weather, GraphCast, FourCastNet) which beat NWP at 10-day lead times.
- *Physics-Informed Neural Networks*: Including governing laws such as Navier-Stokes and energy balance equations as constraint for learning; necessary in cases where observations are unavailable.
- *Graph Neural Networks*: Function on relational data structures – teleconnection network, rivers basins. Core of GraphCast's multi-scale icosahedral mesh.
- *Generative Models*: Statistical downscaling, generating ensembles and quantifying uncertainties in regional climate projections.

KEY APPLICATION DOMAINS

Prediction of Extreme Events: The CNN method manages to detect 98% of atmospheric river events; flood forecasting based on LSTM in the state of Kerala has an accuracy of 89% after 72 hours compared to the 73% accuracy of NWP methods.

Carbon Cycle Observation: Machine Learning techniques can reduce uncertainties in net ecosystem exchange from eddy covariance data, satellite vegetation indices, and atmospheric carbon dioxide columns by 30–40%.

Ocean and Cryosphere: Physical-ML methods cut down the root mean squared error in sea-level projections compared to the CMIP6 ensemble projections by 22%;

GNNs discover new heat transport mechanisms in the South Atlantic.
Climate Model Emulation: Transformer-based surrogate models can reduce decades-long GCM runs into quick ensemble generation processes.

PERFORMANCE BENCHMARKS

Critical conclusion: hybrid Physics-ML systems have the best skill scores (0.95), demonstrating that integrating domain expertise surpasses purely data-driven or physics-based methods.

AI Technique	Best Application	Skill Score
Transformer / GNN	10-day NWP, Teleconnections	0.94
Physics-ML Hybrid	Carbon flux, Earth systems	0.95
LSTM	Streamflow & drought index	0.91
CNN	Precipitation mapping	0.87
GAN	Downscaling & scenarios	0.89

SIX AI ARCHITECTURES TRANSFORMING CLIMATE RESEARCH

There isn’t one way for AI to tackle the issue of climate change. Various architectures perform better in specific areas — recognizing spatial patterns, modeling time dependencies, capturing physics, or building network relationships. These are the six that matter.

Convolutional Neural Networks:
Proficient in recognizing spatial patterns in gridded climate variables. Applications include the identification of atmospheric rivers, cloud classification, and sea surface temperature mapping with skill scores of 87%.

LSTM & Recurrent Networks:
Model multi-decade temporal dependencies. Apply for stream flow prediction, drought indices forecast, and sea level projection at lead times of 12 months.

Transformers:
Global connections modeled by attention mechanisms. Pangu-Weather, GraphCast, and FourCastNet exceed the numerical weather predictions at 10 day horizon with skill scores of 0.94.

***Physics Informed Neural Networks*:**
Navier-Stokes and energy balance equations embedded into loss functions as constraints. The top-performing category achieving skill scores of 0.95 in estimating carbon fluxes.

***Graph Neural Networks* :**
Operate on graph structures - teleconnection graphs, river networks. GraphCast by Google DeepMind uses the multi-resolution icosahedral mesh GNN for the best forecasting results.

***Generative Models*:**
Used for statistical downscaling and probabilistic scenarios creation. GANs generate the km resolution precipitation maps from coarse-grained GCM simulations, diffusion models give uncertainty quantification.

CHALLENGES

- *Interpretability*: Output from black box insufficient for policy science; SHAP and attention visualization require climate-specific approaches
- *Data equity*: Deserts of observations in Africa, South Asia, Arctic; transfer learning and semi-supervised learning hold promise
- *Compute sustainability*: Large model training involves carbon cost; NAS and distillation required
- *Uncertainty quantification*: No unified UQ approach like CMIP ensembles exists
- *Multiscale modeling*: Seamless architecture from global to local scale still under research

CONCLUSION

Artificial intelligence, in particular AI that combines both physics and machine learning, is revolutionizing climate science in terms of monitoring, modeling, and decision-making processes. The pipeline described in the paper provides a framework through which AI can be incorporated into all aspects of scientific research. Data explainability, data fairness, and computing sustainability will be key issues in the next 10 years.

