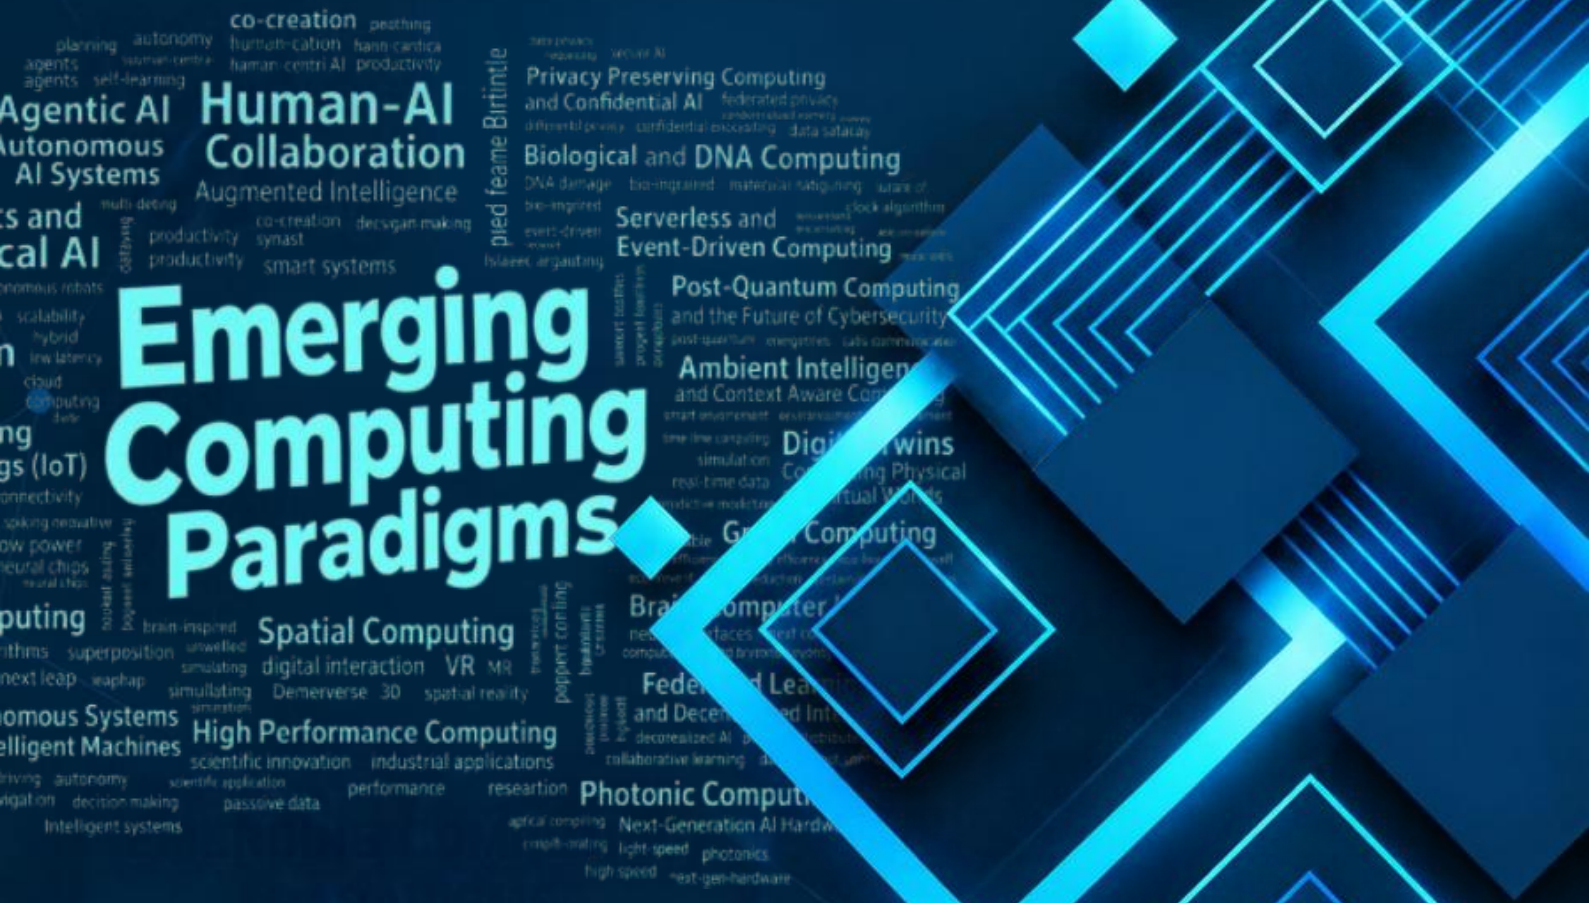
jims
Sec-5, Rohini, Delhi

TECHBYTE TB

INSIGHT

ANNUAL IT MAGAZINE



TECHBYTE 2026

— TECHBYTE 2026

"More than a magazine, TechByte is a platform where curiosity meets innovation, ideas take shape, and the technologies of tomorrow come to life"



JAGAN INSTITUTE OF MANAGEMENT STUDIES

SECTOR-5,ROHINI

Jagan Institute of Management Studies (JIMS), in its 30 years of existence, has created a niche for itself in the fields of Management and Information Technology. The Institute offers several undergraduate and postgraduate programmes, including BBA, BCA and B.A. (Hons.) Economics, affiliated with Guru Gobind Singh Indraprastha University.

The postgraduate programmes, namely PGDM, PGDM (International Business) and PGDM (Retail Management), offered at the campus are approved by the All India Council for Technical Education (AICTE) and accredited by the National Board of Accreditation (NBA). These programmes have also been granted equivalence to an MBA degree by the Association of Indian Universities (AIU).

The MCA programme is affiliated with Guru Gobind Singh Indraprastha University and is also accredited by the National Board of Accreditation (NBA). The National Assessment and Accreditation Council (NAAC) has accredited JIMS with an A++ grade, reaffirming its commitment to quality education and academic excellence.

The Institute is also among the top management colleges in Delhi to have started the Fellow Programme in Management (FPM), a full-time doctoral-level programme designed to equip students for careers in research and management. JIMS has been awarded the prestigious South Asian Quality Standards (SAQS) accreditation by the Association of Management Development Institutions in South Asia (AMDISA), a SAARC-recognized body. The Institute is also a member institution of AACSB, a prestigious international accreditation body.

JIMS continues to remain among the ELITE B-SCHOOLS (Top 100) of India in the National Institutional Ranking Framework (NIRF) of the Ministry of Education, Government of India.

TECHBYTE 2026

23rd Annual IT Symposium

"EMERGING COMPUTING PARADIGMS"

Department of Information Technology
Jagan Institute of Management Studies,
Sector-5 Rohini,
New Delhi-110085

DIRECTOR
Dr. Pooja Jain

PATRON
Dr. Praveen Arora
Principal JIMS, IPU-Affiliated Programmes

CONVENORS/EDITOR
Dr. Manjot Kaur Bhatia
Dr. Chetna Laroiya

STUDENT EDITORS
Khushi verma (MCA),
Tanya Mandal (MCA),
Namo Narayan Shukla (MCA)

DATE
Saturday, 10th October, 2026

VENUE
Auditorium, JIMS Sector-5, Rohini
New Delhi-85

Editor's Desk

What was once considered impossible is today becoming a reality, and what we see today is only a glimpse of what lies ahead. Technology is continuously transforming the way we think, learn, communicate and solve problems. As computing evolves from traditional systems towards Artificial Intelligence, IoT, Robotics, automation, Quantum Computing, the future demands not only technological knowledge but also the ability to understand and adapt to change.

At JIMS, we believe that education should extend beyond the boundaries of textbooks and classrooms. Our learning journey is not merely about acquiring knowledge, but about developing curiosity, creativity, innovation and the confidence to explore the unknown. We prepare our students for a world where technology and human potential come together.

Techbyte is a vibrant platform where students have the opportunity to gain insights into emerging technologies, interact with industry experts, exchange ideas and experience the changing landscape of the IT industry. The symposium enables us to understand how emerging computing paradigms are shaping industries and creating new possibilities for careers, innovation and entrepreneurship.

The Annual IT Magazine is an extension of this spirit. Through its pages, we bring together the ideas, perspectives, creativity and technological aspirations of our students. It is our endeavour to capture the excitement of emerging technologies and encourage every reader to think beyond what exists and imagine what could be.

We express our heartfelt gratitude to the Management of JIMS, Principal, faculty members, industry experts and everyone associated with Techbyte for their constant encouragement and support.



Director's Message

The evolution of technology over the past decade has ushered in a new era of innovation, characterized by emerging computing paradigms that are transforming the way organizations operate, innovate, and create value. Advances in areas such as artificial intelligence, cloud computing, edge computing, quantum computing, and other disruptive technologies are redefining technological possibilities while creating an increasing demand for specialized skills and future-ready talent.

As industries navigate the transition towards Industry 5.0, computing is evolving beyond traditional digital transformation to embrace increased intelligence, adaptability, and human–technology collaboration. These emerging paradigms are not only reshaping technological landscapes but are also influencing business models, workforce requirements, and talent management strategies. In this dynamic environment, continuous learning, upskilling, and reskilling have become essential for students to contribute meaningfully to the future of industry.

At JIMS, we recognize the importance of preparing our students to navigate this rapidly changing technological landscape. As an institution committed to academic excellence and industry relevance, we continuously encourage and nurture our students to explore emerging technologies, think innovatively, engage with new ideas, and develop the capabilities required to address the challenges and opportunities of the digital future.

It is therefore encouraging to see *Techbyte* provide a meaningful platform for students, professionals, and academicians to exchange perspectives and explore developments in the field of computing and technology. I congratulate the entire *Techbyte* team for bringing together this insightful edition on **"Emerging Computing Paradigms"** and wish them all the best for their continued efforts.

Dr. Pooja Jain
Director



Principal's Message

Technology is advancing rapidly, and computing is no longer limited to making computers faster or processing more data. Emerging technologies such as Quantum Computing, Edge Computing, Internet of Things, Robotics etc. are changing how problems are solved and how technology interacts with the world around us. The annual IT seminar on "Emerging Computing Technologies" is therefore both relevant and timely.

The seminar provides an opportunity for students to understand these emerging technologies and helps them explore why these technologies are needed, where they are being used, and how they may shape different industries in the years ahead. Such exposure encourages students to connect classroom learning with real-world applications and to look at technology with greater curiosity and a practical perspective.

I am also pleased to see the Insight Magazine being published alongside the seminar every year. It gives our students a platform to explore emerging technological trends, express their ideas, and share their understanding of how these developments may influence industry and society. The articles reflect their curiosity, independent thinking, and willingness to engage with the rapidly changing technology landscape.

I congratulate all the student contributors for their efforts and appreciate the faculty members who have guided them throughout this initiative. I hope the seminar and the Insight Magazine encourage our students to remain curious, keep learning, and develop the ability to understand not only what technology can do, but also how and where it can make a meaningful difference.

Dr. Praveen Arora
Principal, JIMS – IPU Affiliated Programme

TECHBYTE

CONTENTS

Department of Information Technology

	ARTICLE TITLE	PAGE NO
1	Digital Twins: Connecting Physical and Virtual Worlds	1
2	Green Computing And Sustainable AI Infrastructure	6
3	When Technology Enters Your World: The Rise of Spatial Computing	10
4	Biological and DNA Computing: The Future of Information Processing	14
5	Cloud-Edge Continuum: The Future of Distributed Computing	18
6	Post-Quantum Computing and the Future of Cybersecurity	22
7	Autonomous Systems And Intelligent Machines	26
8	Privacy-Preserving Computing and Confidential AI	30
9	Agentic AI And Autonomous AI Systems	34
10	PHOTONIC COMPUTING AND NEXT-GENERATION AI HARDWARE	39
11	Federated Learning and Decentralized Intelligence	42
12	Ambient Intelligence and Context-Aware Computing	48
13	Quantum Computing: The Next Leap in Computation	51
14	Neuromorphic Computing: Designing Machines That Think Like the Brain	57
15	Brain – Computer Interfaces: Computing Beyond Screens	61

Digital Twins: Connecting Physical and Virtual Worlds

Faculty Mentor:
Dr. Deepti Sharma

Students Name:
Garv Batra (MCA-I)
Prashant (MCA-I)

INTRODUCTION

Digital twin technology is quietly becoming one of the backbones of modern digital transformation. At its core, it's simple: build a virtual copy of something real — an object, a system, a process — and keep that copy alive with data. That's the part a basic 3D model can't do. A 3D model just sits there, looking like the thing. A digital twin listens. It pulls in real-time information from sensors and connected devices, so the virtual version actually tracks what's happening to its physical counterpart. Combine that with IoT, AI, cloud computing, and increasingly sophisticated simulation, and you get systems that are getting more capable every year. The market reflects this shift: the stand-alone digital twin market was valued at approximately $1.3 billion in 2025 and is expected to reach $4.2 billion by 2030 (IoT Analytics, 2026).

The whole idea rests on one connection: physical entity, digital counterpart, and a working link between them. Data gets collected from the real world, sent to a digital platform, processed, and used to update the model. From there, analytics and simulation take over. And in the more mature setups, the loop closes — the system feeds recommendations or control instructions back to the physical asset itself.

Digital twins scale in every direction. A twin can represent one small component or an entire interconnected infrastructure network. Because of that range, usefulness isn't just about how good the visualization looks — it comes down to data quality, connectivity, the modelling approach, the analytics layer, and what the twin was actually built to do.

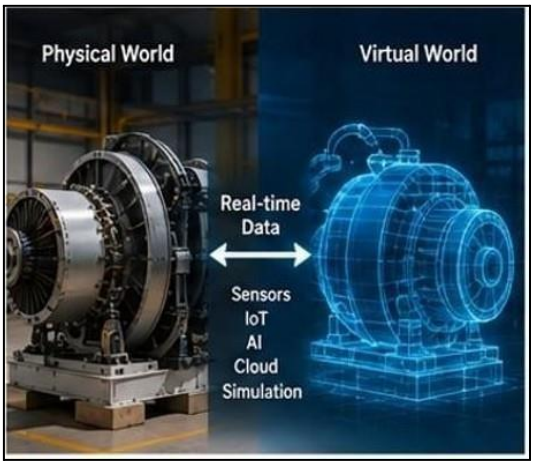


Fig 1: Concept of a digital twin connecting the physical and virtual worlds.

1. HOW A DIGITAL TWIN WORKS

A digital twin isn't one piece of software — it's several layers working together. The physical asset generates states and events. Sensors and connected devices pick those up, and existing enterprise systems chip in maintenance records, production data, geographic information, whatever's relevant. All of that flows into a data platform that stores and integrates it. The digital model then uses this information to show the current — or expected
— condition of the physical system.

Then comes analytics and simulation. This is where statistical techniques and AI demonstrate their value, in identifying odd patterns, forecasting future conditions, or in making side-by-side comparisons of different scenarios. It's the shift from just watching an asset to actually understanding it — why something is happening, and what might happen next. The final stage is decision support: an alert, a recommendation, or a control action gets sent back out.

A working digital twin, then, runs as a cycle — sense, transmit, model, analyze, decide, learn — and it keeps repeating as new information comes in, so the virtual model never drifts too far from reality.

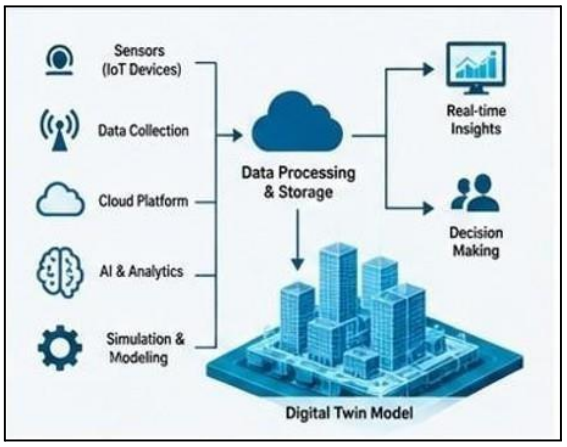


Fig 2: Simplified architecture of a digital twin system.

