INFOLINE

VOLUME XVI ISSUE II OCTOBER 2025

DEPARTMENT OF COMPUTER TECHNOLOGY AND INFORMATION TECHNOLOGY

KONGU ARTS AND SCIENCE COLLEGE
(Autonomous)
Affiliated to Bharathiar University,
Coimbatore.
Accredited with A+ Grade - 3.49 CGPA by NAAC
NANJANAPURAM, ERODE - 638 107

INFOLINE

EDITORIAL BOARD

EXECUTIVE COMMITTEE

CONTENTS

TOP ANTIVIRUS SOFTWARE IN 2025

The antivirus landscape is constantly evolving to combat new and sophisticated threats like ransomware and zero-day exploits. "Next-Generation Antivirus" (NGAV) is a key trend, moving beyond traditional signature-based detection to use advanced techniques like machine learning, behavioural analysis and AI for more proactive and comprehensive protection. Some of the best antivirus software options available today are:

TotalAV: Often considered a top choice, with strong malware detection, ease of use and extra features such as a VPN, password manager and optimization tools. According to SafetyDetectives, it has a 100% detection rate. TotalAV provides a 30-day money-back guarantee.

Norton 360: A reputable suite known for robust malware protection, including phishing protection, an unlimited VPN, password manager, cloud storage and parental controls. SafetyDetectives indicates that Norton 360's scanner achieves 100% protection through heuristic analysis and machine learning.

Bitdefender Total Security: Features an advanced cloud-based scanning engine for effective malware detection with minimal system impact. It includes various protections like web protection, vulnerability scanning and ransomware protection. SafetyDetectives notes that Bitdefender's anti-phishing protection effectively blocked most phishing sites in testing.

McAfee Total Protection: Provides a solid set of internet security features, including malware and anti-phishing protection, a password manager, VPN and identity theft monitoring. It is particularly suitable for families due to its parental controls and unlimited device protection.

Surfshark Antivirus: A good option for users seeking a VPN bundle, offering strong malware protection and customizable scans with low system impact.

ESET NOD32 Antivirus: A lightweight and efficient antivirus that provides advanced protection against complex threats like fileless malware and zero-day exploits.

Avast One: Features an intuitive interface and is a suitable choice for macOS users, including a firewall for that platform.

Avira: Offers a powerful antivirus engine with high detection rates and helpful system optimization tools. According to Cybernews, its free version offers strong malware protection and useful features.

AVG AntiVirus Free: Provides strong web and email protection, along with system performance optimization tools.

Considerations When Choosing the Antivirus Software

Security features: Look for features like real-time protection, malware detection and removal, behavioural analysis, web protection, automatic updates and parental controls.

System impact: Choose software that runs efficiently without slowing down your device.

Ease of use: Look for an intuitive interface with clear instructions.

Compatibility: Ensure the software works with one's operating system and devices.

Updates and support: Regular updates are vital to combat new threats.

Reputation and reviews: Research the software and vendor's reputation.

Pricing and licensing: Consider one's budget and the number of devices one's need to protect.

Free vs. paid antivirus

While free antivirus software can provide basic protection against known threats, paid versions offer a wider array of features and stronger protection against emerging threats like zero-day attacks. Some free antivirus options, like Bitdefender and Avast Free, include real-time protection, which is essential for ongoing security.

Deepa Jothi M

I B.Sc. (Information Technology)

MODERN PROGRAMMING LANGUAGE

Kotlin is a modern, concise and expressive programming language that has gained widespread adoption in recent years. Developed by JetBrains, the creators of popular IDEs like IntelliJ IDEA, Kotlin was designed to address common challenges faced by developers while working with existing languages such as Java. Officially released in 2016, Kotlin is now a preferred language for Android app development and offers robust features that make it suitable for a wide range of applications.

The Evolution of Kotlin

JetBrains introduced Kotlin in 2011 as a statically-typed programming language to improve productivity and reduce boilerplate code. The language gained significant traction after Google announced official support for Kotlin as a first-class language for Android development in 2017. Kotlin is open-source, with its source code available on GitHub, encouraging community contributions and improvements.

Key Features of Kotlin

Conciseness: Kotlin reduces boilerplate code compared to Java, enabling developers to write less code and focus more on logic and functionality.

Null Safety: Kotlin introduces null safety as a first-class feature, addressing the infamous

NullPointerException (NPE) issue common in Java.

Interoperability: One of Kotlin's standout features is its seamless interoperability with Java, allowing developers to use existing Java libraries and frameworks without modification.

Coroutines for Asynchronous Programming: Kotlin supports coroutines, making it easier to write asynchronous and non-blocking code. This feature is particularly beneficial for applications that require high performance and responsiveness.

Extension Functions: Developers can add new functionality to existing classes without modifying their source code, enhancing flexibility and readability.

Applications of Kotlin

Android Development: Kotlin is widely used for developing Android applications due to its concise syntax, null safety and robust tool support. Popular apps like Pinterest, Evernote and Netflix are built with Kotlin.

Server-Side Development: With frameworks like Ktor and Spring Boot, Kotlin is an excellent choice for server-side programming, offering scalability and high performance.

Cross-Platform Development: Kotlin Multiplatform enables developers to share code across multiple platforms, including Android, iOS, and web applications, reducing duplication and increasing efficiency.

Data Science: Kotlin's simplicity and integration with Java libraries make it a suitable language for data science and machine learning applications.

Advantages and Limitations of Kotlin

Advantages

- Improved developer productivity.
- Enhanced code safety and readability.
- Strong tooling support from JetBrains.

Limitations

- Learning curve for developers new to Kotlin.
- Slower compilation times compared to Java in some cases.
- Smaller community compared to more established languages.

With its growing popularity and expanding ecosystem, Kotlin is poised to play a significant role in the future of programming. Its versatility and alignment with modern development practices make it a language worth investing in for developers and organizations alike.

M.Harini

III B.Sc. (Computer Technology)

CUTTING-EDGE MACHINE LEARNING METHODS

Machine Learning (ML) has rapidly evolved over recent years, driven by advances in algorithms, computational power, and the availability of vast datasets. Cutting-edge machine learning methods represent the forefront of this evolution, offering powerful new tools that significantly enhance the ability to model complex data, learn from limited examples and generate human-like outputs. These methods span a wide range of techniques, including deep learning architectures, reinforcement learning innovations, unsupervised and self-supervised learning, and hybrid approaches that blend multiple paradigms. The continuous development of these methods is revolutionizing fields such as natural language processing, computer vision, robotics, healthcare, and many more, making AI more accurate, adaptable and accessible.

One of the most prominent and impactful cutting-edge techniques is deep learning, especially the use of advanced neural network architectures. Deep learning models, inspired by the structure and functioning of the human brain, utilize multiple layers of artificial neurons to automatically extract hierarchical features from raw data. Convolutional Neural Networks (CNNs) remain the backbone of many image and video recognition systems, enabling computers to interpret visual data with near-human accuracy. Meanwhile, Recurrent Neural Networks (RNNs) and their more advanced variants, such as Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRUs), have been instrumental in modelling sequential data like speech, text, and time series. More recently, the introduction of transformer architectures has revolutionized natural language processing (NLP). Transformers leverage self-attention mechanisms to efficiently model long-range dependencies in data, dramatically improving the quality of language understanding and generation. Models like BERT, GPT series and T5 have set new benchmarks in numerous NLP tasks, including translation, summarization, question answering and conversational AI.