REFERENCES

- Reichstein et al., Nature 566, 195–204 (2019) · IPCC AR6 (2021)
- S. Rasp, M. S. Pritchard, and P. Gentine, “Deep learning to represent sub-grid processes in climate models,” Proceedings of the National Academy of Sciences, vol. 115, no. 39, pp. 9684-9689, 2018.
- T. Schneider et al., "Earth System Modeling 2.0: A Blueprint for Models That Learn From Observations and Targeted High-Resolution Simulations," Geophys. Res. Lett., vol. 44, no. 24, pp. 12,396–12,417, 2017.

Sovereign Cloud Architectures for Artificial Intelligence: Principles, Frameworks, and Strategic Imperatives

Ms. Rini Amado
Asst. Prof.,
Dept of Computer Science

ABSTRACT

The quick adoption of large foundational Artificial Intelligence (AI) models has driven the economic change but has brought strategic weaknesses concerning the legal jurisdiction of the data, geopolitical issues, and digital colonialism. The hyper-scalers public clouds, which are usually by default based on the centralized infrastructures crossing legal borders, can bring compliance risks with respect to the GDPR and national security regulations. In this research paper, we define the concepts and architecture of Sovereign Cloud Architectures designed especially to support AI workloads. Technical criteria including data sovereignty of infrastructure, physical autonomy of operations, and software independence are defined. Based on the analysis of current orchestration practices of hybrid architectures, federated computing nodes, and hardware-enforced trusted execution environments , this research paper defines the reliable reference architecture for sustainable digital self-determination of the local area without reducing the state-of-the-art computational capabilities.

INTRODUCTION

The prevalence of Artificial Intelligence (AI) and Large Language Models (LLMs) has turned computing from a widely utilized technological tool into an element of national sovereignty . As modern deep learning algorithms use the processing of massive data sets, which have many layers, the geographic areas in which this data flows, resides, and is being trained have become a matter of contention. Processing proprietary data by public and private organizations in an infrastructure governed by foreign jurisdictions is tantamount to undermining the infrastructure's boundary limits and violating the rules of regional compliance. Sovereign cloud infrastructures address these security flaws. Sovereign clouds consist of isolated execution environments, operational architecture, and governance models, which ensure that the processing of the data is done entirely within a particular jurisdictional limit . Importantly, as AI becomes increasingly dependent on real-time pipelines, the creation of sovereign infrastructures is no longer a question of storage boundaries alone but one of redesigning the entire lifecycle process.

CONCEPTUALIZING SOVEREIGN AI

The concept of sovereign AI refers to the capability of an individual or a state to develop, train and deploy advanced artificial intelligence through their own infrastructure, workforce, and data sources . The process of obtaining sovereign AI necessitates the resolution of three crucial elements of the core:

- **Data Sovereignty**: Making sure all input data, generated datasets, and metadata of the model are entirely confined within jurisdictional boundaries and not subject to extraterritorial claims or inspections.
- **Operational Sovereignty**: Making sure all maintenance, systems engineering, and system administration are done by the local people and not subjected to shutdowns or updates by the foreign provider
- **Technological Sovereignty**: Avoiding the use of any black-box proprietary API and relying only on open-source foundation models and orchestration engine.

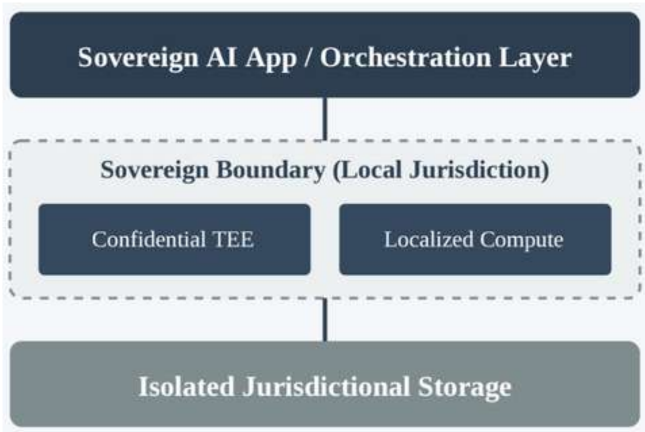

ARCHITECTURAL IMPLEMENTATIONS

For the implementation of sovereign clouds for AI applications, it is necessary to move away from old unencrypted computational methods to zero trust infrastructure approaches.

Confidential computing using trusted execution environments

To secure the process during the training and inference stages, it is necessary to use hardware-based trusted execution environments (TEEs), such as AMD SEV-SNP or Intel TDX . This microarchitecture-based security encrypts the data in system memory lines, preventing even the hypervisor administrators and cloud providers from accessing model parameters and datasets .