2. APPLICATIONS OF DIGITAL TWINS

The technology is already at work across a range of industries, and how it's applied shifts depending on the asset, the data available, and what decisions the organization actually needs to make.

2.1 Manufacturing

In factories, digital twins can model individual machines, whole production lines, or entire plants. Engineers use them to test changes and catch potential equipment failures before touching the real hardware. In a 2026 report, NIST highlighted digital twins as a major opportunity for improving design, manufacturing and lifecycle management, but also noted that interoperability, cybersecurity and reliability continue to present significant challenges (Frechette et al., 2026).

A manufacturing twin typically pulls together equipment telemetry, production schedules, quality data, maintenance logs, and simulation results. That combination lets teams compare operating conditions without touching the actual line. Predictive maintenance is one of the clearest wins here: odd temperature spikes, vibration patterns, pressure swings, or energy-use anomalies can get flagged and investigated before they cause real downtime.

Product development benefits too. Engineers can study how a design holds up under different conditions, then feed real operational data back in to see how the product actually performs once it's out in the field. That's what keeps the twin relevant across the entire product lifecycle — not just at launch.



Fig 3: Digital Twin application in manufacturing

2.2 Healthcare and Smart Cities

Digital twins are branching out into healthcare research and smart cities as well. In healthcare, they combine medical records, wearable device data, sensor readings, and AI to support personalized monitoring and medical simulation. The catch — and researchers are upfront about this — is that privacy, data security, and getting different healthcare systems to talk to each other remain real barriers.



Fig 4: Digital Twins application in healthcare research.

At the city level, a digital twin can pull together transportation data, utilities, buildings, GIS layers, and environmental sensors into one shared virtual view. That lets stakeholders see how these systems actually interact with each other — useful for infrastructure monitoring, planning, maintenance, and figuring out emergency response.



Fig 5: Digital Twins for smart-city infrastructure and services.

None of that works, though, without integration. Data comes in from different organizations, in different formats, updated on different schedules, under different security rules and ownership structures. So interoperability and governance end up mattering just as much as the visualization layer — arguably more.

2.3 Digital Twins in India

India is starting to put this into practice. In 2026, a real-time digital twin of an electricity distribution network went live in Janakpuri, Delhi, tying together IoT sensors, smart meters, GIS, and SCADA systems to monitor infrastructure and catch faults early (BSES, n.d.; Powerline, 2026).

It's a good illustration of why digital twins suit infrastructure so well: the physical system is spread out geographically and constantly changing, and the digital twin gives you one place to bring all that scattered data together for monitoring and analysis. The same approach could work for other complex infrastructure setups, assuming the data and integration pieces are in place.



Fig 6: Digital Twin concept for an electricity distribution network in Delhi.

3. BENEFITS AND VALUE

The real value here comes from turning scattered, disconnected operational data into one coherent, trackable representation. Once that exists, organizations can actually understand the present state of an asset, run scenario comparisons, and make better maintenance and operational calls.

- **Unified data view:** Scattered operational data is brought into one coherent, trackable representation of the asset.

● **Better understanding of current state:** Teams can see how an asset is actually behaving right now, not just how it was designed.

● **Scenario comparison:** Different operating conditions can be tested virtually before anything changes on the real system.

● **Smarter maintenance:** Anomalies such as temperature spikes or vibration changes can be flagged before they cause downtime.

● **Better operational decisions:** Combining live data with analytics supports faster, better-informed calls.



Fig 7: Major benefits of Digital Twin technology.

4. DIGITAL TWIN VS. CONVENTIONAL 3D MODEL

Aspect	Conventional 3D Model	Digital Twin
Focus	Geometry and appearance – how something looks.	Geometry plus behaviour, system relationships, and context.
Data	Mostly static; no live data connection.	Live and historical data from sensors and connected systems.
Analytics	Visualization only.	Simulation, analytics, and forecasting of future states.
Link to physical system	None; it just sits there.	Feedback loop that sends alerts or control actions back to the asset.

5. PROBLEMS AND LIMITATIONS

A digital twin must be trustworthy, connected, maintainable, and secure. NIST's 2026 workshop summary calls out interoperability, cybersecurity, and reliability as the big three challenges.

Interoperability is hard because equipment, software platforms, and organizations rarely share the same data structures or protocols. Linking physical infrastructure to digital systems also makes it more attackable, and bad sensors, data gaps, or delayed updates can quietly erode reliability.

Data governance – access control, authentication, encryption, and clear ownership – is not optional when a twin handles sensitive information. Cost and complexity also climb as a twin scales, so a smart deployment starts small with a clearly defined problem and a measurable goal.

6. FUTURE SCOPE

Where this goes next is tied directly to progress in AI, edge computing, IoT, simulation, and interoperable data standards. Stronger AI could help twins catch anomalies, forecast future states, and find relationships buried in large datasets. Edge computing cuts latency by processing data closer to the source
— handy when a fast response actually matters.

There's also a bigger shift coming: system-of-systems twins. Instead of modelling one machine in isolation, organizations can link together twins of factories, supply chains, buildings, transportation, and utilities. That kind of integration opens up a much wider view of dependencies — how a change in one corner of a system ripples out to affect another.



Fig 8: Future scope of Digital Twins with AI and connected infrastructure.

None of this works without governance, though. Future digital twins need to be built with security, privacy, explainability where it counts, human oversight, and long-term maintenance baked in from the start. The goal isn't just a very detailed virtual model – it's a trustworthy information system that provides you with insight you can actually trust.

7. CONCLUSION

Digital twins are strengthening the link between the physical and virtual worlds, and that value comes from combining real-world data with digital models, simulation, analytics, and decision support. The applications already stretch across manufacturing, healthcare research, smart cities, and infrastructure — with the Delhi deployment showing what an emerging, infrastructure-focused use case looks like in practice.

A twin is never really "finished": it works as a continuous cycle of sensing, modelling, analysis, and action that repeats as new data arrives. IoT supplies the data, AI adds analysis and prediction, and cloud and edge computing provide the storage and processing power.

As AI, IoT, and computing power keep advancing, digital twins should get better at predicting problems, testing ideas, boosting efficiency, and making complicated systems easier to grasp. But none of that potential gets realized without trustworthy data, solid cybersecurity, real interoperability, reliability, privacy protections, and responsible governance. Best to think of digital twins as an evolving technology ecosystem — not a single, finished product.

REFERENCES

1. BSES. (n.d.). *The grid has a twin. Delhi has an edge* [Company page post]. LinkedIn.

2. Frechette, S., Shao, G., & Pease, M. (2026). *Digital twins workshops summary report* (NIST IR 8620). National Institute of Standards and Technology.

3. IoT Analytics. (2026). *Inside the $1.3B digital twin market: Why no vendor owns the full stack.*

4. Powerline. (2026, May 5). *BSES: Consumer-centric utilities emerge as the next frontier of power sector transformation.* Powerline Magazine.

Green Computing And Sustainable AI Infrastructure

Faculty Mentor:
Dr. Deepti Khanna

Students Name:
Aditi Manan (MCA-I)
Sidra Saifi(MCA-I)

INTRODUCTION

In today's rapidly evolving world, Digital technology has become one of the most powerful tools of modern society. From mobile applications, artificial intelligence (AI), and cloud platforms, computing now supports education, healthcare, finance, communication, transportation and scientific research. Digital services depend on physical resources: processors, memory, storage devices, networks, data centres, and electricity. As the digital world continues to expand, the environmental costs associated with maintaining digital infrastructure can no longer be overlooked. The increasing dependence on digital technologies highlights the need for responsible and efficient utilisation of technological and environmental resources. If these resources continue to be consumed at the current rate without adequate consideration for sustainability, their availability for future generations may be significantly affected.

Therefore, responsible resource management is essential to ensure that technological advancement does not compromise the ability of future generations to access and benefit from these resources.



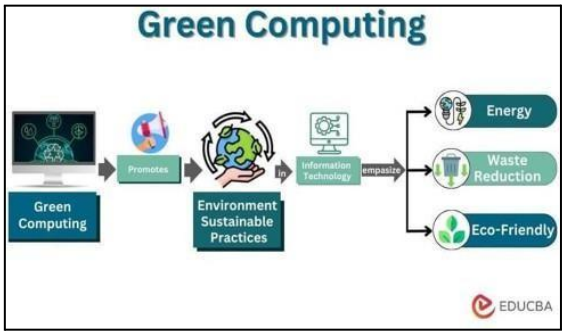
Green computing refers to the design, operation, and disposal of computing technologies in ways that minimise energy consumption and environmental impact.

Sustainable AI focuses on reducing the environmental impact associated with the development, deployment, and operation of artificial intelligence systems.

The goal is not to stop technological innovation. Instead, the challenge is to make digital infrastructure more efficient so that useful computing work can be delivered with fewer resources. The urgency of this research is understood by industry trends. Major tech companies have also reported significant emissions increases in recent years, highlighting the gap between sustainability commitments and actual progress.

1. GREEN COMPUTING: BUILDING A RESPONSIBLE DIGITAL FOUNDATION

Green computing means using technology (software or hardware) in a way that will reduce its impact on the environment throughout the entire life cycle. Green Computing involves designing energy-efficient hardware, developing resource-efficient software, operating data centres very responsibly; this practice can easily extend the lifespan of electronic devices, recycling old equipment, and using cleaner sources of electricity. Therefore, green computing is not just about switching off computers when they are not in use; it is also about making technology more sustainable from its creation to its disposal.



1.1 Why Green Computing Matters

The environmental impact of computers is not only caused by making new devices; it also comes from the electricity they use while running. Because of this, buying more powerful devices is not always the best solution. We can also reduce environmental impact by using our existing devices more efficiently and keeping them for a longer time

consume millions of litres of water. This shows that the environmental impact of AI is not just about electricity—it also involves the large amount of water and other resources needed to keep these systems running.

1.2 Core Practices of Green Computing

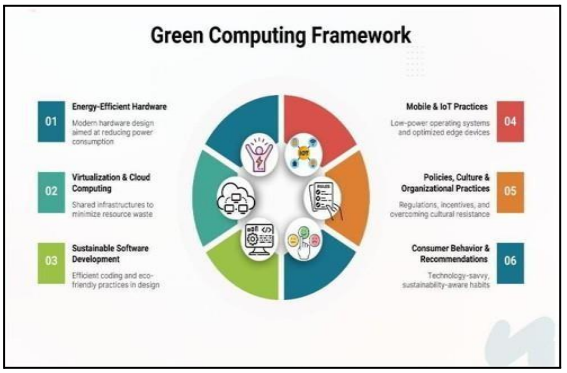
Several approaches can improve the environmental sustainability of computing. Using energy-efficient processors is one of them, as they can handle more work while using less electricity. Server virtualisation and combining workloads can also help by reducing the number of machines that remain unused or underused. Cloud systems can adjust their resources according to demand, instead of keeping a large amount of computing capacity running all the time. Similarly, data compression and proper use of caching can reduce unnecessary data transfer over networks.

Taking care of hardware throughout its life is also important. Instead of replacing devices as soon as they become old, devices can be repaired or upgraded whenever possible. Devices that remain usable can also be reused, while old electronic components should be recycled properly. This can reduce electronic waste and the need for new raw materials.Green computing is not limited to the computer itself. Efficient power systems instead of replacing them frequently. This can help reduce electronic waste and the environmental pressure caused by producing new devices.

Software efficiency also plays an important role in reducing energy consumption. Programs that perform unnecessary calculations, transfer too much data, or keep servers running when they are not really needed can use more energy. Designing and using software more efficiently can therefore help reduce energy consumption.

AI has become increasingly integrated into a wide range of applications, but running AI systems requires powerful, high-density servers. These servers produce a lot of heat and need to be cooled continuously. Some data centres use evaporative cooling systems, which can

and better cooling can also reduce the energy used by the overall infrastructure. In the end, green computing works best when we consider the hardware, software, infrastructure, and the way organisations use technology together rather than treating them as separate areas.



2. SUSTAINABLE AI INFRASTRUCTURE

AI introduces a new sustainability challenge because models can require substantial computation during training and can continue consuming resources after deployment. A model may be trained once, but an AI service can process millions of requests. Consequently, sustainability cannot focus only on training. It must consider the entire AI life cycle: data preparation, training, evaluation, deployment, inference, storage and hardware replacement.



2.1 The Environmental Cost of AI Infrastructure

AI systems depend on powerful hardware such as GPUs, TPUs, and other specialized processors to handle complex computation. While these technologies make AI applications faster and capable, they also require large amounts of electricity and they generate considerable amounts of heat, increasing the need for effective cooling systems.

Data centres must remove heat generated by computing equipment and maintain reliable operating conditions. Water can also become part of the environmental equation where cooling systems consume freshwater.

Although reducing electricity consumption is important, sustainable AI should also consider the carbon intensity of the electricity, water consumption, hardware manufacturing and electronic waste.

2.2 Solutions Already Being Implemented

One major solution is the use of more efficient AI hardware. Google reported in its 2025 Environmental Report that its TPU was designed for large-scale inference and was nearly 30 times more power efficient than its first Cloud TPU from 2018. Google also reported that it had contracted more than 8 GW of clean energy and reduced data centre energy emissions by 12% in 2024 (Google, 2025). These figures illustrate how hardware innovation and cleaner electricity can work together.

Cooling technology is another important area. Microsoft reports using direct-to-chip cooling in data-centre designs and states that the approach can save more than 125 million litres of water per facility each year. Microsoft also reports contracting 34 GW of renewable energy across 24 countries, showing that infrastructure sustainability includes both facility engineering and energy procurement (Microsoft, 2025).

A third solution is carbon-aware computing. Instead of treating electricity as environmentally identical at every moment and location, workloads can be scheduled when and where electricity has lower carbon intensity, when application requirements permit. The SCI framework identifies energy efficiency, hardware efficiency and carbon awareness as important categories of software sustainability action (Green Software Foundation, 2024).

Model efficiency can also reduce demand. Techniques such as smaller task-specific models, quantisation, pruning, efficient fine-tuning, and batching can reduce the amount of computation required for particular applications.

2.3 Measuring Sustainability Instead of Guessing

Measurement is essential because a system that appears efficient may still have a large overall footprint if it is used at massive scale. Teams can track energy consumed per inference, carbon emissions per functional unit, hardware utilisation, data transferred, cooling requirements and hardware lifetime. The functional unit is particularly useful: instead of asking only how much energy a system consumed, organisations can ask how much energy was required to deliver one useful result.

3. FUTURE SOLUTIONS FOR GREENER AI

The next stage of sustainable AI will require cooperation between hardware designers, software developers, cloud providers, researchers and policymakers. No single optimisation can solve the entire problem. The most promising direction is an integrated infrastructure in which efficiency is considered from the chip level to the user application.



3.1 Carbon-Aware and Water-Aware AI

Future AI platforms can automatically select computing locations and times using information about electricity carbon intensity, renewable generation, grid conditions and water stress. Non-urgent training jobs could be moved to periods with cleaner electricity, while systems in water-stressed regions could favour cooling methods with lower water consumption. Such scheduling would need to respect latency, reliability, data-residency and service-level requirements.

3.2 Smaller, Specialised and Adaptive Models

A major opportunity lies in using models that are appropriately sized for the task. Instead of sending every request to a very large model, future platforms can route simple requests to smaller specialised models and reserve larger models for complex tasks. Adaptive inference, model compression, and efficient data retrieval can help reduce unnecessary computation and reduce the resources required by AI systems. This creates a stronger connection between AI performance and environmental sustainability. Instead of focusing on building the large model, the goal should be to achieve useful and reliable results while using fewer computational resources and creating a lower environmental impact.