Beyond supervised learning, self-supervised learning and unsupervised learning methods have emerged as pivotal in addressing the challenge of labeled data scarcity. Self-supervised learning enables models to generate their own supervisory signals from the data itself, effectively learning useful representations without extensive human annotation. For example, contrastive learning techniques train models to recognize different augmented views of the same data point, thereby capturing essential features in a scalable way. Such approaches have gained traction especially in domains like computer vision and speech recognition, where manual labelling can be prohibitively expensive. Similarly, unsupervised learning methods focus

on uncovering underlying data structures, such as clustering and dimensionality reduction, to extract meaningful patterns and insights. Techniques like variational autoencoders (VAEs) and generative adversarial networks (GANs) are at the forefront of generative modelling, producing realistic images, audio, and text by learning data distributions. GANs, in particular, consist of competing neural networks a generator and a discriminator that progressively improve each other, resulting in impressive outputs used in creative arts, data augmentation and anomaly detection.

Another area of intense research is Reinforcement Learning (RL), where agents learn optimal actions through trial and error by interacting with an environment and receiving feedback in the form of rewards or penalties. Cutting-edge RL methods have made remarkable strides, especially in domains requiring sequential decision-making and strategy. Advances such as deep reinforcement learning combine RL with deep neural networks to handle high-dimensional sensory inputs and complex policy spaces. Breakthroughs like Deep Q-Networks (DQN), policy gradient methods and actor-critic algorithms have enabled machines to achieve superhuman performance in games like Go, chess, and real-time video games. Moreover, recent innovations in multi-agent reinforcement learning and hierarchical RL expand the capabilities of agents to collaborate, compete and learn abstract skills transferable across tasks. These developments hold great promise for robotics, autonomous vehicles and personalized recommendation systems, where agents must adapt dynamically to changing environments and objectives.

Hybrid and ensemble methods also represent a cutting-edge trend in machine learning. By combining multiple models or techniques, these methods aim to leverage the strengths of each component to improve overall performance and robustness. For example, integrating symbolic reasoning with neural networks sometimes referred to as neuro-symbolic AI can enhance interpretability and generalization, addressing some of the limitations of purely data-driven approaches. Ensemble learning techniques like boosting, bagging, and stacking aggregate predictions from diverse learners to reduce variance and bias, often leading to superior predictive accuracy. Meta-learning, or “learning to learn, is another exciting approach where models acquire the ability to quickly adapt to new tasks with minimal data, mimicking human learning efficiency. This technique is particularly useful in few-shot or zero-shot learning scenarios, enabling AI systems to generalize better in real-world applications where labelled data is sparse or unavailable.

Scalability and efficiency have become critical focus areas in the development of cutting-edge machine learning methods, as real-world applications demand models that are

not only accurate but also computationally feasible. Techniques such as model pruning, quantization, and knowledge distillation help reduce model size and inference latency, making deployment on edge devices and mobile platforms more practical. Moreover, federated learning has emerged as an innovative paradigm that allows multiple devices or organizations to collaboratively train models without sharing raw data, preserving privacy and security. This is particularly valuable in sensitive domains like healthcare and finance, where data confidentiality is paramount. Advances in hardware, including specialized AI accelerators like GPUs, TPUs, and neuromorphic chips, complement these algorithmic improvements, enabling faster training and real-time inference at scale.

Ethical considerations and explainability are increasingly integrated into cutting-edge machine learning research. As ML systems are deployed in critical areas such as criminal justice, hiring, and healthcare, the demand for transparency, fairness, and accountability grows. New techniques for interpretable machine learning strive to make complex models more understandable to humans, facilitating trust and enabling stakeholders to diagnose and mitigate biases. Explainability methods range from post-hoc analysis tools, like SHAP and LIME, to inherently interpretable model architectures designed to provide insights into decision-making processes. Additionally, fairness-aware learning algorithms aim to detect and correct discriminatory patterns in data and predictions, promoting equitable outcomes. Addressing these challenges is essential for responsible AI adoption and to ensure that the benefits of cutting-edge ML methods are shared widely and ethically.

Cutting-edge machine learning methods encompass a diverse and rapidly advancing set of techniques that push the boundaries of what artificial intelligence can achieve. From deep learning and reinforcement learning breakthroughs to innovative unsupervised approaches and hybrid models, these methods empower machines to understand, generate and act upon data with increasing sophistication. Coupled with ongoing efforts to improve scalability, efficiency, transparency and fairness, cutting-edge ML is transforming industries and society at large. As research progresses and these methods become more accessible, the future of machine learning promises even greater impact, enabling smarter, more adaptable and ethically aligned AI systems.

M.S.K Manassha

III B.Sc. (Information Technology)

THE RISE OF ARTIFICIAL INTELLIGENCE IN DATA ANALYTICS

Machine Learning (ML) has rapidly evolved over recent years, driven by advances in algorithms, computational power, and the availability of vast datasets. Cutting-edge machine learning methods represent the forefront of this evolution, offering powerful new tools that significantly enhance the ability to model complex data, learn from limited examples and generate human-like outputs. These methods span a wide range of techniques, including deep learning architectures, reinforcement learning innovations, unsupervised and self-supervised learning, and hybrid approaches that blend multiple paradigms. The continuous development of these methods is revolutionizing fields such as natural language processing, computer vision, robotics, healthcare and many more, making AI more accurate, adaptable and accessible.

One of the most prominent and impactful cutting-edge techniques is deep learning, especially the use of advanced neural network architectures. Deep learning models, inspired by the structure and functioning of the human brain, utilize multiple layers of artificial neurons to automatically extract hierarchical features from raw data. Convolutional Neural Networks (CNNs) remain the backbone of many image and video recognition systems, enabling computers to interpret visual data with near-human accuracy. Meanwhile, Recurrent Neural Networks (RNNs) and their more advanced variants, such as Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRUs), have been instrumental in modelling sequential data like speech, text, and time series. More recently, the introduction of transformer architectures has revolutionized natural language processing (NLP). Transformers leverage self-attention mechanisms to efficiently model long-range dependencies in data, dramatically improving the quality of language understanding and generation. Models like BERT, GPT series, and T5 have set new benchmarks in numerous NLP tasks, including translation, summarization, question answering, and conversational AI.

Beyond supervised learning, self-supervised learning and unsupervised learning methods have emerged as pivotal in addressing the challenge of labelled data scarcity. Self-supervised learning enables models to generate their own supervisory signals from the data itself, effectively learning useful representations without extensive human annotation. For example, contrastive learning techniques train models to recognize different augmented views of the same data point, thereby capturing essential features in a scalable way. Such approaches have gained traction especially in domains like computer vision and speech recognition, where manual labelling can be prohibitively expensive. Similarly, unsupervised learning methods focus

on uncovering underlying data structures such as clustering and dimensionality reduction, to extract meaningful patterns and insights. Techniques like Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs) are at the forefront of generative modelling, producing realistic images, audio, and text by learning data distributions. GANs, in particular, consist of competing neural networks a generator and a discriminator that progressively improve each other, resulting in impressive outputs used in creative arts, data augmentation and anomaly detection.

