



Vision

To evolve as a centre of academic excellence and advanced research in Computer Science and Engineering discipline

Mission

To inculcate students with profound understanding of fundamentals related to discipline, attitudes, skills, and their application in solving real world problems, with an inclination towards societal issues and research

The Department of Information Technology was started UG Programme in "Information Technology" in the year 2000 with an initial intake of 60 students and the intake is enhanced to 120 in the year 2007. In 2011, the department started a PG Program in "Computer Science and Technology (CST)" with an intake of 36 students. The department started a new The Department was established in 2000 with an initial intake of 60 students and has consistently expanded its academic offerings and capacity in response to the growing demand for quality technical education. The undergraduate intake was enhanced to 120 in 2007, further increased to 180 in 2022, and subsequently revised to 120 in 2025 in accordance with institutional and regulatory requirements. The Department has established a strong legacy of academic excellence, having earned NBA accreditation six times for its undergraduate programme, including accreditation cycles during 2008–09, 2014–17, and 2017–20, with the current Outcome-Based Education (OBE) accreditation granted for a period of six years (July 2022 to June 2028), reflecting its sustained commitment to quality education and continuous improvement. PG program in "Data Science" with an intake of 18 students in the year 2019. NBA Thrice accredited the B.Tech. Programme during 2008-09, 2014-17 and 2017-2020 and extended Accreditation till 30-6-2021 (OBE Tier I).

Dr.M.Suneetha
HOD : IT

EDITORIAL BOARD

Faculty Editors

Mrs.G.Geetha

Student Editors

P.BHAVESH SRI PAVAN(238W1A12G9)

K KUSUMA (238W1A1298)

VOLUME : 1 | ISSUE: 1 | DECEMBER 2025

Table of Contents

Agentic AI: From Generative Models to Autonomous Intelligence	Dr.T.Anuradha	Professor
Zero-Trust Security with AI-Based Threat Intelligence	Dr.Y.Sangeetha	Associate Professor
Digital Twin–Driven Smart Cities and Industrial Systems	Ch.NandaKrishna	Assistant Professor
Vision Transformers for Next-Generation Computer Vision	Y.Kalyan Chakravarthy	Assistant Professor
Foundation Models for Multispectral and Hyperspectral Earth Observation	Dr.V.Satya Sahithi	Assistant Professor
AI-Native 6G Networks and Intelligent Wireless Systems	Venkata Lakshmi	Assistant Professor
Quantum Machine Learning: Algorithms, Architectures and Emerging Applications	T.Sujith Kumar	Assistant Professor
Multimodal Generative AI for Medical Imaging and Precision Healthcare	P.Bhavesh	III/IV B.Tech

Agentic AI: From Generative Models to Autonomous Intelligence

Dr.T.Anuradha
Professor

ARTIFICIAL INTELLIGENCE

- Agentic AI extends generative AI from response generation toward goal-directed planning, tool use and autonomous task execution.
- An agent maintains state, interprets context, selects actions and iteratively evaluates outcomes.
- Multi-agent systems distribute complex tasks among specialized agents such as planner, researcher, coder and verifier.
- Memory, retrieval, tool orchestration and human approval are increasingly important for reliable deployment.

TECHNICAL ARCHITECTURE

EXPERT TECHNICAL OVERVIEW

A contemporary agentic architecture typically combines a foundation model with planning, memory, retrieval and external tools. The controller decomposes a user objective into subtasks and assigns them to specialized agents. Tool calls may access databases, APIs, code execution environments or enterprise systems. A verifier evaluates intermediate results before the controller proceeds to the next step. State and long-term memory preserve relevant context across interactions. Human-in-the-loop checkpoints are retained for sensitive decisions, while observability and evaluation frameworks track tool failures, hallucinations, latency and task completion quality.



Figure 1: Agent orchestration graph with specialized agents, tools, memory and verification

Iterative planning • tool use • verification • memory

Zero-Trust Security with AI-Based Threat Intelligence

Dr.Y.Sangeetha
Associate Professor

INDUSTRIAL IOT / SMART SYSTEMS

- Digital twins maintain a dynamic digital representation of physical assets, processes or infrastructure.
- Sensor observations are synchronized with computational models to support monitoring, diagnosis and prediction.
- Simulation enables what-if analysis before operational decisions are applied to the physical system.
- AI-driven twins are moving from visualization toward predictive and prescriptive decision support.