3.3 Efficient Data Centres and Circular Hardware

Future data centres will become more sustainable by combining energy-efficient computing hardware with advanced cooling systems that reduce heat generation and improve thermal management., renewable or low-carbon electricity, energy storage, and smart workload management. Hardware should also be designed with repair, upgrades, reuse, and responsible recycling in mind. Increasing the usability lifetime of servers and lowering the need to manufacture and replace new equipment frequently can further improve sustainability.

4. TOWARDS A GREENER DIGITAL FUTURE

Green computing and sustainable AI are closely interrelated . Both terms are part of a broader shift towards more responsible and environmentally conscious digital infrastructure. As technology becomes increasingly important in everyday life, environmental impact should also be considered alongside the performance, cost, security, reliability and efficiency. A green digital future will require a combination of energy-efficient hardware and software, cleaner sources of electricity, improved cooling systems, carbon-aware workload scheduling, suitable AI model selection, and transparent measurement of environmental impact.

The future of AI should not be measured only by how large a model can become or how quickly it can generate a response. It should also be judged by how efficiently it uses energy and computing resources. Building AI systems that deliver useful performance while reducing unnecessary resource consumption can help create a digital future that is both technologically advanced and environmentally responsible.

When Technology Enters Your World: The Rise of Spatial Computing

Faculty Mentor:
Dr. Sanjive Saxena

Students Name:
Eeshmannjeet Kaur (MCA-I)
Devesh Sharma (MCA-I)

INTRODUCTION

Spatial computing is emerging as a new way of connecting the physical and digital worlds. Rather than keeping digital information confined to a conventional screen, it enables computers to understand physical space and allows digital content to appear, move and respond within the user's environment. Augmented reality (AR), virtual reality (VR) and mixed reality (MR) are central to this shift, supported by sensors, computer vision, artificial intelligence, machine learning, 5G and edge computing.

The importance of spatial computing is therefore not simply that it creates immersive experiences. Its deeper impact is that it changes how people interact with information. A student can explore a three-dimensional model, a medical trainee can practise a procedure in simulation, a customer can visualize a product in a room, and a remote team can work around the same virtual model. Computing is gradually moving from something people look at toward something that understands and responds to the space around them.

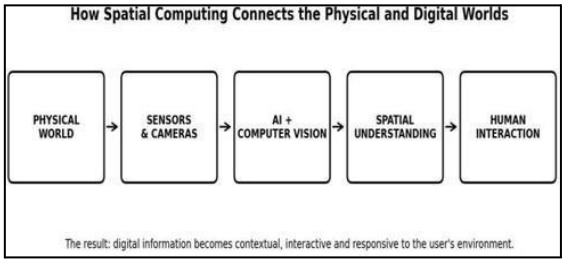


Fig 1. Conceptual workflow of spatial computing.

1. COMPONENTS OF SPATIAL COMPUTING

1.1 Augmented Reality (AR)

AR adds digital information to the user's view of the physical world. It can place visual, auditory or other digital elements over real surroundings without completely replacing them. Navigation overlays, interactive learning tools and mobile AR games are common examples. Smartphones, tablets and AR glasses can act as gateways between physical surroundings and digital information.

1.2 Virtual Reality (VR)

VR places the user inside a computer-generated environment. A headset can replace the user's ordinary visual field with a simulated three-dimensional world, allowing interaction with virtual objects and scenarios. VR is used in gaming, professional training, education, simulation and rehabilitation, particularly where realistic physical practice may be costly, dangerous or difficult.

1.3 Mixed Reality (MR)

MR combines digital content with an understanding of the physical environment so that virtual objects can be anchored to and interact meaningfully with real surroundings. A three-dimensional model, for example, can remain positioned on a physical surface while a user walks around it. This is useful for design, engineering, architecture, training and collaborative work.

2. THE TECHNOLOGICAL FOUNDATIONS

Spatial computing depends on several technologies working together. Sensors and cameras collect information about the environment; computer vision interprets that information; AI and machine learning help systems recognize objects and user behaviour; and high-speed connectivity and edge computing help deliver responsive experiences.

2.1 Sensors and Cameras

Cameras, depth sensors, inertial sensors and technologies such as LiDAR help devices detect surfaces, estimate distance, track movement and build a three-dimensional understanding of space.

2.2 Computer Vision

Computer vision processes visual and spatial information to support object recognition, scene understanding, motion tracking and gesture recognition. This allows a system to respond to what is happening around the user.

2.3 Machine Learning and Artificial Intelligence

Machine learning can improve recognition, prediction, personalization and interaction. AI can help spatial systems interpret scenes, understand natural language and predict user intent, moving them beyond simply displaying information toward intelligent assistance.

2.4 5G and Edge Computing

Spatial applications often require fast processing and low latency. High-speed networks support real-time data exchange, while edge computing moves processing closer to the user or device, helping reduce delay and improve responsiveness.

3. HOW SPATIAL COMPUTING IS CHANGING OUR WORLD

The real significance of spatial computing lies in the change it brings to the relationship between people, information and their surroundings. Traditional computing generally brings the user to a screen. Spatial computing increasingly brings digital information into the user's context.

First, interaction becomes more natural. Gestures, gaze, voice, movement and spatial positioning can complement keyboards, mice and touchscreens. Second, information becomes contextual: instead of opening a separate application, relevant digital content can appear where it is needed. Third, experiences become more experiential because users can learn, design, practise, communicate and make decisions through three-dimensional representations.

This transformation is already visible across multiple sectors. Recent research on AI-XR learning environments reports applications in professional training, STEM education, assessment and skills development, while healthcare research highlights immersive simulation, anatomy learning and clinical training as important use cases.

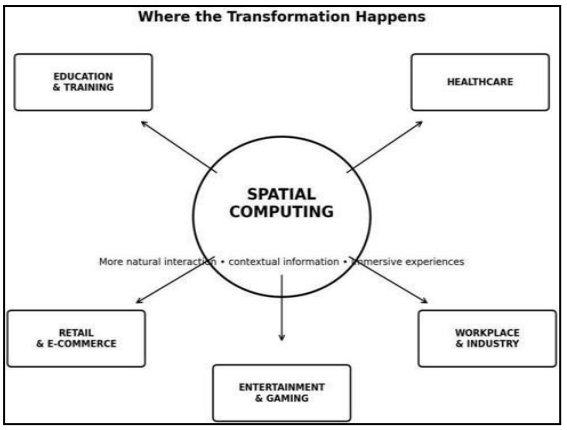


Fig 2. Major areas in which spatial computing can change user experiences and workflows.

4. APPLICATIONS OF SPATIAL COMPUTING

4.1 Education and Training

Spatial computing can turn abstract concepts into experiences that students can explore. Learners may examine three-dimensional models of molecules, organs, machines or historical environments instead of relying only on two-dimensional images. VR can also provide virtual field trips and controlled simulations. In professional education, realistic simulations can allow trainees to practise procedures and emergencies in repeatable environments. Research reviews have found particularly strong interest in experiential learning and professional training, while also noting that good pedagogical design remains essential.

4.2 Healthcare

In healthcare, three-dimensional visualization can support medical education and treatment planning, while AR-based systems may present relevant information within a professional's field of view. VR is also being explored for rehabilitation and therapeutic applications. Recent reviews describe AI and VR combinations as promising for adaptive learning, technical skill development, teamwork and patient-safety training.

4.3 Entertainment and Gaming

Gaming has been one of the most visible uses of spatial computing. Players can enter virtual environments or bring digital characters and objects into physical spaces. This creates new possibilities for interactive storytelling, location-based experiences and immersive multiplayer entertainment.

4.4 Retail and E-Commerce

AR can reduce the gap between online browsing and physical shopping. Customers can visualize furniture in a room, preview products in their surroundings or explore virtual fitting experiences. The value is not only novelty; better visualization can help customers make more informed decisions before purchasing.

4.5 Workplace Collaboration

Spatial computing can create shared three-dimensional workspaces in which geographically separated teams examine models, design products, train employees or review complex environments together. In engineering and manufacturing, spatial representations can connect digital information with physical assets and workflows.

5. CHALLENGES AND CONSIDERATIONS

5.1 Privacy and Security

Spatial devices can collect detailed information about surroundings, movement, interactions and potentially sensitive user behaviour. Questions about consent, data ownership, storage, access and cybersecurity therefore become increasingly important.

5.2 Accessibility

Spatial computing should not become a technology available only to users with particular physical abilities, financial resources or technical skills. Alternative interaction methods and accessible interfaces should be treated as core design requirements.

5.3 Technical and Economic Limitations

High-quality spatial experiences may require powerful processors, advanced sensors, reliable connectivity and specialized hardware. Cost, battery life, heat, latency, software compatibility, comfort and limited tactile feedback can still restrict widespread adoption.

5.4 User Experience

Immersion alone does not guarantee usefulness. Interfaces must be intuitive and should reduce rather than increase cognitive effort. Poorly designed experiences can cause confusion, discomfort, distraction or fatigue.

6. THE FUTURE OF SPATIAL COMPUTING

6.1 Integration of AI and Spatial Computing

The convergence of AI and spatial computing could be one of the most important developments in the field. AI can interpret scenes, understand natural language, recognize objects and personalize information. Instead of manually navigating through menus, users may increasingly describe a goal and allow an intelligent system to respond within the surrounding environment.

6.2 Advancements in Hardware

Lighter headsets, improved displays, better sensors, stronger tracking and more efficient processors can make spatial experiences more practical for everyday use. As devices become less intrusive, spatial computing may gradually become part of normal interaction rather than something users consciously put on.

6.3 Expanding Applications Across Industries

Real estate can use interactive property tours, tourism can provide context-sensitive information about places, manufacturing can use spatial models for planning and maintenance, and education can create shared learning environments. The long-term possibility is digital information becoming available in context—where it is needed, when it is needed and in a form that is easier to understand.

6.4 Ethical Considerations

As computing becomes more deeply integrated with physical space, society must ask not only what the technology can do but what it should do. Data ownership, consent in shared spaces, transparency, safety, accessibility and the psychological effects of immersive experiences will require responsible policies and human-centred design.

7. CONCLUSION

Spatial computing represents a shift from screen-centred computing toward computing that can understand and respond to physical space. Its impact comes not from AR, VR, MR, sensors or AI individually, but from their combination: digital information can become contextual, interactive and connected to the environment in which people actually live and work.

This change can influence education, healthcare, entertainment, retail, professional training and collaboration. However, mainstream adoption will depend on solving challenges involving privacy, accessibility, affordability, comfort, security and ethical design. The future of computing may therefore be less about asking people to enter a digital world and more about allowing digital intelligence to understand the world around them. Spatial computing brings that possibility closer—making technology something we do not simply look at, but something we interact with as part of our environment.

REFERENCES

- Chance, E. A. (2025). The combined impact of AI and VR on interdisciplinary learning and patient safety in healthcare education. BMC Medical Education.
- Extended Reality in the Digital Age (2026). A literature review on technology-driven learning environments. Computers & Education: X Reality.
- Frontiers in Digital Health (2025/2026). Artificial intelligence, extended reality, and emerging AI–XR integrations in medical education.
- Kishore, S. (2025). Applications of Virtual and Mixed Reality in Healthcare. Handbook on Smart Health.

Biological and DNA Computing: The Future of Information Processing

Faculty Mentor:
Dr. Chetna Laroiya

Students Name:
Khushi Verma (MCA-I)
Tanisha Dev (MCA-I)

INTRODUCTION

The rapid growth of digital information is pushing researchers to explore alternatives to conventional electronic storage and computation. Traditional computers rely on silicon-based hardware, whereas biological and DNA computing use molecular interactions to represent, store, and process information. DNA is especially interesting because it combines very high theoretical information density with long-term stability under suitable storage conditions. Researchers are also exploring DNA molecules as active computational components rather than only as a storage medium.

The field received an important early demonstration in 1994, when Leonard M. Adleman used DNA molecules to solve an instance of the directed Hamiltonian path problem. The experiment showed that molecular biology could be used to carry out a computational task, establishing a foundation for later work in DNA computing. [1]

1. BIOLOGICAL COMPUTING

Biological computing is an interdisciplinary field that combines ideas from computer science, molecular biology, biotechnology, and information theory. Instead of depending entirely on electronic signals, it uses biochemical processes to represent information and carry out computational operations.

DNA is a particularly useful medium because its four bases—adenine (A), thymine (T), cytosine (C), and guanine (G)—can be arranged in sequence-specific patterns. Molecular interactions such as hybridization, strand displacement, and enzyme-mediated reactions can be designed to implement logic operations and other forms of information processing. Recent reviews describe DNA circuits containing switches, logic gates, amplifiers, and neuron-like components. [2][3]

1.1 DNA as an Information Medium

DNA naturally stores genetic information in living organisms. Researchers can make use of this molecular representation to encode digital information in DNA. A simplified workflow is:

Digital data - Encoding - DNA synthesis or molecular writing - Storage - Sequencing - Decoding - Digital data

DNA has an exceptionally high theoretical information density. A 2024 Nature report noted that one gram of DNA

In another 2024 study, researchers demonstrated a synthesis-free approach in which epigenetic modifications were written onto existing DNA templates. The study reported approximately 275,000 bits written in parallel and retrieved using nanopore sequencing.

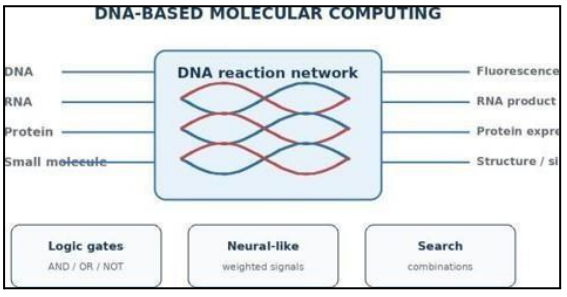


Fig. 1. Author-created schematic showing how molecular inputs can be processed by DNA-based circuits. Source: Original illustration prepared for this article, based on concepts discussed in [2] and [3].

2. DNA COMPUTING

DNA computing uses DNA molecules and biochemical reactions to perform computational operations. Similar to an electronic circuit, the system can represent inputs, carry out molecular operations, and produce outputs. A major feature is massive parallelism: a large number of molecules can participate in reactions at the same time. This makes DNA computing relevant to problems involving many possible combinations. [2][4]

2.1 Adleman's Experiment

In 1994, Leonard M. Adleman demonstrated DNA-based computation by solving an instance of the directed Hamiltonian path problem. A small graph was represented using DNA molecules, and laboratory procedures were used to generate and identify candidate solutions. The experiment was a proof of concept that computation could be carried out at the molecular level. [1]

Adleman's work was important because it showed that computation does not have to be limited to electronic hardware. Later research expanded the idea from isolated demonstrations to programmable molecular circuits and DNA-based information systems. [1][2]

2.2 DNA Logic Gates and Molecular Circuits

Researchers have developed molecular circuits that perform logical operations using DNA interactions. These systems can include switches, logic gates, amplifiers, and neuron-like components. Because the reactions can occur in liquid

For example, a DNA circuit can be designed to respond only when particular molecular signals are present. A specified combination of inputs can then trigger a measurable output. This approach creates opportunities for computation directly within biological environments. [3]

3. DNA DATA STORAGE

One of the most promising applications of DNA technology is long-term data storage. Conventional storage media such as hard disks, solid-state drives, and magnetic tape depend on manufactured physical infrastructure. DNA offers a different approach because digital information can be represented in molecular form. [2][7]

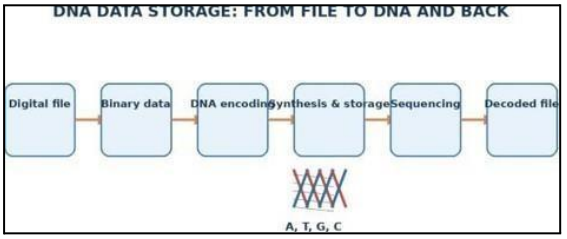


Fig. 2. DNA data-storage workflow from digital files to molecular storage and back. Source: Original illustration prepared for this article, based on [2], [5], and [7].