Another area of intense research is Reinforcement Learning (RL), where agents learn optimal actions through trial and error by interacting with an environment and receiving feedback in the form of rewards or penalties. Cutting-edge RL methods have made remarkable strides, especially in domains requiring sequential decision-making and strategy. Advances such as deep reinforcement learning combine RL with deep neural networks to handle high-dimensional sensory inputs and complex policy spaces. Breakthroughs like Deep Q-Networks (DQN), policy gradient methods and actor-critic algorithms have enabled machines to achieve superhuman performance in games like Go, chess and real-time video games. Moreover, recent innovations in multi-agent reinforcement learning and hierarchical RL expand the capabilities of agents to collaborate, compete, and learn abstract skills transferable across tasks. These developments hold great promise for robotics, autonomous vehicles and personalized recommendation systems, where agents must adapt dynamically to changing environments and objectives.

Hybrid and ensemble methods also represent a cutting-edge trend in machine learning. By combining multiple models or techniques, these methods aim to leverage the strengths of each component to improve overall performance and robustness. For example, integrating symbolic reasoning with neural networks sometimes referred to as neuro-symbolic AI can enhance interpretability and generalization, addressing some of the limitations of purely data-driven approaches. Ensemble learning techniques like boosting, bagging and stacking aggregate predictions from diverse learners to reduce variance and bias, often leading to superior predictive accuracy. Meta-learning or “learning to learn,” is another exciting approach where models acquire the ability to quickly adapt to new tasks with minimal data, mimicking human learning efficiency. This technique is particularly useful in few-shot or zero-shot learning scenarios, enabling AI systems to generalize better in real-world applications where labelled data is sparse or unavailable.

Scalability and efficiency have become critical focus areas in the development of cutting-edge machine learning methods, as real-world applications demand models that are

not only accurate but also computationally feasible. Techniques such as model pruning, quantization and knowledge distillation help reduce model size and inference latency, making deployment on edge devices and mobile platforms more practical. Moreover, federated learning has emerged as an innovative paradigm that allows multiple devices or organizations to collaboratively train models without sharing raw data, preserving privacy and security. This is particularly valuable in sensitive domains like healthcare and finance, where data confidentiality is paramount. Advances in hardware, including specialized AI accelerators like GPUs, TPUs, and neuromorphic chips, complement these algorithmic improvements, enabling faster training and real-time inference at scale.

Ethical considerations and explainability are increasingly integrated into cutting-edge machine learning research. As ML systems are deployed in critical areas such as criminal justice, hiring, and healthcare, the demand for transparency, fairness, and accountability grows. New techniques for interpretable machine learning strive to make complex models more understandable to humans, facilitating trust and enabling stakeholders to diagnose and mitigate biases. Explainability methods range from post-hoc analysis tools, like SHAP and LIME to inherently interpretable model architectures designed to provide insights into decision-making processes. Additionally, fairness-aware learning algorithms aim to detect and correct discriminatory patterns in data and predictions, promoting equitable outcomes. Addressing these challenges is essential for responsible AI adoption and to ensure that the benefits of cutting-edge ML methods are shared widely and ethically.

Cutting-edge machine learning methods encompass a diverse and rapidly advancing set of techniques that push the boundaries of what artificial intelligence can achieve. From deep learning and reinforcement learning breakthroughs to innovative unsupervised approaches and hybrid models, these methods empower machines to understand, generate and act upon data with increasing sophistication. Coupled with ongoing efforts to improve scalability, efficiency, transparency and fairness, cutting-edge ML is transforming industries and society at large. As research progresses and these methods become more accessible, the future of machine learning promises even greater impact, enabling smarter, more adaptable and ethically aligned AI systems.

Sanjay Kumar N

I B.Sc. (Computer Technology)

AI IN DATA ANALYTICS

Artificial Intelligence (AI) is transforming industries across the globe, and one of its most impactful applications is in the field of data analytics. As businesses generate vast amounts of data, traditional methods of analysis struggle to keep up. AI offers powerful tools to extract insights faster, more accurately, and at a scale never seen before. The modern digital ecosystem produces massive volumes of data every second from customer interactions to IoT devices. Traditional analytics methods rely heavily on manual processes and predefined models, which are often limited by human bias and capacity.

AI, particularly Machine Learning (ML) and deep learning, allows systems to:

- Learn patterns from historical data.
- Make predictions without explicit programming.
- Continuously improve with new data inputs.

Key Applications of AI in Data Analytics

a. Predictive Analytics

AI can predict future trends based on historical data. For example, it helps retailers forecast inventory needs or financial institutions predict credit risk.

b. Natural Language Processing (NLP)

With NLP, AI can analyze unstructured data such as emails, social media and customer feedback to understand sentiment and extract insights.

c. Anomaly Detection

AI models can quickly identify outliers or unusual behavior in large datasets critical for fraud detection, network security or equipment failure in manufacturing.

d. Automated Insights

AI-powered tools can automatically generate insights from dashboards, removing the need for users to interpret complex data themselves.

Benefits of AI-Driven Analytics

Speed: Real-time analysis and reporting.

Accuracy: Reduced human error and bias.

Scalability: Handles vast and complex datasets.

Efficiency: Automates repetitive tasks like data cleaning and reporting.

Challenges and Considerations

While AI offers great promise, there are hurdles:

Data Quality: Poor data can lead to inaccurate models.

Bias and Fairness: AI can perpetuate existing biases if not carefully monitored.

Explainability: Some AI models (like deep learning) are "black boxes" and hard to interpret.

Skills Gap: There’s a shortage of professionals skilled in both AI and analytics.

The Future of AI in Data Analytics

AI's role in analytics will continue to grow with advancements in:

Explainable AI (XAI): Making AI models more transparent and interpretable.

Edge AI: Bringing analytics to the edge for real-time decisions in IoT.

AutoML: Tools that automate model creation and deployment, democratizing AI use.

Ultimately, AI will not replace data analysts but augment them, enabling smarter decisions and allowing professionals to focus on strategic tasks. The integration of AI into data analytics is not just a trend it's a fundamental shift in how organizations derive value from data. As AI becomes more accessible and powerful, businesses that harness its potential will gain a significant competitive edge.

Gurucharan R

I B.Sc. (Information Technology)

AUTOMATED MACHINE LEARNING (AUTOML)

Automated Machine Learning (AutoML) refers to the process of automating the complex and iterative tasks involved in the application of machine learning to real-world problems. Traditionally, building a machine learning model requires extensive expertise in data science, statistics and computer science. This includes steps such as data preprocessing, feature engineering, algorithm selection, hyperparameter tuning, model evaluation and sometimes deployment. AutoML aims to democratize machine learning by automating many of these tasks, making it accessible to users with limited ML knowledge while also helping experts work more efficiently.

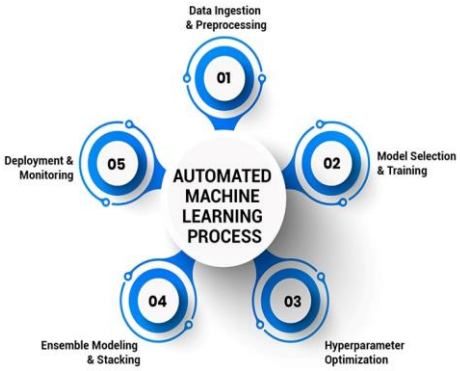

At its core, AutoML seeks to reduce the trial-and-error process that data scientists go through when creating models. One of the first steps in any ML pipeline is data preprocessing. This involves handling missing values, encoding categorical features, normalizing or standardizing numerical values and sometimes reducing dimensionality through feature selection or extraction. AutoML tools automate much of this, allowing users to focus more on their business goals rather than the technical intricacies of preparing data. Moreover, these tools automatically explore a variety of machine learning algorithms such as decision trees, support vector machines, ensemble

methods and even deep learning models, to find the most suitable one for a specific task.