TECHNICAL ARCHITECTURE

EXPERT TECHNICAL OVERVIEW

An advanced digital-twin environment integrates physical assets, sensing infrastructure, communication middleware, data platforms and simulation models. Real-time telemetry updates the virtual representation, while historical information supports trend analysis and machine-learning models. The twin can estimate equipment health, simulate operating conditions and identify deviations between expected and observed behavior. In a closed-loop implementation, validated decisions are transferred back to the physical process through control systems. This architecture is relevant to manufacturing, energy networks, buildings, transportation and large-scale infrastructure.




Figure 1: Digital-twin closed-loop synchronization between physical and virtual environments

Continuous synchronization and closed-loop decision support

Digital Twin–Driven Smart Cities and Industrial Systems

Ch.NandaKrishna

Assistant Professor

EDGE COMPUTING / EMBEDDED AI

- Edge AI executes machine-learning inference close to sensors, reducing dependence on centralized cloud processing.
- TinyML targets highly constrained microcontrollers with limited memory, compute and power budgets.
- Quantization, pruning and knowledge distillation enable compact models suitable for embedded deployment.
- Local inference is valuable where latency, privacy, connectivity and energy efficiency are critical.

EXPERT TECHNICAL OVERVIEW

The Edge AI pipeline begins with physical sensing followed by signal conditioning and local feature extraction. A compact neural model performs inference on a microcontroller, edge gateway or embedded accelerator. Only events, features or selected summaries need to be transmitted to a higher-level platform. Model compression techniques such as integer quantization and structured pruning reduce memory footprint and computation. Continuous or event-driven sensing can therefore support real-time applications such as predictive maintenance, environmental monitoring, agriculture, wearables and industrial safety.

TECHNICAL ARCHITECTURE



Figure 3. TinyML edge stack from sensing hardware to local inference and cloud learning

Local inference → event-driven communication → fleet-level learning

Vision Transformers for Next-Generation Computer Vision

Y.Kalyan Chakravarthy

Assistant Professor

CYBERSECURITY / CRYPTOGRAPHY

- Post-quantum cryptography develops classical cryptographic mechanisms designed to resist attacks from future quantum computers.
- Migration is important because encrypted information captured today may be vulnerable to future decryption.
- Lattice-based, hash-based and code-based constructions are major families of quantum-resistant cryptography.
- Hybrid deployment can combine existing public-key mechanisms with post-quantum algorithms during transition.

TECHNICAL ARCHITECTURE

EXPERT TECHNICAL OVERVIEW

A quantum-safe security architecture begins with cryptographic asset discovery and identification of systems dependent on vulnerable public-key algorithms. During migration, organizations can deploy hybrid key-establishment and signature mechanisms while maintaining compatibility with existing infrastructure. Cryptographic agility allows algorithms to be replaced without redesigning the complete application stack. Key management, certificate infrastructure, software libraries, network protocols and long-lived data archives all need consideration. The transition therefore involves both algorithm selection and enterprise-wide cryptographic inventory and governance.



Figure 4: Post-quantum cryptography migration roadmap with hybrid and crypto-agile stages

Crypto-agility enables staged migration without redesigning the full enterprise stack

Foundation Models for Multispectral and Hyperspectral Earth Observation

Dr.V.Satya Sahithi
Assistant Professor

GENERATIVE AI / FOUNDATION MODELS

- Multimodal foundation models learn relationships across text, images, audio, video and structured information.
- A shared representation enables reasoning across multiple input modalities within a unified interaction.
- Retrieval augmentation connects foundation models to current, domain-specific and enterprise knowledge.
- Evaluation increasingly considers factuality, grounding, robustness, safety and modality-specific performance.

TECHNICAL ARCHITECTURE

EXPERT TECHNICAL OVERVIEW

A multimodal generative architecture uses modality-specific encoders or tokenizers to convert different data types into representations that can interact inside a common model space. A language or multimodal transformer performs contextual reasoning and generates text, structured outputs or other supported modalities. Retrieval components can inject authoritative external information at inference time, while guardrails constrain unsafe or unreliable behavior. In enterprise environments, orchestration layers manage access control, document retrieval, model routing, logging and evaluation. The architecture supports intelligent assistants, document understanding, visual question answering and content generation.