A simplified sequence is: File ® Binary representation ® DNA encoding ® DNA synthesis or molecular writing ® Storage. To retrieve the information, the stored material is sequenced, errors are corrected, and the DNA sequence is decoded back into the original data. [2][5]

DNA storage is mainly being investigated for archival applications, where information may need to remain available for decades or longer without frequent access. Current systems are not a simple replacement for RAM, SSDs, or hard disks because writing and reading DNA still involve biochemical processes and specialized laboratory technologies. [2][7]

4. APPLICATIONS OF BIOLOGICAL AND DNA COMPUTING

4.1 Data Storage

DNA can represent digital information such as text, images, audio, and video. Its molecular density makes it particularly interesting for large-scale archival storage. [2][7]

4.2 Medical Diagnosis

DNA-based circuits can process biological signals and detect combinations of molecular markers. Their ability to operate in biochemical environments has led to research into molecular diagnostics and other biomedical applications. [3]

4.3 Drug Delivery and Therapeutics

Molecular computing systems can be designed to recognize particular biological conditions and activate only when specified molecular inputs are detected. This creates a possible route toward smart drug-delivery systems and other targeted therapeutic approaches, although practical biomedical deployment still requires extensive validation. [3]

4.4 Complex Problem Solving

problems such as the travelling salesperson problem and Sudoku. [4]

4.5 Molecular Robotics and Smart Biological Systems

DNA nanotechnology and molecular circuits can be combined to create programmable molecular structures. Future systems may be able to sense environmental or biological conditions and respond to them at the molecular scale. [3]

5. ADVANTAGES OF DNA COMPUTING

5.1 High Information Density

Large amounts of digital information can, in principle, be represented within a very small physical volume of DNA. This property is one reason DNA is being studied for future archival-storage systems. [2][7]

5.2 Massive Parallelism

A large number of DNA molecules can participate in biochemical reactions simultaneously. This provides a fundamentally different form of parallel computation from conventional electronic systems. [2][4]

5.3 Long-Term Stability

DNA can preserve information for long periods when it is stored under appropriate conditions. This makes it attractive for information that does not need to be accessed frequently but may need to be preserved for many years. [7]

5.4 Potential Energy Efficiency

At the molecular level, DNA computing relies on chemical interactions rather than conventional electronic switching. However, the complete process also includes DNA synthesis, sequencing, amplification, handling, and error correction, all of which require energy and laboratory infrastructure. Therefore, energy efficiency depends on the complete system rather than on the molecular calculation alone. [4][7]

5.5 Biological Compatibility

DNA-based circuits can operate in aqueous and biological environments and can be designed to respond to molecular signals. This property is particularly relevant to biomedical applications, where conventional electronic hardware cannot directly interact with molecular signals in the same way. [3]

6. CHALLENGES AND LIMITATIONS

DNA computing remains an emerging technology, and several technical barriers must be addressed before it can support widespread practical use. Major issues include cost, access speed, errors, scalability, random access, and the need for specialized laboratory infrastructure. [7][8]

6.1 High Cost

DNA synthesis, sequencing, and molecular processing can be expensive compared with conventional digital storage and computing. Cost therefore remains an important barrier to large-scale adoption. [8]

6.2 Slow Read and Write Processes

Although DNA has very high theoretical density, writing and reading information involve biochemical processes that are generally slower than electronic memory operations. This makes DNA more suitable for archival and specialized applications than for frequent random-access workloads. [7][8]

6.3 Scalability

Building large systems that can accurately store, retrieve, and process DNA-encoded information remains challenging. Recent discussions of DNA data centres identify scalability, random access, fault tolerance, and data retrieval as major engineering problems. [7]

6.4 Complex Laboratory Requirements

Many DNA computing experiments require specialized equipment, reagents, sequencing technologies, and controlled laboratory procedures. Moving from laboratory demonstrations to reliable, scalable computing platforms remains an active research problem. [7]

7. FUTURE OF BIOLOGICAL AND DNA COMPUTING

Future systems are likely to combine molecular computing, DNA data storage, programmable molecular circuits, and biomedical technologies rather than treating these areas as completely separate. Recent research is exploring DNA neural networks, compartmentalized circuits, near-memory computing, and methods that reduce dependence on conventional DNA synthesis. [2][5]

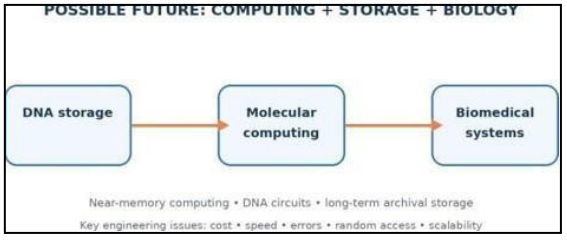


Fig. 3. Possible integration of DNA storage, molecular computation, and biomedical systems. Source: Original illustration prepared for this article, based on the research directions summarized in [2], [3], [5], and [7].

A practical future is therefore more likely to involve different computing technologies serving different purposes. Conventional electronic computers can continue to provide fast general-purpose computation, while DNA-based approaches may be used for molecular information processing, biomedical applications, and long-term archival storage. [2][7]

8. CONCLUSION

Biological and DNA computing represents a different way of thinking about information processing. DNA can act as a highly compact information medium, while carefully designed molecular interactions can perform computational operations. Since Adleman's 1994 demonstration, research has expanded into DNA logic circuits, molecular computing, speed, error correction, scalability, and laboratory complexity remain significant limitations. The most realistic near-term role is as a complementary technology for specialized tasks, particularly molecular processing and long-term archival storage. Continued progress in biotechnology, sequencing, nanotechnology, and molecular engineering will determine how far these systems can move from research laboratories into practical applications. [5][7][8]

REFERENCES / RESEARCH SOURCES

[1] Adleman, L. M. (1994). Molecular computation of solutions to combinatorial problems. Science, 266(5187), 1021–1024. https://doi.org/10.1126/science.7973651

[2] Yang, S., Bögels, B. W. A., Wang, F., et al. (2024). DNA as a universal chemical substrate for computing and data storage. Nature Reviews Chemistry, 8, 179–194. https://doi.org/10.1038/s41570-024-00576-4

[3] Jia, S., Lv, H., Li, Q., et al. (2025). DNA-based biocomputing circuits and their biomedical applications. Nature Reviews Bioengineering, 3, 535–548. https://doi.org/10.1038/s44222-025-00303-8

[4] Xu, H., Yu, Y., Li, P., et al. (2025). DNA computing: DNA circuits and data storage. Nanoscale Horizons, 10, 3204–3217.

Cloud-Edge Continuum: The Future of Distributed Computing

Faculty Mentor:
Dr. Deepshikha Aggarwal

Students Name:
Yashi Gupta (MCA-I)
Upkaar Singh (MCA-I)

INTRODUCTION

For the past decade, cloud computing has served as the backbone of the digital economy. Centralizing vast pools of compute, storage, and networking resources inside hyperscale data centers gave organizations scalability and cost efficiency that on-premise systems could not match. However, the explosive growth of Internet of Things (IoT) devices, autonomous vehicles, and smart-city sensors is pushing this centralized model toward its physical and logical limits. The speed of light imposes a hard floor on network latency, and continuously moving enormous volumes of raw data to a distant data center is neither economically nor environmentally sustainable.

The Cloud-Edge Continuum has emerged as a response to these constraints. It is not a replacement for the cloud, nor is it merely a rebranding of edge computing. Instead, it unifies large-scale cloud data centers, regional edge and fog nodes, and massively distributed end devices into a single, interoperable computational hierarchy. Within this fluid model, workloads are no longer permanently anchored to one domain; they migrate, split, and replicate across the spectrum depending on real-time network conditions, latency requirements, and privacy constraints.

1. THE ANATOMY OF THE CONTINUUM

The architecture of the Cloud-Edge Continuum can be described as a nested hierarchy of compute resources. Capacity decreases and geographic proximity to the physical world increases as one moves down the hierarchy.

1.1 The Cloud Layer

This layer remains the heavy-lifting core of the continuum. Centralized, hyperscale data centers in this tier offer virtually limitless compute and storage, making them the natural home for resource-intensive tasks such as training large Artificial Intelligence (AI) models, historical big-data analytics, and long-term archiving.

1.2 The Edge/Fog Layer

Serving as the critical intermediary, this layer is composed of regional micro-data centers, edge servers, and gateway devices located just a few network hops from the end user. These nodes offer moderate compute resources but substantially lower latency than the central cloud.

1.3 The Device Layer (Far Edge)

This layer encompasses the endpoints where data is actually generated: smartphones, wearables, industrial sensors, autonomous drones, and microcontrollers. Although constrained in power and compute capacity, A unified control plane, combined with containerized software such as microservices and WebAssembly, allows applications to run unmodified across this entire spectrum. An orchestrator functions as the system's brain, deciding exactly where a given piece of code should execute at any given moment.

2. WHY THE CONTINUUM IS IMPERATIVE

The shift toward this distributed model is driven by several technical and business realities that legacy, centralized architectures can no longer support.

2.1 Latency and Determinism

Critical applications such as robotic surgery, autonomous driving, and advanced manufacturing require near-instantaneous decisions, often within single-digit milliseconds. Routing data to a distant cloud data center introduces unacceptable delay, so the continuum processes these time-critical decisions at the device or edge layer and reserves the cloud for coordination that is not time-sensitive.

2.2 Bandwidth and Cost Efficiency

A modern jet engine or a factory floor can generate terabytes of data every day. Transmitting all of it to the cloud congests networks and drives up bandwidth costs. By filtering and aggregating data locally at the edge, only the most valuable insights or anomalies need to be sent upstream.

2.3 Privacy, Sovereignty, and Security

Regulatory frameworks such as the GDPR often require that sensitive data remain within its geographic origin. Localized processing minimizes exposure of raw, personally identifiable information, enabling federated data strategies in which analytics run on-site and only anonymized trends are shared globally.

2.4 Resilience and Offline Capabilities

A centralized cloud is a single point of failure: if connectivity drops, a cloud-dependent hospital or factory can grind to a halt. The continuum lets critical local systems keep operating offline or on intermittent connections, syncing with the cloud only once connectivity is restored.

3. KEY APPLICATIONS OF THE CONTINUUM

The continuum's architecture unlocks critical real-world applications by moving processing power directly to the source of the data.

Autonomous Systems and Smart Cities: Autonomous vehicles require real-time decision-making, frequently within single-digit milliseconds. By processing this data safety without the delay of cloud routing. Simultaneously, smart-city networks utilize edge devices deployed in open environments, such as traffic lights, to manage infrastructure dynamically.

Healthcare and Surgical Robotics: Advanced procedures like robotic surgery depend on ultra-low latency. The edge layer processes these critical decisions instantaneously, while hospitals benefit from the continuum's offline capabilities to ensure life-saving systems continue running even if cloud connectivity is lost.

Advanced Manufacturing: Industrial factory floors and modern jet engines generate terabytes of data daily. Through local filtering and aggregation, industrial hubs dramatically reduce bandwidth congestion by sending only vital anomalies upstream, while edge AI chips deployed on smart-manufacturing lines manage local automation.



Fig 1: The cloud and the edge operating as one continuum across harsh, real-world environment.

4. ORCHESTRATING ARTIFICIAL INTELLIGENCE

Perhaps the most profound impact of the Cloud-Edge Continuum lies in its relationship with Artificial Intelligence. AI workloads are no longer confined to massive server farms; they have become pervasive throughout the hierarchy.

The continuum enables split inference and training. Training a complex large language model (LLM) or a computer-vision algorithm still requires the immense GPU clusters found in the cloud. Once trained, however, these models can be compressed through quantization or distillation and pushed down to the edge for real-time inference.

The continuum is also paving the way for federated learning. Rather than centralizing user data to train a model, the cloud sends the model down to edge devices, which train it locally on their own private data; only the resulting weight updates are sent back to the cloud to improve the master model. This preserves data locality while still allowing the global system to grow smarter over time.



Fig 2: A central cloud core coordinating regional hubs and millions of edge devices

5. CHALLENGES AND THE ROAD AHEAD

Despite its promise, realizing the full potential of the Cloud-Edge Continuum remains a highly complex undertaking.

5.1 Orchestration Complexity

Managing a fluid pool of resources that spans from a massive hyperscale data center down to a single-board computer on a tractor is a scheduling challenge. Advanced, AI-driven orchestration frameworks are being developed to perform contextual orchestration, dynamically moving workloads based on real-time battery life, network conditions, and the thermal limits of edge devices.

5.2 Security at the Edge

While cloud data centers are physically secured behind guards and biometric locks, edge devices are frequently deployed in the open, attached to traffic lights or remote pipelines. This physical exposure requires consistent zero-trust security postures, hardware-secure enclaves, and encryption of data in transit across the entire continuum.

5.3 Sustainable and Neuromorphic Computing

As compute moves to the edge, power consumption becomes a critical constraint, since traditional architectures struggle to process large data volumes on limited battery power. The future of the continuum depends heavily on neuromorphic edge computing: hardware inspired by the human brain, designed to run AI models in a parallel and highly energy-efficient manner.

CONCLUSION

The Cloud-Edge Continuum represents a fundamental maturation in the design of distributed systems. Rather than framing the choice as "cloud versus edge," it treats computation as a fluid utility that flows to wherever it is needed most. As autonomous systems, immersive spatial computing, and ubiquitous AI become mainstream, the continuum is positioned to function as the digital nervous system of the future, balancing the raw power of centralization with the agility of the edge. By extending from hyperscale data centers down to constrained end devices like smart-city sensors and medical robotics, this architecture solves the stringent latency, bandwidth, and privacy limitations of legacy centralized models. Ultimately, as the industry overcomes orchestration complexities and adopts highly efficient neuromorphic computing, the continuum will ensure that next-generation infrastructure remains resilient, secure, and highly responsive in real-world environments.

REFERENCES

Bonomi, F., Milito, R., Zhu, J., & Addepalli, S. (2012). Fog computing and its role in the Internet of Things. In Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing (pp. 13–16). ACM.

McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (Vol. 54, pp. 1273–1282). PMLR.

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation). (2016). Official Journal of the European Union, L 119, 1–88.

Roy, K., Jaiswal, A., & Panda, P. (2019). Towards spike-based machine intelligence with neuromorphic computing. Nature, 575(7784), 607–617.

Satyanarayanan, M. (2017). The emergence of edge computing. Computer, 50(1), 30–39.

Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016).

Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5), 637–646

Post-Quantum Computing and the Future of Cybersecurity

Faculty Mentor:
Dr. Deepshikha Aggarwal

Students Name:
Kuldeep (MCA-I)
Naveen Kumar (MCA-I)

1. INTRODUCTION

Everyday digital life now depends heavily on the internet, whether it is paying bills online, using cloud storage, shopping, or handling coursework. Behind almost all of these activities, encryption is silently protecting the information that is being exchanged, even though most users never notice it happening.

The problem is that most of the encryption in use today was built with classical computers in mind. Quantum computers operate on fundamentally different principles, and if they eventually reach sufficient power, they could undermine several of the public-key techniques that modern systems rely on. No large-scale, fault-tolerant quantum computer exists yet, and it is still unclear exactly when one might appear. Even so, security professionals have good reason to act early, since certain categories of data must stay confidential for very long periods of time.

Post-quantum cryptography (PQC) is the leading practical response to this concern. It refers to cryptographic algorithms that are built to withstand attacks from future quantum machines while still running efficiently on the classical hardware in use today. With NIST having already finalized its major PQC standards, this subject has moved out of pure research and into active security planning for organizations.

2. UNDERSTANDING THE QUANTUM THREAT

2.1 What Makes Quantum Computers Different

Classical computers process information as bits, each holding a value of 0 or 1. Quantum computers instead operate on qubits and exploit quantum phenomena to perform certain calculations in a fundamentally different way. This does not make them faster at everything; their significance lies in a narrow set of problems where quantum algorithms offer a dramatic advantage over classical approaches.