Another critical area of automation in AutoML is hyperparameter optimization. This refers to the process of finding the best combination of parameters for a given algorithm to maximize performance. AutoML systems often employ strategies like grid search, random search, or more advanced methods such as Bayesian optimization to navigate the hyperparameter space efficiently. After training, the models are evaluated on appropriate metrics like accuracy, precision, recall or mean squared error depending on the task at hand, whether it's classification, regression, or clustering. Some AutoML frameworks also implement model ensembling, which involves combining predictions from multiple models to improve robustness and accuracy.

Numerous platforms and libraries have been developed to support AutoML. Google Cloud AutoML provides a suite of tools for vision, language and structured data tasks, while Microsoft Azure AutoML integrates well with other Azure services for enterprise-grade solutions. Open-source libraries such as H2O.ai's AutoML, Auto-sklearn, TPOT (Tree-based Pipeline Optimization Tool), and AutoKeras are widely used by the community for various applications. These tools differ in their flexibility, user-friendliness and scope. For instance, Auto-sklearn builds on top of the popular scikit-learn library and is well-suited for traditional ML tasks whereas AutoKeras is built for deep learning and neural networks.

Despite its many advantages, AutoML is not without limitations. One of the primary concerns is the lack of transparency and interpretability in the models it generates, especially when complex ensembling techniques or deep learning models are involved. In high-stakes fields such as healthcare, finance or criminal justice, understanding how a model arrives at a decision is crucial. Moreover, while AutoML is designed to reduce human involvement, it is not a substitute for domain knowledge. Understanding the context of the data and problem remains essential to building meaningful ML solutions. Additionally, AutoML systems can be computationally intensive, as they often train and evaluate a large number of models in search of the best one.

AutoML has found applications in a wide range of industries and use cases. In finance, it helps detect fraudulent transactions by automatically building models that flag suspicious patterns. In marketing, it can predict customer churn or segment audiences based on behavioral data. In healthcare, AutoML is being used to assist in diagnostic predictions, such as identifying diseases from medical images or patient records. The ability of AutoML systems to rapidly prototype and

deploy models makes them especially useful in environments that require quick, data-driven decisions.

Automated Machine Learning is transforming the way organizations and individuals interact with machine learning technology. By automating many of the technical tasks traditionally handled by data scientists, AutoML lowers the barrier to entry and accelerates the development of AI-driven solutions. However, it is not a silver bullet. Successful use of AutoML still requires thoughtful data preparation, clear problem definition and careful interpretation of results. As AutoML continues to evolve, its role in augmenting human intelligence rather than replacing it is becoming increasingly clear, pointing to a future where collaboration between humans and machines becomes the norm in data science and analytics.

P.Logesh

III B.Sc. (Information Technology)

APACHE KAFKA AND EVENT STREAMING

Apache Kafka is an open-source distributed event streaming platform originally developed by LinkedIn and later open-sourced through the Apache Software Foundation. Today, it stands as one of the most robust and scalable technologies for handling real-time data feeds. Thousands of companies major players like Uber, Netflix, LinkedIn, Cisco and Goldman Sachs rely on Kafka to manage vast flows of real-time information across systems and applications. Kafka's primary function is to enable high-throughput, low-latency transmission of messages between producers (sources of data) and consumers (applications or services that use the data). Kafka was designed to address the need for durable and scalable messaging systems, particularly where traditional message queues such as RabbitMQ or ActiveMQ fall short in terms of fault tolerance, throughput and horizontal scalability.

Kafka's core abstraction is the concept of an "event" a record of something that happened, represented as a key, value and timestamp. Events are organized into topics which are like channels or categories that help segment and direct streams of data to the appropriate consumers. Producers publish events to these topics, and consumers subscribe to them in real time or retrospectively. A powerful feature of Kafka is its retention policy; rather than deleting messages immediately after consumption (as traditional message queues do), Kafka retains data for a configured duration, allowing multiple consumers to reprocess data at different times without disrupting the system. This makes Kafka not only a message broker but also a kind of distributed commit log, which can be

invaluable for building fault-tolerant, state full applications.

Another critical concept in Kafka is its distributed architecture. Kafka runs as a cluster of brokers, each of which is responsible for handling read and write operations for a portion of the data. Topics are split into partitions, which allow Kafka to horizontally scale by distributing workload across multiple brokers. Partitions also enable parallelism, as consumers can read from different partitions simultaneously. Kafka guarantees message ordering within a partition, and with its replication mechanism, each partition can be mirrored across several brokers to ensure high availability and fault tolerance. This architecture makes Kafka suitable for mission-critical applications that require strong reliability such as payment processing, system monitoring, log aggregation and fraud detection.

Kafka’s ecosystem includes several additional components that enhance its functionality. Kafka Streams is a lightweight library that allows developers to build real-time, stream-processing applications directly on top of Kafka without requiring an external processing framework like Apache Flink or Spark Streaming. Kafka Streams supports complex event processing such as joins, aggregations and windowed operations while maintaining exactly-once semantics a key requirement for financial and transactional systems. Another important tool in the Kafka ecosystem is Kafka Connect, a framework for integrating Kafka with external systems such as databases, key-value stores, file systems and cloud services. Kafka Connect allows developers to easily build and manage scalable and fault-tolerant ingestion pipelines using pre-built connectors or custom plugins.

Kafka's flexibility enables a wide range of use cases. In e-commerce platforms, Kafka can be used to track user activity in real time, allowing for dynamic personalization and targeted advertising. In finance, it enables the real-time analysis of market data and transaction streams. Logistics and transportation companies use Kafka to process IoT sensor data from fleets of vehicles, optimizing delivery routes and predicting maintenance issues. In media and entertainment, Kafka powers recommendation engines, audience analytics and live content personalization. Kafka is also widely used in micro services architectures as a central backbone for inter-service communication, promoting decoupling, scalability and resilience.

Despite its many advantages, Kafka is not a silver bullet. Deploying Kafka in production requires careful planning around topics, partitions, disk usage and network configurations. Operational complexity can increase with scale and ensuring proper monitoring and tuning is crucial to maintaining

performance and reliability. Additionally, Kafka is not a traditional message queue in the sense of providing guaranteed per-message delivery semantics like at-most-once or at-least-once in all scenarios. Developers must design applications carefully to manage duplicate messages, retries and message ordering across partitions.

The rise of event-driven architecture has placed Kafka at the heart of many digital transformation strategies. Event-driven systems react to changes in real time, allowing organizations to become more agile, responsive, and data-driven. Kafka's ability to ingest, store, process and forward events in a reliable and scalable way makes it an essential tool for building such architectures. By decoupling producers and consumers and enabling asynchronous communication, Kafka fosters a modular approach to system design, where new components can be added or changed with minimal impact on the rest of the infrastructure.

Apache Kafka is more than just a messaging system it is a comprehensive platform for building real-time data pipelines and streaming analytics applications. With its distributed, fault-tolerant, and high-throughput design, Kafka is well-suited for the demands of modern enterprises that rely on real-time data to gain competitive advantage. Its broad ecosystem, support for stream processing, and growing community adoption make it a cornerstone technology in today's event-driven world. As data continues to be generated at unprecedented volumes and velocities, the need for scalable and resilient event streaming platforms like Kafka will only continue to grow.