Federated learning across jurisdictional nodes

In cases where it is not possible to create a shared data repository because of the regulatory requirements, the topology of federated learning is implemented . In such a scenario, the global parameters of the model are securely transferred to the regional nodes for further computations. Only the gradients are returned to the server by means of cryptographic aggregation methods .

EMPIRICAL EVALUATION AND TRADE-OFFS

Moving to sovereign infrastructure entails observable tradeoffs in terms of processing latency, overheads, and training speed [3]. In order to measure such differences, various benchmarks were done on three major platforms: standard decentralized public cloud, shared hybrid multi-region cloud, and pure sovereign node using hardware memory encryption mechanisms.

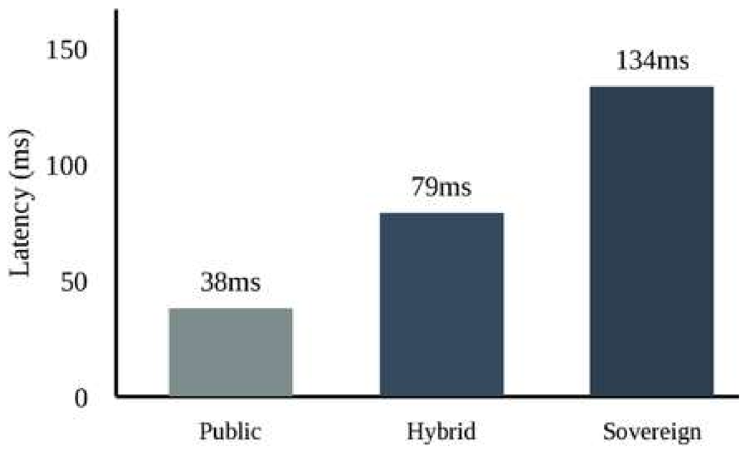

However, according to empirical measurements, although public instances offer the least latency in terms of operations (38ms), they fail to provide compliance checks and sufficient isolation capabilities that are mandated by the rigid national security frameworks. Pure sovereign instances suffer from higher latency levels, which go up to 134ms. The reason for such high latencies is attributed to encryption of operations and routing of traffic through local nodes.

CONCLUSION

Architectures of sovereign clouds are evolving from the specialized type to become the basic building block of a secure environment for corporate and state AI systems. Leveraging memory protection through hardware means, regionalized data architecture, and open-source foundation models, it is possible to protect against digital threats . The performance costs that were found in the experiment are justified by complete digital self-sovereignty.

REFERENCES

- L. Heim et al., "Governing through the cloud: The intermediary role of compute providers in AI regulation," arXiv preprint arXiv:2403.08501, 2024.
- A. Shoker, "Digital sovereignty strategies for every nation," Appl.Cybersecurity & Internet Governance, vol. 1, no. 1
- W. D. Cloud and J. Doe, "Principles of regional digital sovereignty in distributed processing platforms," IEEE Trans. Cloud Comput.,vol. 11, no. 2, pp. 142–155, 2024.
- R. Amado, "An analysis of zero-trust security postures within distributed infrastructure deployment architectures," J. Sovereign Software Eng., vol. 4, pp. 210–225, May 2026.
- E. Albaroudi and M. A. Hatamleh, "Empowering data sovereignty through artificial intelligence: A framework for sustainable smart energy systems"

Green Computing: A Sustainable Approach in Computer Science

Ms. Anila A J
Asst. Prof.,
Dept of Computer Science

ABSTRACT

rapid advancement of Information and Communication Technology (ICT) has significantly transformed modern society. However, the increased use of computing devices, data centers, and communication networks has resulted in high energy consumption and environmental concerns. Green Computing, also known as Green Information Technology (Green IT), focuses on designing, manufacturing, using, and disposing of computing resources in an environmentally responsible manner. This paper discusses the concept, components, technologies, benefits, challenges, and future trends of Green Computing. The study highlights how sustainable computing practices contribute to energy conservation, reduction of electronic waste, and environmental protection while maintaining computational efficiency.