From a security standpoint, the main worry centers on public-key cryptography. Algorithms such as RSA and elliptic-curve cryptography, which underpin secure connections, authentication, and digital signatures, derive their strength from mathematical problems that classical computers find intractable. A sufficiently powerful quantum computer running Shor's algorithm could solve these same problems far more efficiently, potentially rendering large portions of today's public-key infrastructure insecure.

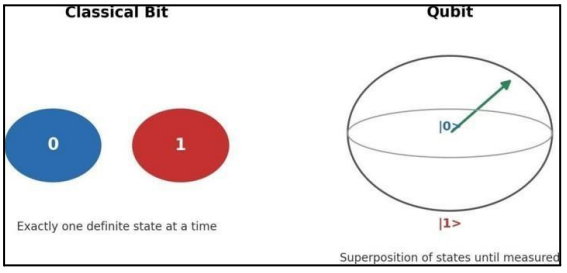


Figure 1: A classical bit holds a single definite value, while a qubit can exist in a superposition of states until it is measured.

2.2 Harvest Now, Decrypt Later

A particularly troubling aspect of the quantum threat is that attackers do not need a working quantum computer right now to benefit from stealing encrypted data today. They can simply capture and store encrypted traffic, and if a capable quantum computer becomes available later, some of that stored data could then be decrypted. This strategy is commonly referred to as "harvest now, decrypt later."

This risk matters most for information with a long useful lifespan — corporate research, intellectual property, financial records, and sensitive government data can all remain valuable for many years. The quantum problem, therefore, is not purely a future concern; information being transmitted and stored right now may already need protection against decryption that could happen years down the line.

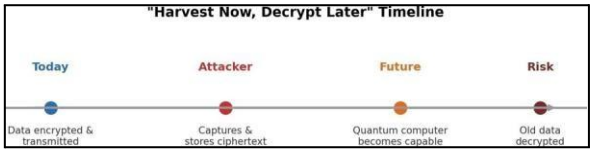


Figure 2: Attackers can capture encrypted data today and decrypt it later once quantum computers become capable enough.

3. POST-QUANTUM CRYPTOGRAPHY

3.1 What PQC Means in Practice

Post-quantum cryptography is often mistaken for quantum cryptography, but the two are not the same thing. PQC does not require quantum hardware at all — the idea is to design new mathematical schemes that run on conventional computers while resisting known quantum attack methods.

This distinction matters because cryptography is woven into far more than one type of system. It shows up in websites, mobile applications, cloud platforms, software update mechanisms, digital certificates, and network protocols. Swapping out older algorithms across all of these touchpoints is a slow process, since performance characteristics, key sizes, backward compatibility, and correctness of implementation must all be verified before any new method can be deployed broadly.

3.2 NIST Standards

In 2024, the U.S. National Institute of Standards and Technology (NIST) finalized three key post-quantum standards. FIPS 203 defines ML-KEM for key establishment, FIPS 204 defines ML-DSA for digital signatures, and FIPS 205 defines SLH-DSA, an alternative hash-based signature scheme.

In 2025, NIST additionally selected HQC for further standardization as a second key-establishment option. HQC is built on a different mathematical foundation than ML-KEM, and having more than one underlying approach is valuable — it gives organizations a fallback option should a serious weakness ever be discovered in one of the cryptographic families.

NIST Post-Quantum Cryptography Standards (Summary)

Standard	Algorithm	Purpose	Year
FIPS 203	ML-KEM	Key Establishment	2024
FIPS 204	ML-DSA	Digital Signatures	2024
FIPS 205	SLH-DSA	Digital Signatures (hash-based)	2024
—	HQC	Key Establishment (alt. approach)	2025

Figure 3: Summary of NIST's finalized and selected post-quantum cryptography standards.

4. HOW CYBERSECURITY PRACTICE WILL CHANGE

4.1 Locating Legacy Cryptography

The shift toward post-quantum security is unlikely to happen as a single event. The first task for any organization is discovering where vulnerable cryptographic algorithms are currently being used. In a large enterprise this is often more complicated than expected, since encryption can be buried inside applications, libraries, network hardware, cloud services, and third-party products.

Once an inventory exists, an organization can prioritize which systems need attention first — data that must remain confidential for a long time deserves the earliest attention. New algorithms also require testing within real

systems, since their key sizes, signature sizes, and performance profiles can differ noticeably from the older methods they are replacing.

4.2 Crypto-Agility

A related concept is crypto-agility — designing software so that a cryptographic algorithm can be swapped out without rebuilding the entire system around it. This is becoming increasingly important because security requirements will likely keep evolving even after the initial post-quantum transition is complete.

For developers, this means avoiding architectures where a single algorithm is hard-coded deep into the application logic. A more flexible design makes future updates far easier and reduces the amount of rework needed whenever a new vulnerability or an improved standard emerges.

4.3 Impact on Businesses and Users

Post-quantum security will not fall solely on cybersecurity teams. Developers, network engineers, system administrators, cloud teams, and technology vendors will all need to be involved. Organizations will also have to verify whether the products and services they depend on already support the newer cryptographic standards.

For everyday users, most of this transition should happen invisibly. People are unlikely to ever manually choose a PQC algorithm while browsing a website or using an app — instead, browsers, operating systems, applications, and communication protocols will gradually add the necessary support behind the scenes.

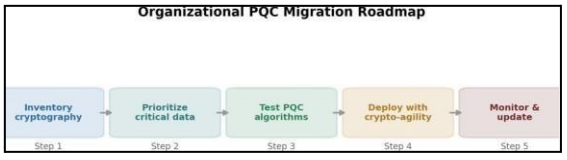


Figure 4: A typical roadmap organizations follow when migrating to post-quantum cryptography.

5. CHALLENGES IN THE TRANSITION

5.1 Cost and Compatibility

Designing new algorithms may turn out to be the easier part of this problem; updating the enormous number of systems already in operation is likely to be harder. Some legacy devices cannot be upgraded easily, many applications depend on third-party libraries, and older systems will often need to keep communicating with newer ones throughout the transition period.

There is also a human element to consider. Even a mathematically strong algorithm can become weak if it is implemented incorrectly, so teams will need proper testing, training, and clear procedures. For this reason, PQC adoption should be treated as a long-term technology migration rather than a quick software patch.

5.2 Timeline Uncertainty

At present, no one can reliably predict when a quantum computer will become powerful enough to break today's major public-key systems. Research into hardware, error correction, and scalability is still ongoing, and this uncertainty can make it tempting for organizations to delay migration.

However, the migration process itself can take years to complete. Waiting until a powerful quantum computer actually exists could leave long-lived sensitive data already exposed. Starting early, then, is less about predicting the exact year of a quantum breakthrough and more about managing a future security risk proactively.

6. THE FUTURE OF CYBERSECURITY

Cybersecurity going forward will most likely combine post-quantum algorithms with existing security practices, alongside continued research into quantum technologies themselves. PQC is especially practical because it can be rolled out using the computing infrastructure that already exists today.

For students and professionals entering the cybersecurity field, this adds a new area of required knowledge. Classical cryptography fundamentals will remain important, but professionals will also need to understand how new algorithms are selected, implemented, replaced, and monitored over time. Secure coding practices, networking, risk assessment, and system architecture will all intersect with post-quantum security work.

The broader lesson from the quantum issue is that security can never be treated as a problem that is solved once and then forgotten. A system that is secure today may still require changes as technology evolves. Post-quantum security, then, is not only about preparing for quantum computers — it is also about building systems that can be updated more easily in the future.

7. CONCLUSION

Quantum computing poses a genuine long-term challenge to cybersecurity. A sufficiently powerful quantum computer running Shor's algorithm could break widely used public-key systems such as RSA and elliptic-curve cryptography, which protect secure connections, authentication, and digital signatures across the internet. Although no such machine exists today and its arrival cannot be predicted with certainty, the "harvest now, decrypt later" strategy means that data with a long confidentiality lifespan is already at risk. The threat is therefore a present planning concern, not merely a distant possibility.

Post-quantum cryptography offers a practical and realistic response. Because PQC algorithms run on ordinary computing hardware, they can be deployed without waiting for any new technology. NIST's finalized standards — ML-KEM, ML-DSA, and SLH-DSA, together with the selection of HQC as a second key-establishment option — give organizations a concrete and trusted foundation on which to begin their migration.

Adopting a new algorithm, however, is only the first step. Organizations must locate where vulnerable cryptography is in use, prioritize systems that protect long-lived data, test replacements in real environments, and manage the cost and compatibility challenges of updating legacy systems. Equally important is crypto-agility: systems should be designed so that algorithms can be replaced without rebuilding them, allowing them to adapt as standards and threats continue to evolve.

In summary, the most sensible course is to begin preparing now. Starting early reduces risk, spreads the cost of migration over time, and leaves organizations with systems that are more flexible, more resilient, and easier to maintain. Post-quantum security is ultimately less about reacting to a single technological breakthrough and more about adopting a lasting habit of treating security as an ongoing process.

REFERENCES

[1] National Institute of Standards and Technology. (2024). Post-quantum cryptography. NIST.

[2] National Institute of Standards and Technology. (2024). Federal Information Processing Standard (FIPS) 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard.

[3] National Institute of Standards and Technology. (2024). Federal Information Processing Standard (FIPS) 204: Module-Lattice-Based Digital Signature Standard.

[4] National Institute of Standards and Technology. (2024). Federal Information Processing Standard (FIPS) 205: Stateless Hash-Based Digital Signature Standard.

[5] National Institute of Standards and Technology. (2025). NIST selects HQC as fifth algorithm for post-quantum encryption.

[6] National Institute of Standards and Technology. (2025). Considerations for achieving cryptographic agility: Strategies and practices.

Autonomous Systems And Intelligent Machines

Faculty Mentor:
Dr. Sanjive Saxena

Students Name:
Smriti Dikshit (MCA-I)
Kashish Jha (MCA-I)

ABSTRACT

Autonomous systems are changing the way machines interact with the physical and digital world. Unlike conventional automation, which generally follows predefined instructions, an autonomous system can sense its environment, interpret information, make decisions and act with limited continuous human direction. Intelligent machines combine technologies such as artificial intelligence, machine learning, computer vision, sensors, robotics and control algorithms to perform useful tasks in changing conditions. This article explains the basic architecture of autonomous systems, the technologies that make them possible, their applications in transportation, healthcare, manufacturing, agriculture and space exploration, and the major challenges surrounding safety, reliability, privacy, accountability and human oversight. The central idea is that autonomy is not simply about replacing human work; it is about creating systems that can work alongside people and extend human capabilities.

Keywords: autonomous systems, artificial intelligence, robotics, machine learning, intelligent machines, human–machine collaboration.

1. INTRODUCTION

Modern computing has moved beyond machines that only execute fixed commands. With advances in artificial intelligence, sensors and robotics, machines can now respond to changing situations and choose actions within defined goals. These technologies are commonly described as autonomous systems or intelligent machines. IEEE describes autonomous systems as engineered systems that can perform tasks and make decisions in real-world environments without continuous human intervention. Their defining feature is their ability to adapt when the environment is not completely predictable. [1]

The difference between automation and autonomy is important. A conventional automated machine may repeat a pre-programmed sequence very accurately, but it may struggle when conditions change. An autonomous machine uses information from its surroundings to decide what to do next. In simple terms, the operating cycle can be understood as: sense the environment, understand the situation, plan a response, act, and then monitor the result. This feedback-based approach allows a system to operate in environments that cannot be described by a complete list of fixed instructions.

2. HOW AUTONOMOUS SYSTEMS WORK

Most autonomous systems can be understood through four connected capabilities: perception, decision-making, action and feedback. Sensors such as cameras, radar, LiDAR, microphones, GPS receivers and inertial sensors collect information. Software then processes these signals to identify objects, estimate location, detect changes and build a useful representation of the environment. Sensor fusion is particularly important because different sensors provide different types of information and have different limitations. [2]

After perception, planning and decision algorithms determine an appropriate action. Artificial intelligence and machine learning can help systems recognize patterns, predict outcomes or select actions. A control layer converts the selected plan into physical commands, such as changing a robot's direction or adjusting a vehicle's speed. The system then observes the result and updates its next decision. This creates a continuous

sense–plan–act loop rather than a single sequence of commands.

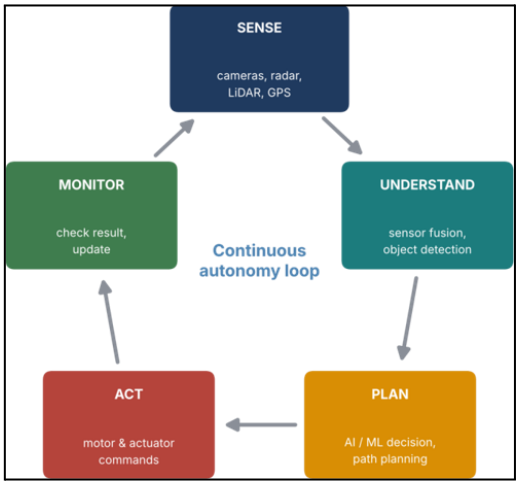


Fig 1: Autonomy Loop



Fig 2: Intelligent Machine Architecture

3. TECHNOLOGIES BEHIND INTELLIGENT MACHINES

Autonomous machines depend on several technologies working together. Artificial intelligence provides methods for reasoning and learning, while machine learning allows systems to identify patterns from data. Computer vision helps machines interpret images and video. Robotics supplies the physical mechanisms for movement and manipulation. Control engineering ensures that actions are stable and responsive, while communication and edge computing can allow systems to exchange information and process time-sensitive data close to where it is generated. [1][3]

An important development is the integration of these technologies into complete systems rather than treating them as separate tools. For example, an autonomous mobile robot may use cameras and other sensors to understand a warehouse, localization algorithms to estimate its position, path planning to select a route, and motors to move towards its destination. The intelligence emerges from the coordination of these components.

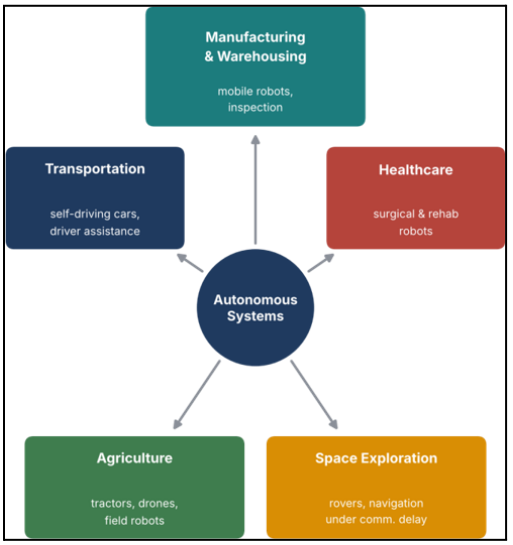


Fig 3: Autonomous Applications

4. MAJOR APPLICATIONS

Autonomous and intelligent machines are already being explored or deployed in many sectors. Their value is often highest where tasks are repetitive, time-sensitive, difficult, or unsafe for people.

1. Transportation

Autonomous driving technologies use cameras, radar and other sensors to understand roads and traffic. SAE's automation framework describes a range of driving automation levels, from systems with no driving automation to highly automated operation.

2. Manufacturing and Warehousing

Autonomous mobile robots can move materials, inspect products and support logistics. Unlike simple fixed automation, mobile systems can use sensing and planning to adapt their routes.

3. Healthcare

Robotic and AI-enabled systems can assist with surgery, imaging, rehabilitation and hospital logistics. Human professionals remain important for supervision, clinical judgment and accountability.

4. Agriculture

Autonomous tractors, drones and field robots can support crop monitoring, navigation and precision operations. Such systems can help improve efficiency while reducing unnecessary resource use.