A.P.Anbu

II B.Sc. (Information Technology)

CLOUD DATA WAREHOUSES

A Cloud Data Warehouse (CDW) is a modern data storage solution that enables organizations to consolidate and analyze vast amounts of data by leveraging cloud computing infrastructure. Unlike traditional on-premises data warehouses, which require significant upfront hardware investments and ongoing maintenance, cloud data warehouses provide scalable, flexible and cost-effective platforms hosted by cloud service providers. This shift to the cloud offers businesses the ability to dynamically allocate resources, manage data storage and run complex analytics without the constraints of physical infrastructure. Popular examples of cloud data warehouses include Amazon Redshift, Google BigQuery, Snowflake and Microsoft Azure Synapse Analytics. These platforms have revolutionized the way enterprises handle data, allowing them to perform large-scale data analytics with speed and agility, driving better decision-making and innovation.

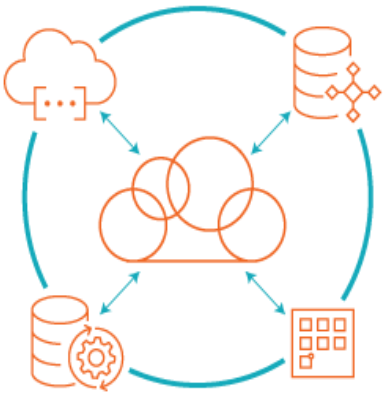

One of the fundamental advantages of cloud data warehouses is their scalability**.** Organizations can easily increase or decrease their computational and storage resources based on demand, paying only for what they use. This elasticity contrasts sharply with traditional warehouses, which often require over-provisioning to handle peak workloads, leading to wasted resources and higher costs during normal periods. Additionally, cloud data warehouses separate compute and storage layers, enabling simultaneous scaling of each independently. This architectural feature means companies can handle large volumes of data ingestion and complex querying efficiently, without bottlenecks. For instance, during heavy analytical workloads, additional compute nodes can be spun up temporarily and then scaled down when the demand subsides, optimizing both performance and cost.

Cloud data warehouses also support a variety of data types and sources, facilitating the integration of structured, semi-structured, and even unstructured data from diverse origins such as transactional databases, logs, IoT devices and third-party APIs. They provide built-in tools for data ingestion, transformation, and orchestration, often compatible with ETL (Extract, Transform, Load) or ELT (Extract, Load, Transform) workflows. Modern cloud warehouses often offer native support for JSON, Avro, Parquet, and other semi-structured formats, enabling seamless analytics on data beyond traditional relational tables. This versatility allows organizations to consolidate multiple data silos into a single repository, fostering more comprehensive and insightful analysis. Moreover these platforms are cloud-native, they integrate well with other cloud services such as machine learning frameworks, business intelligence tools and data lakes, creating robust ecosystems for end-to-end data processing and analytics.

Another key benefit of cloud data warehouses is performance optimization through advanced query engines and Massively Parallel Processing (MPP). By distributing queries across multiple nodes in a cluster, these systems can handle large datasets and complex SQL queries with remarkable speed. Many cloud warehouses leverage columnar storage formats that optimize data compression and input/output efficiency, accelerating read operations for analytical queries. Additionally, features like materialized views, result caching and automatic query optimization further enhance performance. For example, Snowflake’s unique multi-cluster architecture allows for concurrent workloads without

contention, supporting numerous users and applications simultaneously. Furthermore, cloud providers continuously update their platforms with new optimizations and features, ensuring users benefit from the latest advancements without manual upgrades or downtime.

Security and compliance are paramount concerns for organizations dealing with sensitive or regulated data, and cloud data warehouses address these needs robustly. Providers implement strong encryption protocols both at rest and in transit, identity and access management controls, and audit logging to ensure data privacy and integrity. Compliance certifications such as SOC 2, HIPAA, GDPR, and FedRAMP are common among leading cloud data warehouse services, giving organizations confidence that their data meets stringent regulatory standards. Additionally, role-based access controls and data masking features help protect sensitive information while enabling appropriate data sharing and collaboration within an organization. Since cloud platforms are managed by specialized teams, security patches and updates are deployed automatically, reducing the operational burden on internal IT staff.

Cost management is another critical aspect where cloud data warehouses provide significant advantages. The pay-as-you-go pricing model allows organizations to avoid large capital expenditures and reduce total cost of ownership. Instead, expenses are operationalized as ongoing service fees based on actual usage, which improves budget predictability. Many platforms also offer cost monitoring and alerting tools to help users track spending and optimize resource utilization. By eliminating the need for dedicated hardware, reducing maintenance costs, and enabling rapid deployment of new projects, cloud data warehouses empower businesses of all sizes from startups to large enterprises to innovate faster and stay competitive in data-driven markets.

The rise of cloud data warehouses has also impacted the role of data engineers and analysts. With infrastructure management abstracted away, data professionals can focus more on extracting insights and developing analytical models rather than worrying about system uptime or capacity planning. The availability of SQL-based querying interfaces alongside integration with programming languages like Python and R allows for flexible and powerful data analysis workflows. Moreover, cloud data warehouses facilitate collaboration across departments by providing centralized, real-time data access, breaking down traditional data silos. They also enable advanced analytics such as predictive modelling, AI and machine learning by integrating seamlessly with cloud-native AI services, expanding the scope and impact of data initiatives.

Cloud data warehouses represent a transformative approach to data storage and analytics, offering scalability, performance, security and cost efficiency that traditional systems struggle to match. By harnessing the power of the cloud, organizations gain the agility to adapt to changing business needs, unlock deeper insights from diverse data sets and accelerate innovation. As data volumes continue to grow and analytics become increasingly central to business strategy, cloud data warehouses are poised to remain foundational components of modern data architectures for years to come.

Sanjay Kumar N

I B.Sc. (Computer Technology)

EDGE AI

Edge AI refers to the deployment of artificial intelligence (AI) algorithms and models directly on edge devices such as sensors, smartphones, gateways or other IoT hardware instead of relying solely on centralized cloud servers for processing. This approach brings the power of AI closer to where data is generated, enabling real-time analysis, faster decision-making and reduced dependency on constant internet connectivity. By processing data locally, edge AI significantly reduces latency, lowers bandwidth usage and improves privacy since sensitive data doesn't need to be sent to the cloud.

In IoT ecosystems, edge AI is transforming how devices operate by allowing them to perform tasks like anomaly detection, predictive maintenance, image recognition and voice processing autonomously. For instance, a smart camera with embedded edge AI can detect security threats immediately and trigger alerts without waiting for cloud processing. This not only speeds up responses but also enhances reliability in environments with intermittent connectivity.

Advancements in hardware such as specialized chips called Neural Processing Units (NPUs) and energy-efficient microcontrollers have made it feasible to run complex AI models on resource-constrained devices. Edge AI is also essential for applications requiring real-time insights or operating in remote locations where cloud access is limited or too slow. Overall, edge AI is a critical enabler of smarter, faster and more secure IoT systems, supporting a wide range of industries including healthcare, manufacturing, transportation and smart cities.

Importance of Edge Computing

1. **Low Latency and Real-Time Processing**

   Many IoT applications like autonomous vehicles, industrial automation or healthcare monitoring require instant responses. Waiting for data to travel to the cloud for processing can cause

delays that are unacceptable in these scenarios. Edge AI allows devices to analyse data immediately as it is collected, making split-second decisions possible.