INTRODUCTION

Information Technology has become an essential component of modern life, supporting education, healthcare, communication, business, and scientific research. The increasing demand for computing resources has led to higher electricity consumption and the generation of large quantities of electronic waste (e-waste). According to environmental studies, the ICT sector contributes significantly to global carbon emissions through energy-intensive computing infrastructures.

Green Computing emerged as a solution to address these environmental challenges. It promotes sustainable practices that reduce the ecological footprint of computing systems while ensuring optimal performance. The objective of Green Computing is to maximize energy efficiency, minimize waste generation, and encourage environmentally responsible technology usage.

CONCEPT OF GREEN COMPUTING

Green Computing refers to the efficient and eco-friendly use of computing resources throughout their lifecycle. It involves reducing power consumption, utilizing renewable energy sources, minimizing hazardous materials, and promoting recycling and reuse of electronic equipment.

The concept is based on four major principles:

- Green Design
- Green Manufacturing
- Green Usage
- Green Disposal

These principles collectively contribute to sustainable development in the field of Computer Science.

COMPONENTS OF GREEN COMPUTING

***Green Design*:**
Green Design concentrates on creating efficient hardware and software products. Green designers try to lower the energy consumption and environmental footprint at all stages of product lifecycle.

***Green Manufacturing*:**
This part concentrates on green manufacturing processes, decreasing the amount of hazardous substances and industrial waste.

***Energy Efficiency*:**
Energy-efficient products use less power but have no impact on their functionality. Low-energy processors, SSD drives and power management software play an important role in energy conservation.

***Green Usage*:**
Green Usage is about encouraging users to save electricity, turning computers into sleep mode, printing less and shutting them down when they are not used.

E-Waste Management:
Appropriate recycling and discarding of outdated computer equipment can protect the environment from contamination and recycle valuable resources.

Sustainable Disposal:
Sustainable disposal will make sure that the products can be recycled or reused to avoid contamination of hazardous substances in landfills.

TECHNOLOGIES SUPPORTING GREEN COMPUTING

Virtualization:
Virtualization allows multiple virtual machines to run on a single physical server, reducing hardware requirements and energy consumption.

Cloud Computing:
Cloud computing optimizes resource utilization through shared infrastructure, reducing operational costs and power usage.

Renewable Energy Integration:
Modern data centers increasingly use renewable energy sources such as solar and wind power to minimize carbon emissions.

Smart Power Management:
Operating systems and hardware have power saving features that ensure energy consumption is minimized.

AI for Efficient Resource Utilization:
Artificial intelligence can be used to optimize the use of energy within computing infrastructure.

BENEFITS OF GREEN COMPUTING

There are various advantages of Green Computing to the environment, economics, and society.

Environmental Benefits

- Carbon emission reduction.
- Natural resource conservation.
- Electronic waste reduction.
- Reduced environmental pollution.

Economic Benefits

- Cost reduction of electricity.
- Maintenance cost reduction.
- Long life of computing gadgets.

- Effective use of resources.

Social Benefits

- Awareness of the environment promotion.
- Good public health due to pollution reduction.
- Sustainable development support.

CHALLENGES IN IMPLEMENTING GREEN COMPUTING

Despite its advantages, Green Computing faces several challenges:

- High initial investment costs.
- Lack of awareness among users.
- Rapid increase in electronic waste generation.
- Complex recycling processes.
- Energy consumption versus performance.

These challenges can be overcome through collaboration between different sectors, including government, industry, education sector, and end-users.

FUTURE TRENDS

The future of Green Computing is closely associated with advancements in Artificial Intelligence, Edge Computing, Internet of Things (IoT), and Renewable Energy Technologies. Smart data centers powered by renewable energy and AI-based energy management systems are expected to become increasingly common.

Also, research related to carbon-neutral computing infrastructure and sustainable hardware production methods is going on.

CONCLUSION

Green Computing is one of the significant fields that have been gaining increasing popularity in the Computer Science discipline. With the constant expansion of technologies in computing, the sustainable methods are needed for minimizing environmental impact while still allowing for advancements. Energy-saving technology, cloud computing, virtualization, renewable energy usage, and efficient e-waste management are all ways through which Green Computing helps promote sustainability in our environment.

REFERENCES

- S. Murugesan, "Harnessing Green IT: Principles and Practices," IT Professional, vol. 10, no. 1, pp. 24–33, 2008.

- Kansal and F. Zhao, "Fine-Grained Energy Profiling for Power-Aware Application Design," ACM SIGMETRICS, pp. 26–31, 2008.