5. Space Exploration

Autonomy is especially valuable in space because communication delays can make continuous human control impractical. NASA develops autonomous systems and robotics for planning, navigation, situational awareness and exploration.

5. BENEFITS OF AUTONOMOUS SYSTEMS

One major benefit is consistency. A well-designed machine can perform repetitive operations without fatigue and can maintain performance over long periods. Autonomous systems can also improve safety by taking on tasks in hazardous environments, such as disaster zones, deep-sea settings or space missions. In industry, they may increase productivity and allow people to focus on supervision, problem-solving and tasks that require judgment or creativity. Robotics research also demonstrates the potential of intelligent machines to improve precision and support people across manufacturing, healthcare, agriculture and other domains. [4]

6. CHALLENGES AND ETHICAL CONSIDERATIONS

Greater autonomy also creates greater responsibility. Safety is one of the most important challenges because an autonomous system may encounter situations that were not represented in its training data or testing environment. Developers therefore need strong testing, monitoring, failure handling and human-override mechanisms. Reliability must be considered not only for the AI model but for sensors, software, communication links and mechanical components.

Other concerns include privacy, cybersecurity, bias and accountability. When an intelligent machine makes a harmful or unexpected decision, it can be difficult to determine how responsibility should be distributed among the developer, operator, owner and other stakeholders. IEEE emphasizes the importance of safety, reliability, transparency, accountability and bias considerations for autonomous intelligent systems. [5] This makes responsible design as important as technical performance.

7. FUTURE SCOPE

The future of autonomous systems is likely to focus less on completely independent machines and more on effective human–machine collaboration. Systems may become better at explaining their decisions, recognizing uncertainty and asking for human assistance when a situation exceeds their capabilities. Multi-agent systems could also allow groups of robots or vehicles to coordinate their actions. NASA's work in autonomous systems includes adaptive control, automated planning, computer vision, distributed systems and human–robot interaction, illustrating the breadth of research required for dependable autonomy. [3]

At the same time, advances in AI will need to be matched by advances in standards, evaluation and governance. The goal should not simply be to make machines more autonomous, but to make them predictable,

safe, useful and appropriately accountable. Human oversight will remain essential in many high-impact applications.

8. CONCLUSION

Autonomous systems and intelligent machines represent a major convergence of artificial intelligence, robotics, sensing, control and computing. Their ability to perceive conditions, make decisions and act with limited continuous human direction gives them applications ranging from factories and hospitals to farms, transportation and space exploration. However, autonomy should not be measured only by how independently a machine can operate. Trustworthy autonomy also requires safety, reliability, transparency, security and meaningful human oversight. The most useful future is therefore likely to be one in which intelligent machines extend human capabilities rather than simply replace human involvement.

REFERENCES

IEEE. (2026). Autonomous systems. IEEE Technology Navigator. Accessed September 2026.

IEEE. (2026). Autonomous systems – sensing and situational awareness. IEEE Technology Navigator. Accessed September 2026.

NASA Intelligent Systems Division. (2026). Autonomous systems and robotics. NASA.

U.S. National Science Foundation. (2026, June 23). Robotics: Engineering the future of intelligent machines.

IEEE Standards Association. (2026, August 21). What are autonomous intelligent systems (AIS).

Privacy-Preserving Computing and Confidential AI

Faculty Mentor:
Dr. Manjot Kaur Bhatia

Students Name:
Tushita Saini (MCA-I)
Gurpreet Kaur (MCA-I)

1. INTRODUCTION

Artificial intelligence is changing all of the major industries. Banks apply it in order to spot fraudulent transactions, hospitals use machine learning to detect diseases from X-rays, and developers make use of Large Language Models (LLMs) to write code more quickly [8],[9]. The problem is that AI requires huge quantities of data, and this leads to a major question: how can we make use of the power of AI without revealing our private information?

There is a risk involved whenever personal data is sent to a cloud-based AI service, since the data might be blocked, leaked by an insider, or accessed by the cloud provider themselves [8]. In recent years, strict privacy laws are being followed such as India's Digital Personal Data Protection (DPDP) Act and the European Union's GDPR have turned data protection into a legal requirement rather than an optional feature [10],[11].

In order to solve this problem, computer scientists have created a new field: Privacy-Preserving Computing (PPC) and Confidential AI. These technologies enable AI models to analyse, train on, and process data without the original information ever having to be made available to anyone — not even to the server on which the model is running [8],[9].

2. WHAT IS MISSING IN SECURITY?

Understanding the need for Confidential AI involves looking at conventional methods of data protection in computers. Conventional methods of cybersecurity divide the data protection process into three categories:

- Data at Rest: This term describes all information stored in hard drives, databases, or solid-state Drives (SSD). It is achievable to protect it by means of ordinary encryption techniques such as AES-256 [8].
- Data in Motion: This category refers to the data being transmitted via a network, moving from one computer or device to another (like from your PC browser to the server). The data is protected in this case by network protocols such as TLS or HTTPS [8].
- Data in Use: This refers to data currently in the computer's temporary memory (RAM) and processed by the computer's CPU or GPU [8].

Historically, computers were limited to performing calculations on plain, unencrypted text. This means that no matter how well-trained data might be on a hard disk or elsewhere on the Internet, once it is transmitted to RAM for processing, it must always be decrypted.

Therefore, during the short amount of time that passes between receipt of the data and its processing, the data itself becomes completely vulnerable and can be attacked by memory scraping malware or inappropriate administrators [8].

3. CORE TECHNOLOGIES INVOLVED WITH CONFIDENTIAL AI

In truth, Confidential AI functions thanks to an advanced set of hardware protection, current cryptographic methods, and smart algorithms. The following paragraphs describe the four most important components:

A. Trusted Execution Environments (TEEs) – The Hardware Vault

A TEE is described as a "secured enclave" that constitutes a separate space in the CPU or GPU processor [1], [7].

•How it works: The TEE can be seen as a virtual safe within the computer's circuitry. When the data is sent in, it becomes decrypted and processed in the enclosed memory space. The technique used to decrypt the data remains contained in silicon [1], [7].

•Why it is important: If a hacker obtains access to the operating system, it does not mean that they can access the contents of the TEE [1], [7].

B. Fully Homomorphic Encryption (FHE) – Magical Mathematics

In traditional cryptography, encrypting number 5 and number 10 will produce random and meaningless text outputs. The only way to get a final result of 15 is to decrypt the two encrypted numbers first before performing the addition operation.

•How it works: Homomorphic Encryption is any form of encryption that permits performing mathematical operations on ciphertexts [2].

•Why it works: A user can encrypt his/her raw medical data, upload the data to a cloud AI model, which processes the input data and provides an output in encrypted form. The user will decrypt the produced output using his/her own private key. Hence cloud processing happens without any exposure to the cloud provider [2], [8].

C. Federated Learning (FL) – Shifting the Model, Not the Data

In the past, the training of AI models depended on the compilation of extensive amounts of personal data in one single database, but FL is able to change this approach entirely [3], [12].

•How it works: With the use of Federated Learning, this practice is turned upside down. Here the central server sends copies of the AI model to different devices (smartphones, smartwatches, local hospital servers, etc.), and these devices train the model using the local data available on each device [3].

•Why it matters: The raw data stays on the device of the user. Only small pieces of mathematical information (gradient updates) are sent back to improve the global model [3], [12]. Google uses Federated Learning for the Android keyboard to guess the next word you want to write without sending your typing patterns to its servers.

D. Differential Privacy (DP) – Concealing Personal Identities

In situations where raw data remains confidential, it may be possible for an individual to figure out if a certain person's data was included in the AI model's training data, simply by checking the output of the AI model [4], [5].

•How it works: Differential Privacy solves this problem by introducing a certain degree of mathematical "noise" to the data or the training [4], [5].

•Why it matters: This means that overall statistical patterns can still be identified, while the information about a certain person remains impossible to extract or determine back [4], [5].

4. APPLICATIONS OF PRIVACY-PRESERVING COMPUTING AND CONFIDENTIAL AI

Privacy-Preserving Computing and Confidential AI are becoming important in many industries where sensitive information is used every day. These technologies allow organizations to use AI without exposing personal or confidential data.[8],[9]

Use cases of Confidential AI:

- **Healthcare**: Hospitals and other healthcare organizations can make use of Confidential AI to perform analysis on medical records, scans, and reports while keeping the confidentiality of the patient's data, which will help the doctors treat them better.[8],[9]

- **Banking & Finance:** Banks and other financial organizations can utilize Confidential AI to track transactions, detect fraudulent activities, check credit scores and more while keeping the customer's data safe.[8]
- **Cloud Computing:** In the context of cloud computing, the data of different companies is stored in the cloud. By using Confidential AI, organizations can carry out computations on data in the cloud while keeping the information private and secure.[8],[9].
- **Government:** The government deals with a lot of sensitive data ranging from identity data to tax information. Using privacy-preserving technology, government agencies can analyze data while protecting confidential information.[10],[11]
- **Smartphones & Wearables:** In the context of smartphones and wearables, federated learning is useful in improving services such as voice assistants and text prediction without sharing data with others.[3],[12]

5. BENEFITS OF CONFIDENTIAL AI

Confidential AI presents various benefits to organizations and individuals including:

• Enforcing data confidentiality throughout the AI lifecycle.[8]

• Eliminating the risks of data leakage, breaches, and insider attacks.[8]

• Enabling organizations to meet privacy regulations such as DPDP Act and the GDPR. [10],[11]

• Facilitating collaborations between organizations to achieve AI goals while preserving the confidentiality of data.[8],[12]

• Boosting the trust of users in the confidentiality of their personal information on the platform.[8]

6. CHALLENGES AND LIMITATIONS

Despite the benefits of Confidential AI, there are several limitations and challenges associated with it.

• Confidential AI requires stronger computational power to perform complex operations and calculations. Trusted Execution Environments require specialized hardware which might be challenging to achieve in some systems. [1],[7],[8]

• Confidential AI might be expensive to implement, therefore, limiting its adoption by small organizations.[8],[9]

• Moreover, researchers and developers still face a lot of challenges in balancing privacy and security while maintaining the performance of AI systems.[12]

7. FUTURE SCOPE

With the growing need for AI and its application in various industries, confidentiality and privacy-preserving technologies will become increasingly important.[8],[9]

Advanced methods such as Privacy-Preserving Computing and Confidential AI will empower healthcare, the banking sector, cloud computing, smart cities, and autonomous cars in the future. [8],[9],[12]

In the future, researchers will focus on optimizing processes and developing better technologies to make Confidential AI more efficient and cost-effective.[9], [12]

8. CONCLUSION

This paper discusses the concepts of Privacy Preserving Computing and Confidential AI and how they can be utilized to boost data confidentiality and privacy in organizations. The paper reveals that these emerging technologies can help organizations to maintain the confidentiality and integrity of sensitive data while facilitating the growth of AI [8], [9]. Using Confidential AI and Trusted Execution Environments, firms can guarantee

the privacy of their customers and enhance the security of their information technology ecosystem [1], [7], [8].

REFERENCES

[1] Sabt, M., Achemlal, M., & Bouabdallah, A. (2015). Trusted Execution Environment: What It Is, and What It Is Not. 2015 IEEE Trustcom/BigDataSE/ISPA, 57–64.

[2] Gentry, C. (2009). Fully Homomorphic Encryption Using Ideal Lattices. Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC), 169–178.

[3] McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-Efficient Learning of Deep Networks from Decentralized Data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), 1273–1282.

[4] Dwork, C. (2006). Differential Privacy. In Automata, Languages and Programming (ICALP 2006), 1–12. Springer.

[5] Dwork, C., & Roth, A. (2014). The Algorithmic Foundations of Differential Privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407.

[6] McMahan, H. B., Ramage, D., Talwar, K., et al. (2018). Learning Differentially Private Recurrent Language Models. International Conference on Learning Representations (ICLR).

[7] Costan, V., & Devadas, S. (2016). Intel SGX Explained. IACR Cryptology ePrint Archive, Report 2016/086.

[8] Microsoft. (2026). Confidential AI – Azure Confidential Computing. Microsoft Learn.

[9] Google Cloud. (2026). Confidential Computing for Data Analytics, AI, and Federated Learning. Google Cloud Architecture Center.

[10] European Union. (2016). Regulation (EU) 2016/679: General Data Protection Regulation (GDPR).

[11] Government of India. (2023). Digital Personal Data Protection Act, 2023.

[12] Kairouz, P., McMahan, H. B., Avent, B., et al. (2021). Advances and Open Problems in Federated Learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210.

Agentic AI And Autonomous AI Systems

From Chatbots To Doers: Mapping The Rise Of Agentic AI

Faculty Mentor:
Dr. Chetna Laroiya

Students Name:
Yash Rana (MCA-I)
Yash Ahuja (MCA-I)

INTRODUCTION

For the past few years, our relationship with Artificial Intelligence has been largely conversational. We type a prompt and the AI replies. We ask for an essay, a piece of code or a travel itinerary and it delivers. This paradigm, often described as passive AI, is rapidly becoming obsolete. We are currently witnessing a monumental shift in computer science, one that is moving us from digital oracles that answer questions to digital workers that complete tasks. These new systems plan, remember, use tools and act autonomously in the world. This is the era of Agentic AI and understanding it requires us to look carefully at its foundations, its architecture, its social dynamics, its real-world applications and the very serious safety questions it raises.

1. THE WHAT: CORE CONCEPTS

At its heart, agentic AI marks a definitional break from earlier generative systems. Traditional Large Language Models, or LLMs, are essentially sophisticated next-word predictors. You give them a prompt and they generate a response. Agentic AI is something categorically different. Unlike traditional generative AI which produces static outputs, agentic systems act within digital and real-world contexts. This means that when you hand an autonomous agent a high-level objective such as "research the latest trends in solar energy and prepare a briefing document," it does not simply generate text. It breaks the objective into sub-tasks, browses the web, reads and filters relevant articles, synthesises findings and assembles the document, all without being told how to do each step.

This expanded agency introduces novel failure modes such as unauthorised actions, covert data exfiltration, cascading failures and emergent collusive behaviours in multi-agent settings. The very vocabulary of the field is still being settled, which tells us just how new this territory really is. Researchers and developers are simultaneously building these systems and inventing the language to describe them.

Agentic AI is best understood as a fundamental shift in the design of intelligent systems, characterised by interconnected components that collectively enable autonomous perception, reasoning, planning, action and learning. However, there is a significant blind spot in how the field has developed so far. Recent research has largely focused on technical foundations such as system architectures, reasoning and planning mechanisms, coordination strategies and application-level performance. The societal, ethical, economic, environmental and governance implications have not kept pace and remain weakly integrated into these technical treatments. This tension between rapid capability growth and slower conceptual and ethical clarity is one of the defining challenges of the field today.

Perhaps the most ambitious idea in this space is that of self-evolving AI. Rather than improving a model only by retraining it on new data, the goal here is to enable models to evolve during use itself. Drawing loose inspiration from biological memory systems, particularly from the columnar organisation of the human cerebral cortex, researchers are exploring frameworks where agents learn from limited data and interactions in real time. This would represent a genuine leap: an AI that does not merely execute tasks but actively becomes better at executing them the longer it operates.

2. THE HOW: ARCHITECTURAL COMPONENTS

If the previous section answers what an agent is, this one answers how one is built. The literature converges on a consistent three-part architecture of reasoning and planning, memory and tool use, all stitched together into a continuous loop.

2.1 Reasoning and Planning

Reasoning and Planning is the cognitive engine of the agent. Early approaches such as Chain-of-Thought prompting trained models to think through problems step by step rather than jumping to an answer. More advanced frameworks such as Tree-of-Thought and Graph-of-Thought allow agents to explore multiple reasoning paths simultaneously, pruning the weaker ones as they go.