Figure 5: Multimodal foundation model with converging modality streams and grounded generation

Multiple modalities converge into a shared reasoning space before grounded generation


AI-Native 6G Networks and Intelligent Wireless Systems

Venkata Lakshmi
Assistant Professor

REMOTE SENSING / GEOSPATIAL AI

- Geospatial foundation models learn transferable representations from large volumes of satellite, aerial and environmental observations.
- Multitemporal and multispectral information provides richer representations than single-date image analysis.
- Self-supervised pretraining can reduce dependence on large manually labelled remote-sensing datasets.
- Fine-tuned models support land-cover mapping, change detection, agriculture, disaster response and environmental monitoring.

TECHNICAL ARCHITECTURE



EXPERT TECHNICAL OVERVIEW

A modern geospatial AI pipeline combines optical, radar, elevation and temporal observations into a common representation space. Large-scale self-supervised pretraining learns spectral, spatial and temporal patterns before the model is adapted to a specific downstream task. Patch or token embeddings can be processed by transformer-based architectures, while geospatial metadata provides additional context. Fine-tuning or lightweight adaptation then supports classification, segmentation, object detection and change analysis. This paradigm enables reuse of learned Earth-observation representations across regions and applications.

Quantum Machine Learning: Algorithms, Architectures and Emerging Applications

T.Sujith Kumar

Assistant Professor

CYBERSECURITY / AI

- AI-driven security platforms analyze large volumes of endpoint, identity, network and application telemetry.
- Behavioral analytics can identify deviations that are difficult to detect using static signatures alone.
- Security orchestration combines detection, correlation, investigation and automated response.
- Explainable alerts and analyst feedback are important for controlling false positives and improving operational trust.

TECHNICAL ARCHITECTURE

EXPERT TECHNICAL OVERVIEW

An enterprise AI security architecture collects telemetry from endpoints, identity providers, network infrastructure, applications and cloud environments. A feature and event-processing layer normalizes this data before behavioral models assign risk scores to users, devices and activities. A correlation engine connects related events into incidents and enriches them with threat intelligence. Security orchestration platforms can then automate selected actions such as session revocation, device isolation or credential challenges. Analyst feedback closes the learning loop and helps improve detection quality while maintaining human oversight for high-impact actions.



Figure 7: AI-enabled SOC workflow from telemetry collection to automated response and feedback

Telemetry → correlation → behavioral detection → automated response → analyst feedback

Multimodal Generative AI for Medical Imaging and Precision Healthcare

P.Bhavesb
III/IV B.Tech

HEALTHCARE AI

Image–Text Fusion, Retrieval Grounding & Clinical Reasoning

- Multimodal clinical AI integrates medical images, radiology or pathology reports, laboratory observations and structured clinical variables.
- Vision encoders transform high-dimensional medical images into representations that can be aligned with language and clinical concepts.
- Retrieval-augmented generation can ground model responses in validated guidelines, institutional protocols and relevant prior evidence.
- Structured clinical outputs can explicitly separate findings, supporting evidence, uncertainty and suggested follow-up actions.
- Human validation is essential because clinical datasets exhibit distribution shift, demographic imbalance, annotation variability and incomplete observations.

CLINICAL AI REASONING

A clinically responsible multimodal architecture separates perception, evidence retrieval, reasoning and decision presentation. Medical images are processed through domain-specific vision encoders, while clinical notes and structured observations are converted into compatible representations. A multimodal fusion model reasons over these heterogeneous evidence sources. Retrieval services provide validated clinical guidelines and institutionally approved knowledge, reducing dependence on model memory alone. The response layer is constrained to produce structured findings, supporting evidence and uncertainty indicators. An audit and provenance layer records which evidence influenced each recommendation. The final interface is designed as clinician decision support rather than autonomous diagnosis, with human review retained for ambiguous, high-risk or clinically consequential cases.

- Research emphasis includes multimodal alignment, calibration, privacy-preserving learning, provenance, explainability and prospective clinical validation.




Figure 8: Multimodal clinical AI architecture with evidence retrieval, reasoning and clinician validation