- G. Dhiman and V. Sharma, Green Computing and Environmental Sustainability, Springer, 2017.

- R. Buyya, C. Yeo, and S. Venugopal, "Market-Oriented Cloud Computing: Vision, Hype, and Reality," Future Generation Computer Systems, vol. 25, no. 6, pp. 599–616, 2009.

- U.S. Environmental Protection Agency, "Green Computing Initiatives and Practices," EPA Publications, 2024.

TECH NEWS

MICROSOFT LAUNCHES MAI-CODE-1-FLASH ON GITHUB COPILOT

Microsoft has introduced MAI-Code-1-Flash, a new artificial intelligence model designed specifically for software development tasks and integrated into GitHub Copilot. The model is part of Microsoft's MAI (Microsoft AI) family of models and aims to provide developers with faster and more efficient coding assistance.

MAI-Code-1-Flash is optimized for low-latency performance, allowing it to generate code suggestions, complete programming tasks, and assist with debugging in real time. By focusing on speed and responsiveness, the model helps developers maintain productivity during coding sessions without long waiting times for AI-generated responses.

The new model is being made available to organizations using GitHub Copilot Business and Enterprise plans. It can support a wide range of software development activities, including code generation, code completion, bug fixing, refactoring, and explanation of complex code segments. Its integration into GitHub Copilot enables developers to access these capabilities directly within their development environments.

This launch highlights Microsoft's ongoing investment in developing its own AI technologies. While GitHub Copilot has historically relied on multiple AI models, the introduction of MAI-Code-1-Flash demonstrates Microsoft's efforts to build specialized in-house models tailored for specific use cases such as software engineering. The move also reflects the growing competition among technology companies to create advanced AI systems that can improve workplace productivity and automate complex tasks.

The introduction of MAI-Code-1-Flash is expected to benefit software developers by reducing development time, improving coding efficiency, and enhancing the overall programming experience. As AI-assisted coding tools continue to evolve, models like MAI-Code-1-Flash are likely to play an increasingly important role in modern software development workflows.

Source : ***Microsoft AI***

NVIDIA RAISES ANTE IN AI PC MARKET WITH RTX SPARK

Nvidia has introduced RTX Spark, a next-generation AI computing platform aimed at accelerating the development and adoption of AI-powered personal computers. Announced during Computex 2026, RTX Spark combines Nvidia's latest AI hardware and software technologies to enable users to run sophisticated AI models directly on their PCs without relying entirely on cloud-based services.

RTX Spark is built around a powerful architecture that integrates Nvidia's advanced Blackwell GPU technology with high-performance processing capabilities. The platform is designed to support AI agents, large language models, content creation tools, software development applications, and other AI-driven workloads. By processing AI tasks locally, RTX Spark can offer faster response times, improved privacy, and reduced dependence on internet connectivity.

One of the major goals of RTX Spark is to transform Windows PCs into intelligent AI workstations capable of understanding user commands, automating routine tasks, generating content, and assisting with productivity. Nvidia envisions a future where AI agents can perform complex actions such as managing schedules, analyzing documents, generating code, creating multimedia content, and interacting with software applications on behalf of users.

The platform also benefits developers by providing tools and frameworks that simplify the creation and deployment of AI applications. Through Nvidia's AI ecosystem, developers can optimize models for local execution and take advantage of GPU acceleration to improve performance. This could lead to a new generation of AI-powered software that runs efficiently on consumer devices rather than requiring expensive cloud infrastructure.

RTX Spark arrives at a time when major technology companies are competing to establish leadership in the emerging AI PC market. As artificial intelligence becomes increasingly integrated into everyday computing, hardware manufacturers are focusing on delivering systems capable of handling AI workloads natively. Nvidia's announcement strengthens its position in this rapidly growing sector and expands its influence beyond traditional gaming and data-center markets.

Source : *TechNewsWorld*

INDIA'S INDIGENOUS SUPERCOMPUTER, PARAM RUDRA, OFFICIALLY ENTERED THE JUNE TOP 500 GLOBAL RANKINGS

India marked a major achievement in the field of high-performance computing with the inclusion of PARAM Rudra in the June 2026 TOP500 global supercomputer rankings. The accomplishment reflects the country's growing technological capabilities and its commitment to developing advanced computing infrastructure through indigenous innovation.