2. **Bandwidth**
   Transmitting massive amounts of raw data to the cloud can overload networks and increase costs. Edge AI reduces the amount of data sent upstream by performing initial filtering, summarizing or even making decisions locally, thus lowering network traffic.
3. **Enhanced Privacy and Security**
   Sensitive data, such as personal health information or surveillance footage, may be too confidential to send to cloud servers. Processing this data locally helps protect privacy and reduces risks of data breaches.
4. **Reliability in Connectivity-Limited Environments**
   In remote or mobile settings where internet access is unreliable or unavailable (e.g., on factory floors, rural areas, or vehicles), edge AI enables continuous operation without dependency on cloud connectivity.

Technical Enablers of Edge AI

- **Specialized Hardware**: Recent advances have led to powerful, low-power chips designed for AI inference on edge devices. Neural Processing Units (NPUs), Graphics Processing Units (GPUs), and Field Programmable Gate Arrays (FPGAs) optimized for AI accelerate computations while maintaining energy efficiency.
- **Model Optimization**: AI models must be optimized to run efficiently on devices with limited memory and processing power. Techniques like model pruning, quantization, and knowledge distillation help shrink model size without sacrificing accuracy.

B.Manju Bashini
II B.Sc. (Computer Technology)

EMPOWERING ROBOTS WITH HUMAN-LIKE PERCEPTION TO NAVIGATE UNWIELDY TERRAIN

As robots are increasingly deployed in complex, real-world environments ranging from disaster zones to planetary surfaces they must be capable of perceiving and interpreting challenging terrain much like humans do. This work explores how to integrate human-inspired perceptual models with advanced robotic control systems to improve navigation in unstructured and unpredictable environments. The goal is to bridge the gap between current robotic sensing capabilities and the nuanced perception strategies employed by humans, particularly in difficult terrains such as rubble, dense forest or rocky landscapes.

Human-Like Perception

- **Multi-modal sensory integration** (vision, tactile feedback, inertial sensing).
- **Semantic understanding** of terrain types (e.g., "slippery," "unstable").
- **Predictive modelling**: anticipating movement outcomes based on past experience.

Technological Enablers

- Deep learning for visual terrain classification and affordance learning.
- Simultaneous Localization and Mapping (SLAM) enhanced by semantic cues.
- Imitation learning and reinforcement learning from human demonstrations.
- Event-based vision and neuromorphic **sensing** for real-time adaptation.

Robotic Platforms

- Legged robots (e.g., Boston Dynamics' Spot, MIT's Mini Cheetah).
- Hybrid systems with active perception mechanisms.
- Soft robotics for adapting to terrain deformation.

Challenges

- Sensor noise and degradation in harsh environments.
- Real-time processing constraints.
- Generalizing across terrains without retraining.

Case Studies / Applications

- Urban search-and-rescue robots navigating collapsed buildings.
- Martian rovers with terrain-adaptive perception.
- Autonomous quadrupeds navigating forest trails.

Potential Research Directions

- Cognitive models for terrain perception.
- Self-supervised learning from terrain interactions.
- Adaptive control based on terrain difficulty estimates.
- Fusion of physics-based simulation with real-world sensory data.

S.Dharshini

II B.Sc. (Information Technology)

TYPESCRIPT (TYPED SUPERSET OF JAVASCRIPT)

TypeScript is a typed superset of JavaScript developed by Microsoft that compiles to plain JavaScript. It enhances JavaScript with static typing**,** type inference and modern object-oriented features like interfaces and classes. The main goal of TypeScript is to provide developers with better tooling and early error detection, especially for large-scale JavaScript applications. Since TypeScript code is transpiled to standard

JavaScript, it runs anywhere JavaScript runs browsers, Node.js or embedded systems.

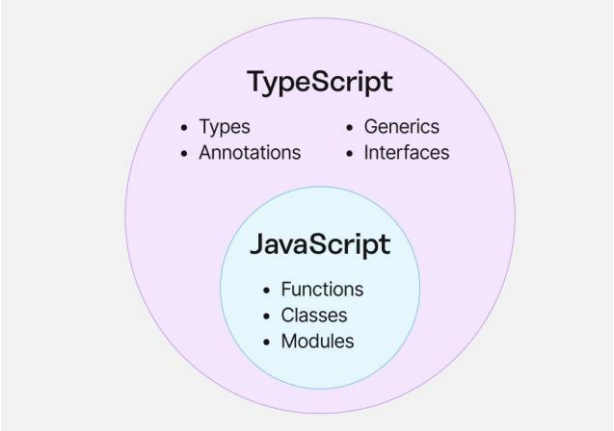

One of the key features of TypeScript is its type system. Type annotations can be added to variables, function parameters, and return values. For example, a variable can be declared as let age: number = 25, which ensures that age only holds numeric values. Common types include string, number, boolean, any, unknown, void, null and undefined. TypeScript also supports arrays (string[]) and tuples ([string, number]) for fixed-length collections with defined types.

TypeScript supports enums**,** which are a way of defining named constants. For instance, enum Direction { Up, Down, Left, Right } allows the use of named directions instead of numeric values. Functions can include parameter and return type annotations, such as function add(x: number, y: number): number. Optional parameters are marked with ?, and default values can be assigned like in JavaScript. For example, function greet (name: string = "Guest") provides a default name when none is passed.

In TypeScript, objects can be strongly typed. You can define an object with a specific structure or shape, such as let user: { name: string; age: number }. For more reusable code, **interfaces** are used. An interface defines a contract for objects and classes, describing required properties and methods. For example, interface Person { name: string; age: number; greet(): void } can be implemented by any object or class that fits this structure. Interfaces also support extension and optional properties, making them very flexible for defining APIs.

Another powerful feature of TypeScript is its support for classes**,** which bring object-oriented programming concepts such as inheritance, access modifiers (public, private, protected), and constructors. A class like class Animal { constructor (public name: string) {} } automatically creates and assigns class properties. Subclasses can extend base classes using the extends keyword and override methods using super.

Generics are another cornerstone of TypeScript, enabling the creation of reusable components that work with a variety of types. A generic function, like function identity<T>(arg: T): T, can be used with any type without losing type information. This is especially useful in data structures like lists, stacks, or service wrappers where the type may

vary. one can also create generic interfaces, such as interface Box<T> { value: T }, which enforce type safety across a wide range of use cases.

Type aliases allow one to create a custom name for a type. For example, type ID = number | string means a variable of type ID can be either a number or a string. This is a clean way to manage union types and increase readability. TypeScript also includes built-in utility types like Partial<T>, Readonly<T>, and Record<K, T>, which make it easy to transform existing types in useful ways. For example, Partial<T> marks all properties of T as optional.

TypeScript supports type narrowing, which refines the type of a variable within a control structure. For instance, if a parameter has a union type (number | string), one can use typeof or instanceof checks to narrow the type and access type-specific methods. This helps avoid runtime errors and improves type safety. Modules in TypeScript follow the ES6 module system. You can use export and import statements to share code across files. For example, a function can be exported with export function add() and imported with import { add } from './math'.

To configure a TypeScript project, the tsconfig.json file is used. It defines compiler options such as the target JavaScript version ("target": "es6"), module system ("module": "commonjs"), and strictness settings ("strict": true). TypeScript integrates well with modern development tools, enabling powerful IDE features like auto completion, type inference and inline documentation. TypeScript brings structure and safety to JavaScript development while remaining closely aligned with JavaScript syntax and behaviour. Its powerful type system, tooling support and modern programming features make it an excellent choice for building robust, maintainable applications at scale.