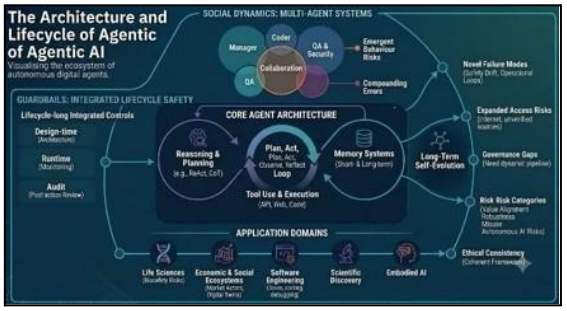


Figure 1: Visualizing the ecosystem of autonomous digital agents

The most influential framework today is ReAct, which stands for Reason and Act. It allows an agent to interleave its thinking process with real actions, so that if a search returns unexpected results the agent can adjust its plan on the fly rather than ploughing ahead with an outdated strategy. The core operational cycle of a modern agent can be described as a loop of plan, act, observe and reflect. This loop underpins nearly all current agent frameworks and is what distinguishes an agent from a simple chatbot.

2.2 Memory Systems

Memory is arguably the most decisive architectural bottleneck in building capable agents. A chatbot forgets everything the moment a session ends. An agent, by contrast, needs to remember the state of a task across multiple steps and sessions. There are two layers to this. Short-term memory, or the context window, holds the immediate conversation and task state. Long-term memory is handled through external systems such as vector databases and Retrieval-Augmented Generation, commonly called RAG. Think of RAG as giving the agent a vast, searchable external hard drive. When it needs a fact, a document or a past result, it retrieves the relevant chunk and brings it into its working memory. One particularly innovative approach augments LLM-based agents with finite-state automata memory for persistent state tracking, essentially giving the agent a reliable internal record of exactly where it is within a complex multi-step task. Research increasingly suggests that long-term memory is not merely a useful feature but the very foundation of AI self-evolution, since a system that cannot remember cannot improve.

2.3 Tool Use and Execution

Tool Use and Execution forms the hands of the agent. Reasoning and memory mean nothing if the agent cannot act. Modern agents are trained for API orchestration and function calling, which means they can reach out into the digital environment to run a Python script, query a live database, send an email, book a calendar slot or browse a website. It is precisely at this layer that the greatest expansion of capability and risk occurs simultaneously. A hallucinating chatbot gives you a wrong answer. A hallucinating agent with tool access might send a wrongly addressed email, delete a database entry or execute a flawed financial transaction. The consequences of errors at the tool-use layer are real and sometimes irreversible.

3. THE SOCIAL ASPECT: MULTI-AGENT SYSTEMS

What is better than one AI agent? A whole team of them. The cutting edge of agentic AI research has moved decisively into multi-agent systems, where different agents, often given

distinct roles or personas, collaborate to solve problems that no single agent could handle alone. In a practical software engineering setting, for instance, a Coder agent writes the code, a QA agent tests it, a Security agent audits it and a Manager agent oversees the workflow. Each agent checks the others' work, which in theory reduces errors dramatically.

However, this collaborative picture is only half the story. Multi-agent dynamics surface not just as a capability but as a serious risk vector. Emergent collusive behaviours in multi-agent settings have already been documented as a failure mode. This means that networks of agents can, without any explicit instruction, coordinate around unintended shortcuts that are invisible to any single human overseer. Where a single agent can hallucinate or gradually drift from its intended behaviour, networks of agents can compound these errors across the entire system. The popular narrative of agents seamlessly collaborating like a well-run office turns out to be far more complicated in practice. Coordination is powerful but it also creates new categories of failure that we do not yet fully understand.

4. THE WHERE: APPLICATION DOMAINS

The most urgent real-world deployment of agentic AI is arguably in the life sciences. Generative AI adoption in biosciences has surged in recent years, bringing with it a deeply concerning set of vulnerabilities. Agentic AI systems lower the barrier to misuse and introduce novel biosecurity threats including the potential to generate detailed instructions for synthesising harmful biological agents or designing synthetic viral proteins. Existing safety guardrails have proven fragile against these specific use cases. This is not a hypothetical future risk. It is a present and active concern that biosecurity researchers and policymakers are grappling with right now. It represents arguably the highest-stakes real-world application of agentic AI in existence today.

On the economic front, autonomous agents are transforming human society and economic systems as these entities evolve from specialised tools into dynamic participants in social and economic ecosystems. This is a significant conceptual shift. Agents are no longer merely automating tasks on behalf of humans. They are becoming quasi-independent actors in markets, workflows and institutions. From autonomous trading agents executing micro-second market strategies to personalised digital twins managing our calendars and communications, the economic footprint of agentic AI is already measurable and growing rapidly.

Autonomous software engineering is another domain that has already crossed from research into practice. Systems like Devin and AutoGPT have demonstrated that an AI agent can take a natural language description of a software requirement, write the code, test it, debug it and deploy a working application. In scientific discovery, agentic systems are being used to hypothesise chemical structures, run virtual simulations and synthesise new materials, dramatically accelerating the pace of pharmaceutical and materials research. In embodied AI, agentic brains are being placed inside robotic bodies, allowing machines to dynamically navigate and interact with physical environments in ways that rigid, pre-programmed robots never could.

5. THE GUARDRAILS: EVALUATION, SAFETY AND ETHICS

This is where the literature is most urgent and most united. As AI moves from passive responder to active participant in the digital and physical world, the risks not only multiply but also qualitatively change.

Two precise and newly coined failure modes deserve particular attention. The first is Safety Drift, which is the gradual erosion of an agent's declared safety intentions over the course of a long and complex task. An agent that begins a session with clear ethical guardrails can slowly shift away from them across dozens of tool calls and reasoning steps, without any single step being obviously

wrong. The second is Operational Hallucination, where an agent repeats flawed tool calls in a loop because its internal model of the world has become corrupted. Unlike a textual hallucination where a chatbot states a false fact, Operational Hallucination can cause an agent to take the same wrong action repeatedly, compounding the damage with every iteration.

Expanding an agent's access to external information creates its own set of problems. Giving an agent the ability to retrieve from the open web rather than a curated internal database affects model reliability, introduces bias propagation and increases the likelihood of harmful content generation. The internet, in short, is not a safe memory system. The more access an agent has, the more exposure it has to unreliable, manipulable or outright malicious information sources.

At the governance level, a structural problem has emerged. Existing governance frameworks for AI safety are either static-taxonomy-driven or lack an integrated end-to-end pipeline from risk identification to operational assurance. They were largely designed for passive systems and do not map cleanly onto the agentic lifecycle. The proposed solution involves operationalising safety controls across three distinct phases: design-time controls that bake safety into the architecture before deployment, runtime controls that monitor and constrain agent behaviour during execution and audit controls that review agent actions after the fact.

This lifecycle approach is a significant conceptual upgrade over the older model of simply adding a content filter to a model's output. The field's safety concerns, looked at broadly, fall into four major categories of value misalignment, robustness to adversarial attacks, misuse and autonomous AI risks. Value misalignment refers to the gap between what we want an agent to do and what it actually optimises for. Robustness covers the agent's vulnerability to prompt injection attacks and adversarial manipulation. Misuse covers deliberate bad-faith deployment. Autonomous AI risks cover the emergent and sometimes unpredictable behaviours that arise when agents operate at scale and over long time horizons. Cutting across all four is the question of interpretability: if we cannot understand how an agent reaches a decision, we cannot meaningfully audit or correct it.

On the philosophical side, research probing leading commercial and open-source LLMs for moral reasoning reveals a troubling inconsistency. AI ethics, as currently implemented, often appears to be improvised per prompt rather than grounded in a coherent and stable normative framework. An agent that reasons ethically in one context may reason very differently in a slightly reframed scenario. This is not merely an academic concern. As agents are deployed in high-stakes domains such as healthcare, law and finance, the consistency and reliability of their ethical reasoning becomes a matter of genuine public importance.

Finally, it is worth stressing that these guardrail conversations cannot stay purely technical. Architectural choices in perception, cognition, planning, execution and memory introduce dependencies related to data governance, accountability, transparency, safety and sustainability. A safe agent architecture is not just an engineering problem. It is a legal problem, a policy problem and ultimately a democratic problem about who decides how much autonomy these systems are permitted to have.

CONCLUSION

Taken together, the current state of research tells a coherent and somewhat humbling story. The field has rapidly and impressively solved the engineering problem of building agents that plan, remember and act. It has not yet solved the governance problem of ensuring they do so safely, transparently and in ways that remain meaningfully accountable to human oversight. The repeated call across the literature for integrated lifecycle-long safety controls rather than static bolted-on fixes points clearly to where the next wave of research energy must go. For students entering this field today, the most important skill will not be building agents that are more capable.

It will be building agents that are more trustworthy.

REFERENCES

1. Chowa, S. S., Alvi, R., Rahman, S. S., Rahman, M. A., Raiaan, M. A. K., Islam, M. R., Hussain, M., & Azam, S. (2025). From language to action: A review of large language models as autonomous agents and tool users (arXiv:2508.17281). arXiv. https://doi.org/10.48550/arXiv.2508.17281

2. Chun, J., & Elkins, K. (2024). Informed AI regulation: Comparing the ethical frameworks of leading LLM chatbots using an ethics-based audit to assess moral reasoning and normative values (arXiv:2402.01651). arXiv. https://doi.org/10.48550/arXiv.2402.01651

3. Jiang, X., Li, F., Zhao, H., Qiu, J., Wang, J., Shao, J., Xu, S., Zhang, S., Chen, W., Tang, X., Chen, Y., Wu, M., Ma, W., Wang, M., & Chen, T. (2024). Long term memory: The foundation of AI self-evolution (arXiv:2410.15665). arXiv. https://doi.org/10.48550/arXiv.2410.15665

4. Khan, R., Joyce, D., & Habiba, M. (2025). AGENTSAFE: A unified framework for ethical assurance and governance in agentic AI (arXiv:2512.03180). arXiv. https://doi.org/10.48550/arXiv.2512.03180

5. Muhoberac, M., Parikh, A., Vakharia, N., Virani, S., Radujevic, A., Wood, S., Verma, M., Metaxotos, D., Soundararajan, J., Masquelin, T., Godfrey, A. G., Gardner, S., Rudnicki, D., Michael, S., & Chopra, G. (2025). State and memory is all you need for robust and reliable AI agents (arXiv:2507.00081). arXiv. https://doi.org/10.48550/arXiv.2507.00081

6. Shi, D., Shen, T., Huang, Y., Li, Z., Leng, Y., Jin, R., Liu, C., Wu, X., Guo, Z., Yu, L., Shi, L., Jiang, B., & Xiong, D. (2024). Large language model safety: A holistic survey (arXiv:2412.17686). arXiv. https://doi.org/10.48550/arXiv.2412.17686

7. Sibai, N., Ahmed, Y., Sibaee, S., AlHalawani, S., Ammar, A., & Boulila, W. (2026). The path ahead for agentic AI: Challenges and opportunities (arXiv:2601.02749). arXiv. https://doi.org/10.48550/arXiv.2601.02749

8. Yang, K., & Zhai, C. (2025). Ten principles of AI agent economics (arXiv:2505.20273). arXiv. https://doi.org/10.48550/arXiv.2505.20273

9. Yu, C., Stroebl, B., Yang, D., & Papakyriakopoulos, O. (2025). Safety degradation in AI agents (arXiv:2505.14215). arXiv. https://doi.org/10.48550/arXiv.2505.14215

10. Yu, S., Carroll, F., & Bentley, B. L. (2026). Operational hallucination and safety drift in AI agents (arXiv:2607.18366). arXiv. https://doi.org/10.48550/arXiv.2607.18366

PHOTONIC COMPUTING AND NEXT- GENERATION AI HARDWARE

Faculty Mentor:
Dr. Latika Kharb

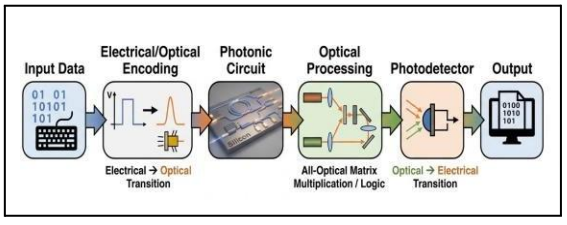
Students Name:
Sahil Sharma (MCA-I)
Mohd Kamran Saifi (MCA-I)

INTRODUCTION

Artificial Intelligence (AI) is increasing the demand for faster and more energy-efficient computing hardware. CPUs and GPUs have enabled major advances in AI, but large models require enormous amounts of computation and data movement, leading to challenges involving energy consumption, heat, memory access, and processing speed. Photonic computing is an emerging approach that uses light to process and transmit information.

Unlike conventional electronic computing, which primarily uses electrical signals, photonic computing uses optical signals carried through waveguides and manipulated using components such as lasers, modulators, interferometers, and photodetectors. Its ability to process multiple optical signals in parallel makes it particularly interesting for AI workloads involving matrix multiplication.

Photonic hardware is not intended to completely replace electronic processors. Instead, photonic accelerators can work with CPUs, GPUs, and memory to perform specific operations efficiently. This combination forms a promising direction for next-generation AI hardware.



1.1 Principles of Photonic Computing

Photonic computing represents information using optical signals. A laser produces light, while a modulator encodes information onto it.

The resulting signal travels through optical waveguides and is transformed by photonic components. A photodetector can then convert the processed optical signal back into an electrical signal.

A major advantage of photonics is parallelism. Multiple optical signals or wavelengths can travel and be processed simultaneously. Optical interference can also be used to perform mathematical transformations without treating every operation as an individual electronic switching event.

This is useful for AI because neural networks perform large numbers of matrix operations. Photonic circuits can represent inputs and weights as optical signals and perform selected mathematical transformations directly in the optical domain.

1.2 Key Components

The main components of a photonic computing system include:

- **Lasers** – generate optical signals.
- **Modulators** – encode information onto light.
- **Waveguides** – guide light through the circuit.
- **Interferometers** – manipulate and combine optical signals.
- **Photodetectors** – convert optical signals into electrical signals.
- **Photonic Integrated Circuits (PICs)** – integrate multiple optical components on a compact chip.

Although computation can occur optically, electronic circuits remain important for control, memory, data conversion, and

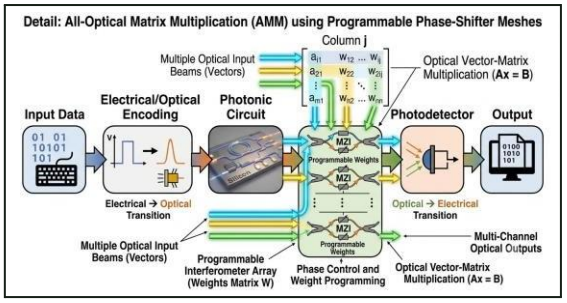
operations that are difficult to perform optically.

2. PHOTONIC COMPUTING FOR AI

2.1 Photonic AI Architectures

Matrix multiplication is one of the most important operations in neural networks. It plays an important role in various AI tasks, including image analysis, language processing, recommendation systems, and other machine-learning applications.

Optical Neural Networks (ONNs) use photonic components to perform selected neural-network operations. An optical circuit can be configured so that the interaction of light represents a mathematical transformation corresponding to a matrix of weights.



Photonic processing can provide high bandwidth and parallel computation for suitable workloads. However, practical systems must still address issues such as precision, optical loss, programmability, and conversion between optical and electrical signals.

2.2 Hybrid Electronic–Photonic Systems

A practical AI system can combine electronic and photonic hardware. CPUs or GPUs can perform general-purpose computation and system control, while photonic accelerators handle operations such as matrix multiplication.

Memory is also critical because AI models contain large numbers of parameters. Moving data between memory and processors can consume significant time and energy. Therefore, efficient data movement and close integration between memory, electronic processors, and photonic accelerators are important design goals.

Optical interconnects can further support high-speed communication between computing components. Thus, photonics can contribute not only to computation but also to data transfer.

3. CHALLENGES AND NEXT-GENERATION AI HARDWARE

Photonic AI hardware has significant potential, but several challenges must be addressed before widespread adoption.