PARAM Rudra was developed under the National Supercomputing Mission (NSM), a flagship initiative jointly implemented by the Ministry of Electronics and Information Technology and the Department of Science and Technology. The mission aims to strengthen India's computational capabilities and support cutting-edge research across various scientific and industrial domains.

The supercomputer is built using Indian-designed HPC (High Performance Computing) server technology and a domestically developed software ecosystem. By relying on indigenous hardware and software components, PARAM Rudra reduces dependence on imported technologies while promoting self-reliance in a strategic sector. The project aligns with India's broader vision of technological sovereignty and the Atmanirbhar Bharat initiative.

High-performance computing systems such as PARAM Rudra are essential for solving complex computational problems that cannot be handled by conventional computers. These systems are used in a wide range of fields including weather forecasting, climate modeling, artificial intelligence, machine learning, drug discovery, aerospace engineering, astrophysics, disaster prediction, and national security research. The inclusion of PARAM Rudra in the prestigious TOP500 list demonstrates that India's indigenous technology can compete with some of the world's most powerful computing systems.

The TOP500 ranking is published twice a year and evaluates supercomputers based on their performance in the LINPACK benchmark, a standard test used to measure computational power. Entry into this list is regarded as a significant international recognition because it places a system among the world's leading supercomputers. PARAM Rudra's presence in the rankings highlights the progress made by Indian researchers and engineers in designing and deploying world-class computing infrastructure.

Beyond scientific research, PARAM Rudra is expected to contribute to the growth of India's digital economy and innovation ecosystem. Universities, research institutions, startups, and government agencies will be able to leverage its computational resources to accelerate research and development activities. The system will also support the training of advanced AI models and the analysis of large-scale datasets, which are becoming increasingly important in modern scientific and industrial applications.

Experts believe that achievements such as PARAM Rudra's entry into the TOP500 rankings will encourage further investment in indigenous semiconductor technologies, advanced computing architectures, and next-generation AI infrastructure. As countries worldwide compete to develop more powerful computing systems, India's success demonstrates its emergence as a significant player in the global supercomputing landscape.

Source : ***PIB Press Release***

INSIDE HYUNDAI'S METAPLANT AND THE FUTURE OF MANUFACTURING

The Hyundai Motor Group Metaplant America (HMGMA) in Ellabell, Georgia, is one of the most advanced automotive manufacturing facilities in the world. Established as Hyundai's newest U.S. production hub, the factory represents the future of smart manufacturing by combining artificial intelligence (AI), robotics, automation, and digital technologies. The facility began full-scale production in 2024 and currently manufactures electric vehicles such as the Hyundai Ioniq 5 and Ioniq 9, with plans to expand production to Kia and Genesis models. transport parts throughout the facility. Hyundai also uses AI-driven quality control systems, drones, digital twin technology, and robotic inspection tools to improve efficiency and reduce manufacturing errors.

What makes the Metaplant unique is its hyper-flexible manufacturing system. Unlike traditional factories that are designed for a single type of vehicle, Hyundai's facility can produce electric vehicles, hybrid vehicles, and potentially other powertrain types on the same production line. This flexibility allows Hyundai to adapt quickly to changing market demands and consumer preferences.

The factory relies heavily on advanced robotics and AI-powered systems. Hundreds of robotic arms assist in welding, assembly, inspection, and material handling, while autonomous mobile robots

A key feature of the Metaplant is its focus on collaboration between humans and machines. Employees receive specialized training through Hyundai's on-site training center, where they learn to work alongside advanced technologies. Rather than replacing workers entirely, automation is used to handle repetitive and physically demanding tasks, allowing employees to focus on supervision, problem-solving, and system management.

Experts believe facilities like the Metaplant represent the next stage of industrial evolution. Future manufacturing plants are expected to use more intelligent AI systems capable of coordinating supply chains, managing production schedules, and enabling build-to-order manufacturing. Advances in robotics and large-scale 3D printing may further transform how vehicles are designed and produced over the coming decades.

The Hyundai Metaplant demonstrates how AI, automation, and digital technologies are reshaping modern manufacturing. By creating a flexible, intelligent, and highly automated production environment, Hyundai is setting a new benchmark for the automotive industry and providing a glimpse into the future of global manufacturing.

Source : ***TechNewsWorld***

EDITORIAL BOARD