M.S.K Manassha
III B.Sc. (Information Technology)

GENERATIVE AI

In recent years, Generative AI has emerged as a transformative force in the field of artificial intelligence, ushering in novel possibilities for content creation, data generation and complex problem-solving. At the heart of generative AI lies the ability of models to autonomously create new data that mimics the patterns, styles and structures of real-world examples. The most notable advancements within this domain have come in the form of Generative Adversarial Networks (GANs) and more recently, Diffusion Models both of which are reshaping the creative and scientific landscapes.

Generative AI: Paving the Path for Creativity

Generative AI refers to a class of machine learning techniques designed to generate new data. Unlike traditional AI systems, which are primarily trained to recognize patterns in data (i.e., discriminative tasks like classification or regression), generative models learn the underlying distribution of data and use this understanding to produce new instances that resemble the input data. These models are especially useful in applications where new, realistic content is needed, such as art creation, content generation and even drug discovery.

The most famous of the early generative models is Generative Adversarial Networks (GANs), introduced by Ian Goodfellow in 2014. GANs consist of two neural networks a generator and a discriminator that are trained in a competitive setting. The generator creates fake data, while the discriminator tries to distinguish between real and fake data. Over time, both networks improve, with the generator becoming increasingly adept at producing data that is indistinguishable from real-world examples. GANs have been pivotal in generating high-quality images, music and even deep fakes, showcasing their potential in creative industries. In recent years, however, Diffusion Models have gained prominence as the next frontier in generative AI, offering significant improvements over GANs in terms of stability, image quality and ease of training.

Diffusion Models: A New Era of Image and Data Generation

Diffusion models are a newer class of generative models that operate on the principle of simulating a diffusion process the gradual transformation of data into random noise and then learning to reverse this process to regenerate the original data. These models work by taking a sample of data (such as an image) and gradually adding noise to it in a series of steps until it becomes pure noise. The model then learns to reverse this noisy process, step by step, until it reconstructs the original data. The goal is for the model to learn how to denoise and recover the original data distribution.

One of the defining features of diffusion models is their remarkable ability to produce high-fidelity images with greater consistency and fewer artifacts compared to GANs. Early examples of diffusion models, such as DDPM (Denoising Diffusion Probabilistic Models) and Score-based Models, demonstrated that these models could generate photorealistic images with better diversity and more stability in training. In recent years, models like Stable Diffusion, DALL·E 2 and MidJourney have gained widespread attention for their ability to generate high-quality images from textual descriptions showing the vast

potential of diffusion models in creative applications.

Unlike GANs, which are prone to training instability and mode collapse (where the generator fails to produce diverse samples), diffusion models are much more stable and easier to train. This is largely due to the way they model data transformation, using a more gradual, continuous process that doesn't rely on the adversarial setup of GANs. The model learns to denoise in multiple stages, making it less prone to failure and improving its ability to generate a diverse range of outputs.

Impact on Creative Industries

The advent of diffusion models has revolutionized creative fields, enabling artists, designers, and content creators to leverage AI tools for generating high-quality images, videos, music and more. Tools like DALL·E 2**,** Stable Diffusion and MidJourney allow users to input textual prompts and generate images with remarkable accuracy, often indistinguishable from human-created artwork. These models have democratized creativity by making high-end design accessible to non-experts, thus opening new avenues for creative expression and production.

Beyond artistic applications, generative AI and diffusion models are being used in industries such as fashion, film and architecture. In fashion, generative models can create new clothing designs based on trends or specific parameters, helping designers quickly prototype and iterate on concepts. In film and animation, diffusion models can generate backgrounds, characters and even entire scenes significantly reducing production costs and time. In architecture, generative models are used to visualize building designs and experiment with novel architectural forms.

Moreover, generative models are also enhancing game development. Game designers use AI to create procedurally generated landscapes, character models and in-game objects, thereby creating dynamic ever-changing environments without needing to manually design every element.

Applications in Science and Technology

The impact of generative AI and diffusion models extends far beyond the creative industry. In the healthcare and life sciences sector, generative models are being explored for drug discovery and protein folding. By generating molecular structures that adhere to certain properties, these models can suggest novel compounds that might be overlooked by traditional methods. For example, diffusion models could be used to generate new candidate drugs by simulating the molecular structures of compounds that are likely to bind with a particular protein or receptor.

In robotics and automation, generative models are being used to simulate

environments and design robotic systems that can adapt to a wide range of tasks. Robots can learn through simulation, with generative models providing synthetic data to train robots on tasks that might be difficult or dangerous to perform in real life. Similarly, in autonomous vehicles generative AI can be used to create realistic driving scenarios for simulation-based training, improving safety and decision-making algorithms.

In natural language processing (NLP), models like GPT-4 have shown how generative models can be used to generate human-like text. Combining these models with diffusion-based approaches for text-to-image or text-to-video generation can lead to rich, multimodal content creation tools. The ability to generate not just text, but also accompanying images or videos from simple text prompts, holds the potential to change how we produce content for education, marketing, and entertainment.

Ethical Considerations and Challenges

While generative AI and diffusion models offer incredible possibilities, they also raise significant ethical concerns. Deepfakes and other manipulated media generated by AI models have sparked debates around privacy, security and misinformation. With the ability to generate highly convincing images, videos and voices, there are growing concerns about the misuse of these technologies for malicious purposes, such as spreading false information, creating harmful content, or impersonating individuals.

Moreover, the potential for bias in generative models remains a critical issue. These models learn from vast datasets, which may contain inherent biases reflecting societal inequalities. If these biases are not carefully mitigated, generative AI could reinforce stereotypes, perpetuate discrimination, or generate harmful content unintentionally.

Another challenge is the intellectual property concerns surrounding generative models. As these models generate content based on data from existing sources, the question arises as to who owns the content generated by AI. Legal frameworks surrounding the copyright of AI-generated works are still evolving, and there is a need for clearer definitions and protections in place.

Generative AI and diffusion models represent a leap forward in the ability of machines to create, simulate, and innovate in ways that were once thought to be exclusive to human creativity. These innovations are not just changing how we make art or design products but are also unlocking new potential in scientific discovery, healthcare and automation. However, as we move forward, it is essential that the ethical challenges and risks associated with these technologies are addressed thoughtfully and responsibly. By harnessing the power of generative models,

while mitigating their risks, we can ensure that these innovations benefit society in meaningful ways. The future of generative AI is bright and as diffusion models continue to evolve, they will likely play an even more central role in shaping the next wave of technological advancement.

K.Barath

II B.Sc. (Computer Technology)

ROBOT AS A SERVICE (RaaS): A REVOLUTIONARY BUSINESS MODEL

In recent years, the concept of Robot as a Service (RaaS) has emerged as a transformative model in the world of automation and robotics. RaaS is a cloud-based service where companies and individuals can lease robotic systems instead of investing in expensive hardware and software upfront. This subscription-based model enables businesses of all sizes to access advanced robotic technologies without the financial burden of purchasing, maintaining and upgrading them. RaaS leverages the power of cloud computing and IoT (Internet of Things) to connect robots to centralized platforms, where users can manage and control the robots remotely, monitor performance and receive real-time data analytics. This flexibility has made RaaS particularly attractive for industries such as manufacturing, logistics, healthcare, agriculture and even hospitality.

The business benefits of RaaS are vast and wide-reaching. For businesses, the cost-efficiency of leasing robots as opposed to owning them allows for better allocation of capital. Organizations can scale up or scale down their robotic needs depending on demand, without the risk of long-term commitment to a single, depreciating asset. Additionally, the subscription model ensures that businesses always have access to the latest advancements in robotics and AI technology, as updates and maintenance are typically handled by the RaaS provider. This eliminates the need for in-house expertise and minimizes downtime due to technical issues, ensuring maximum uptime and productivity. Furthermore, the ability to leverage robots for specific tasks such as automation of repetitive hazardous or complex processes enables businesses to improve efficiency, reduce human error and optimize operations.

Another key advantage of RaaS is its role in democratizing automation for small and medium-sized enterprises (SMEs). In the past, access to cutting-edge robotics technology was

mostly reserved for large corporations with deep pockets. However, the flexible, pay-as-you-go nature of RaaS has leveled the playing field, enabling SMEs to integrate automation solutions into their operations with minimal upfront investment. This democratization of robotics has the potential to increase productivity, enhance safety and allow smaller businesses to compete more effectively in their respective industries.

As the robotics market continues to mature, RaaS is expected to play an even more significant role in shaping the future of automation. The ability to rent robots for specific use cases also opens the door for increased experimentation and innovation. Companies can experiment with different robotic solutions without committing to the high costs of ownership. Moreover, industries like healthcare, where robots can assist in surgery or caregiving, benefit from the RaaS model by lowering entry barriers to high-tech medical solutions. As RaaS providers continue to innovate and expand their offerings, we can expect to see more industries adopt this model, pushing the boundaries of what robots can achieve in our daily lives.

Robot as a Service (RaaS) represents a paradigm shift in how businesses approach automation and robotics. By offering cost-effective, flexible and scalable solutions, RaaS is enabling companies of all sizes to integrate cutting-edge robotic systems into their operations, improving efficiency, reducing costs, and enabling new capabilities. As technology continues to evolve, RaaS will likely become an increasingly essential tool for businesses looking to stay competitive in an ever-changing market. With its combination of accessibility, affordability and innovation, RaaS is poised to play a central role in the future of automation across industries.

S.Dharshini

II B.Sc. (Information Technology)

PROGRESSIVE WEB APPLICATION

A progressive web application (PWA), or progressive web app is a type of application software delivered through the web, built using common web technologies including HTML, CSS, JavaScript and Web Assembly. It is intended to work on any platform with a standards-compliant browser, including desktop and mobile devices. Since a progressive web app is a type of webpage or website known as a web application, it does not require separate bundling or distribution. Developers can simply publish the web application online, ensure that it meets baseline installation requirements and ensure that users will be able to add the application to their home screen.

When one visit a website in the browser, it's visually apparent that the website is "running in the browser". The browser UI

provides a visible frame around the website, including UI features like back/forward buttons and a title for the page. The Web APIs one's website calls are implemented by the browser engine.

PWAs typically look like platform-specific apps they are usually displayed without the browser UI around them but as a matter of technology, still websites. This means they need a browser engine, like the ones in Chrome or Firefox, to manage and run them. With a platform-specific app, the platform OS manages the app, providing the environment in which it runs. With a PWA, a browser engine performs this background role, just like it does for normal websites.

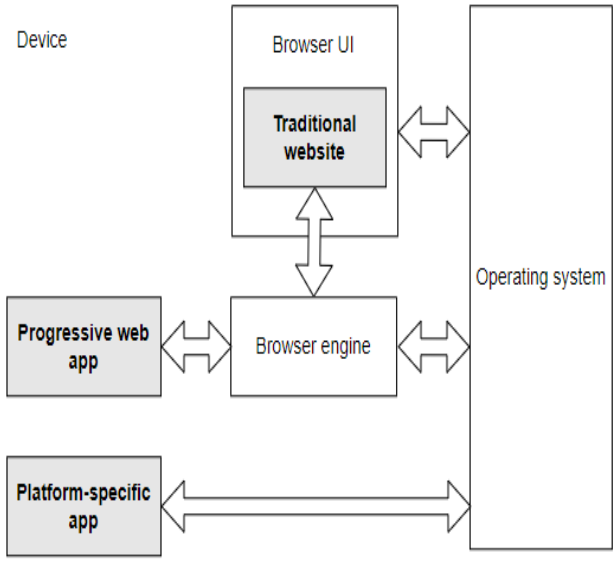

The browser starts a PWA's service worker when a push notification is received. Here, the browser's activity is entirely in the background. From the PWA's point of view, it might as well be the operating system that started it. For some systems, such as Chromebooks, there may not even be a distinction between "the browser" and "the operating system."

Advantages of PWA over traditional web and native apps

- The progressive web apps can work on any device with a browser. As a result, they eliminate the need for separate apps for different platforms.
- These apps work offline or with poor connectivity. As a result, they are providing a seamless user experience.
- The progressive web apps are faster and more responsive than traditional ones. Therefore, leading to higher user engagement.
- Users can access them directly from the web without downloading from an app store.

Sindhu M

I B.Sc. (Computer Technology)

AMBIENT IoT

Ambient IoT (Internet of Things) refers to the next evolution of IoT systems where smart sensing and connectivity are seamlessly embedded into the physical environment, operating quietly and continuously with minimal human interaction. Unlike traditional IoT, which often involves discrete devices like smart thermostats or fitness trackers, Ambient IoT uses ultra-low-

power or even battery-free sensors that are integrated directly into everyday objects packaging, furniture, infrastructure and more.

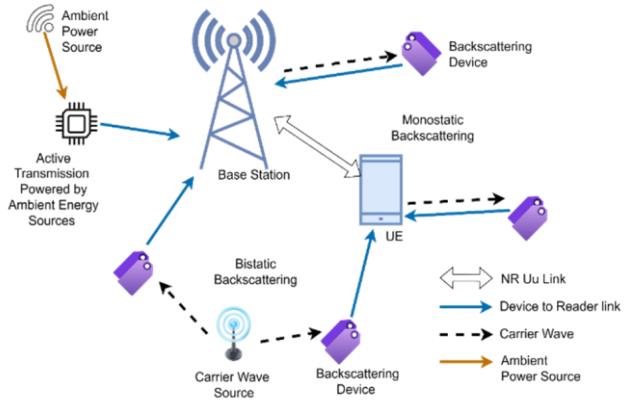

These sensors collect and transmit data about their surroundings using technologies such as RFID, Bluetooth Low Energy (BLE), Wi-Fi, backscatter communication, and even energy harvesting methods. The goal is to create an intelligent environment where physical objects can communicate autonomously, enabling real-time monitoring, automation and analytics at an unprecedented scale.

Ambient IoT is gaining attention due to its scalability, low energy consumption, and potential to revolutionize industries like supply chain, retail, healthcare and smart cities. For example, in a smart supply chain, products could tell warehouses and retailers where they are, their temperature history or whether they've been tampered with all without the need for bulky trackers or active maintenance. The technology is also aligned with 6G visions, as it supports massive connectivity, pervasive sensing and integrated intelligence across physical spaces. However, the deployment of Ambient IoT also raises new challenges, such as ensuring security, interoperability and standardization, especially as devices become more passive and decentralized. Ambient IoT represents a shift from smart devices to smart environments, where intelligence is ambient, invisible and deeply integrated into our surroundings.

V.B Krishna Prabu

III B.Sc. (Computer Technology)

"IT'S NOT THAT WE USE TECHNOLOGY, WE LIVE TECHNOLOGY."

- GODFREY REGGIOCFS