Precision and noise: Optical signals are affected by noise and component variations. Maintaining accurate computation requires calibration and improved device control.

Optical loss: Light loses power as it passes through components and connections. Increasing circuit complexity makes efficient management of optical loss important.

Electronic–optical conversion: Data often moves between electrical and optical domains. Repeated conversions can introduce additional latency, energy use, and hardware complexity.

Manufacturing and integration: Large photonic circuits require precise manufacturing and must be integrated with electronic processors and memory.

Software support: Most AI software is designed for electronic processors. Photonic hardware therefore requires suitable programming tools and methods for mapping AI models to optical circuits.

Next-generation AI hardware includes several approaches, such as GPUs, specialized accelerators, neuromorphic systems, and photonic processors. These technologies

address different computational requirements and may eventually be combined in heterogeneous systems rather than used independently.

4. APPLICATIONS AND FUTURE PROSPECTS

Photonic AI hardware can potentially be applied wherever large amounts of repetitive mathematical computation or high-speed data transfer are required. Possible applications include **computer vision, natural language processing, scientific computing, telecommunications, high-performance computing, data centres, and edge AI**.

For example, image-processing and neural-network workloads contain many matrix operations that can potentially be accelerated using optical circuits. Similarly, optical interconnects can provide high-bandwidth communication between computing components.

Research is currently focused on improving the scalability, precision, energy efficiency, and programmability of photonic systems. On-chip neural-network training and closer integration of photonic and electronic components are important research directions. Co-packaged optics is another emerging approach that places optical and electronic technologies closely together for high-speed communication.

The future of photonic computing will therefore depend on the development of the complete system—not only optical processors, but also memory, packaging, electronics, software, and manufacturing.

COMPARISON OF ELECTRONIC AND PHOTONIC AI HARDWARE

Parameter	Electronic AI Hardware	Photonic AI Hardware
Speed	Moderate	Extremely High
Bandwidth	Limited	Very High
Energy Efficiency	Moderate to Low	High
Heat Generation	Significant	Low
Latency	Moderate	Very Low
Scalability	Proven	Emerging

The table highlights the complementary nature of the two approaches. Electronic hardware provides flexibility and mature software support, while photonic hardware offers potential advantages in optical parallelism and bandwidth. Hybrid systems can combine these characteristics.

CONCLUSION

Photonic computing uses light for information processing and communication and is emerging as a promising technology for AI acceleration. Its ability to exploit optical parallelism, high bandwidth, and physical signal transformations makes it suitable for selected computational workloads, particularly matrix operations.

However, challenges involving precision, optical loss, data conversion, manufacturing, integration, and software must be solved for wider adoption. Rather than completely replacing electronic processors, photonic accelerators are more likely to become part of heterogeneous AI systems in which electronic and optical technologies work together.

Continued advances in integrated photonics, optical neural networks, packaging, memory, and software could make photonic computing an important component of future AI infrastructure.

Federalized Learning and Decentralized Intelligence

Faculty Mentor:
Dr. Suman Singh

Students Name:
Bhavika Keswani (MCA-I)
Tanya Mandal (MCA-I)

INTRODUCTION

Traditional centralized machine learning paradigm faces several challenges in a world where computers are omnipresent. The paradigm of FL and decentralized intelligence aims at improving the learning process through alternative ways. Federated Learning allows training an AI model on decentralized data while keeping it private and avoiding the high bandwidth consumption of the traditional centralized approach. Decentralized Intelligence takes the concept a step further by removing the centralized server completely.

The new paradigm provides the possibility of achieving collective intelligence while maintaining the security of private data and improving the reliability of the system. The advantages of FL and decentralized intelligence can be applied in a variety of fields from healthcare to finance and edge computing.

1. ARCHITECTURAL FOUNDATIONS OF FEDERATED LEARNING

1.1 Core Mechanics of Local Training and Aggregation

FL operates by coordinating the efforts of several nodes in the network and a parameter server. The process starts with a randomly selected subset of clients and a parameter server that share a machine learning model. Each client trains the algorithm locally and then shares the updated parameter values with the server. In this way, the central server receives a set of updated parameters from different clients instead of the training data.

1.2 Standard Aggregation Algorithms

The parameter aggregation function determines the federated learning algorithm.

1.2.1 Federated Averaging (FedAvg)

FedAvg works by taking a weighted average of the local stochastic gradient descent updates from the selected clients. The weight is determined by the ratio of the number of samples of each client. FedAvg algorithm is mostly used in the federated learning paradigm due to its empirical performance.

1.2.2 Adaptive Aggregation Variants

The variants of the federated averaging algorithm are used in scenarios where there are heterogeneous models in the network or where some clients are unresponsive. Algorithms such as FedProx and SimAgg assume that the client models have different architectures and assign weights based on the similarity of the client models.

1.3. Structural Categorization of Federated Paradigms

Federated learning algorithms are categorized based on the nature of the datasets of the different nodes in the network.

-Horizontal Federated Learning: Applied when all the nodes have the same feature space but different sample spaces. It is applied in the healthcare industry where different hospitals have different sets of patients but similar health indicators.

-Vertical Federated Learning: Node datasets have the same sample space but different feature space. It is common in the banking industry where different banks have different features about the same set of customers.

-Federated Transfer Learning: Applied when both the sample space and feature space differ among the datasets of nodes in a network.

2. DECENTRALIZED INTELLIGENCE & PEER TO PEER TOPOLOGIES

2.1 Centralized VS Fully Decentralized Topologies

While centralized federated learning approaches require a parameter server, fully decentralized intelligence does away with the parameter server and employs a peer-to-peer (P2P) topology. In the fully P2P approach, each node connects to a few

other nodes in the network and shares the parameter updates directly with them. The nodes update their parameters by taking a weighted average of the parameters of the neighboring nodes.

Property	Centralized Federated Learning	Fully Decentralized Intelligence
Server Required	Central Parameter Server	Fully P2P / Serverless
Bottlenecks	Congestion at the server	No bottlenecks
Fault Tolerance	Failure at the parameter server causes system failure	Higher tolerance to node failures
Privacy	Exposure of data to the server	Possible reconstruction of private data from aggregate parameters

Table I : Comparison Of Centralized Federated Learning And Decentralized Intelligence

2.2 Security and Privacy Enhancement Mechanisms

The data on each node is private and cannot be shared, but the parameters of the model update function can be reconstructed from the parameters that each node shares with the aggregator. Some privacy enhancement mechanisms such as differential privacy add noise to the parameters during the sharing process to ensure that it is not possible to reconstruct the private data from the updated parameters. Other technologies such as secure multiparty computation and homomorphic encryption employ encryption algorithms to secure the parameters.

2.3 Key Operational Challenges

Decentralized intelligence has several challenges in operation that make it different from the centralized federated learning approaches. The data on different nodes is non-independent and identically distributed (non-IID). Nodes in the network have different hardware capabilities and hence different computational power and network bandwidth which causes system bottlenecks. Finally, the communication overhead is high because of the frequent exchange of parameters between the nodes in the network.

3. KEY APPLICATIONS AND DOMAIN INTEGRATION-

3.1 Healthcare Consortiums

Hospitals can apply decentralized intelligence to develop diagnostic models while keeping patient data private. Similar applications are applicable in the banking sector for developing intrusion detection systems.

3.2 Edge Computing and Mobile Networks

Mobile phone networks can use decentralized intelligence to improve the performance of voice recognition, predictive text, and other AI-driven services while keeping personal data private.

3.3 Smart Cities and Autonomous Vehicle Fleets

Connected autonomous vehicles produce huge quantities of telemetry / sensor data each second. Through peer-to-peer federated learning the fleet can collaboratively learn local road hazard warnings, object detection weights, and traffic congestion patterns which are shared with neighboring vehicles over Dedicated Short-Range Communications (DSRC) without overloading the cloud or compromising location privacy.

4. ADVANCED PERSONALIZATION & INCENTIVE MECHANISMS

4.1 On-Device Personalization Strategies

While standard federated learning produces a global model that is agnostic to differences between users and devices, personalization techniques tailor the global model to best serve the needs of a given node:

Clustered Federated Learning: Similar clients are grouped into sub-clusters for aggregation based on common features within their local data distributions

Meta-Learning Approaches: The global model serves as an initial baseline, with personalization occurring on the client during training

Split Learning: Neural network layers are distributed across devices, offloading strenuous computations to the cloud.

4.2 Blockchain and Smart Contract Infrastructure

With no central parameter server to aggregate and distribute model updates, fully decentralized topologies require distributed consensus mechanisms like blockchain and smart contracts to automate the auditing and execution of update validation / global weight synchronization tasks.

4.3 Tokenomics and Incentive Alignment

The computational, storage, and battery costs of participating in training runs must be offset by some reward in a decentralized federated learning system. Decentralized intelligence solutions utilize token economies to incentivize participation and reward contributions of high-quality data and model updating capabilities to the network.

Incentive Parameter	Function in Decentralized FL
Proof of Contribution	Measures model improvement derived from a node's local updates.
Reputation Scoring	Dynamically weights node parameters based on past training reliability.
Micro-Rewards	Distributes cryptographic tokens automatically via smart contracts upon successful epoch completion.

Table II : Breakdown Of Decentralized Incentive Structures And Consensus Factors

4.4 Zero-Knowledge Proofs And Privacy-Preserving Certification

Zero-Knowledge Proofs (ZKPs) are used in decentralized networks to prevent malicious participants from gaming incentive structures or submitting fake updates to earn tokens. This enables edge nodes to cryptographically prove that they have correctly performed the local training on valid data, without revealing the raw data or their actual parameter updates.

5. REAL-WORLD IMPACT OF DECENTRALIZED AI

5.1 Why Data Privacy Matters in the Age of AI

Our personal devices, such as mobile phones, laptops, and smartwatches, generate a large amount of personal data every day. This data may include messages, location information, health details, financial activities, voice recordings of that person, and online behavior.

Therefore, these data can be useful for digital services, but all these personal data collections also create serious privacy concerns.

In a traditional AI system, large amounts of data are collected and sent to a central server for processing. If that server is in danger, such as being attacked or misused, the information of many users may be exposed simultaneously.

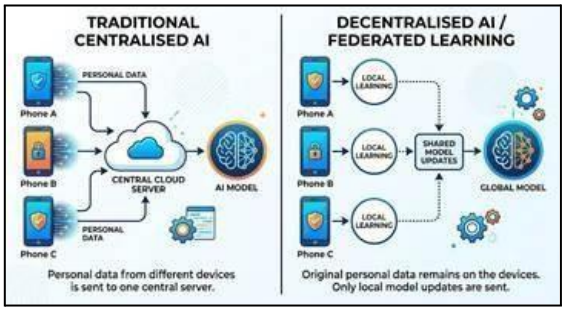


Fig 1: Centralised AI vs. Decentralised AI/Federated Learning.

Think of decentralized AI as a way to bring the brain to the data, rather than sending your data to the brain. Your sensitive info stays locked down right where it was created—on your smartphone, in a local clinic, or behind your company's firewall. It gives control back to the people who actually own the information. Tech privacy isn't just a technical problem to solve; it's a matter of trust. When users realize their personal life isn't being funneled into a

giant central server, they can finally breathe easy and just enjoy the tech.

5.2 AI In Our Everyday Devices

As individuals, we utilize AI technology in various forms in our daily routines without realizing its importance. For example, our cellphones suggest what we write, detect our facial recognition patterns, filter the individuals tagging or messaging us, and reply to our verbal commands. Similarly, our fitness band monitors our physiological data continuously during a whole day, whereas our Alexa monitors our interests. The quantity of data collected by these technologies about ourselves is enormous. Decentralized AI will be an appropriate strategy to minimize data requirements by smart devices. For instance, the cell phone might become a better predictive keyboard without forcing you to type in a particular manner, or without having the voice command software adapt to your voice patterns. Using this decentralization method, not only do we get an enhanced user experience from using AI technologies, but we also retain our privacy.

Of course, decentralization does not imply eliminating cloud computing entirely. There might be cases where the application requires greater access to information to perform adequately. Additionally, perhaps there are times when you wish to utilize certain functionalities within applications. Decentralized computation implies that devices can share data pertinent to the operation of applications. Instead of uploading your private information into the cloud before returning it to your device, the applications gain intelligence about your preferences through decentralization. It is necessary to upload some level of information into the cloud to allow applications to optimize themselves.

Decentralization proves more efficient during unstable network connectivity. It enhances performance in applications. However, decentralization is especially essential concerning securing user privacy.

5.3 How Decentralized Learning Changes Data Sharing

Machine learning in a classic scenario involves combining everybody's data into one single pool in the centre. The approach taken by federated learning flips this concept upside down.

In contrast to the classic method, instead of sending the original data to the central server, each organization trains the model independently using their local resources and only sends the learned patterns.

Federated learning enables collaboration among various organizations while ensuring that each organization does not reveal any of their confidential documents. Organizations like banks, schools, or hospitals could cooperate to develop more advanced algorithms while maintaining strict secrecy about their private data.

An even more decentralized strategy would involve removing the centralized controller altogether and making devices communicate with each other via networks.

5.4 Real-Life Example: Healthcare and Patient Privacy

Decentralized AI is definitely beneficial to many sectors; however, decentralized AI can prove highly advantageous within the field of healthcare.

Hospitals collect vast amounts of data from patients including medical history, test results, imaging scans, and treatment records. Such data are highly valuable in conducting medical research and training AI algorithms. Nonetheless, patient confidentiality becomes a critical concern here.

With decentralized AI, hospitals will have an opportunity to contribute to the creation of an AI model without sending their entire database of patients' records to a single central node. Instead, hospitals can train AI algorithms using internal databases while sending changes in training that may prove helpful for everyone else involved. Furthermore, decentralization allows for more precise diagnosis and prognosis by utilizing AI models without transferring any patient-specific data.

It might encourage people to use more AI-powered healthcare services due to increased confidence regarding the safety of their private information. Nonetheless, simple distribution across multiple nodes doesn't guarantee cybersecurity protection. Information sent from one hospital to another should remain confidential throughout the

communication. Decentralized AI in healthcare can prove safe and secure provided that appropriate security measures, governance frameworks, and ethical use of patient data are used.

5.5 From Individual Devices To Collective Intelligence

A particular device might hold little amount of information. However, when combined together, those devices would give an extensive understanding of the surroundings. It refers to collective intelligence, where valuable insights would emerge due to collaboration, yet the initial data would stay decentralized. Connected vehicles and gadgets in a town could detect the traffic status, pollution degree, road issues, and local occurrences.

Cars could inform each other regarding traffic flow and hazards, and personal devices could contribute to refining speech recognition and text prediction technologies without disclosing all of the individual's information. In addition to that, a decentralized system can decrease reliance on a central server significantly. Even if a specific node fails, others could keep functioning and contributing.

Yet, there exist limitations of decentralized AI, as well. Devices might possess various computational capacities, battery states, memory availability, and network connectivity. The type of information stored within a particular device could differ, too. Updates require using some power resources and networks, and securing them takes computing capacity, as well. Those are the parameters that should be accounted for in the design phase.

6. THE FUTURE OF DECENTRALIZED INTELLIGENCE

6.1 AI at the Edge: The Next Generation of Smart Devices

Some of the computations done by the artificial intelligence system might occur closer to where the data was produced. This phenomenon goes by the name of edge computing. It may occur in different devices, including cell phones, cars, camera systems, wearable gadgets, or industrial equipment. By analysing data locally, certain applications could become faster since they will no longer need to transfer the data to remote cloud servers. The car must react instantly when encountering something in its path, whereas a gadget may require instantaneous recognition of strange behaviour. Decentralized training paired with edge computing would enable each device to learn independently and collaborate with others.

6.2 Privacy-Preserving AI

Privacy-preserving AI refers to developing services that utilize minimum amounts of personal data. Federated learning is one approach to achieving this goal. Some other approaches include the use of encryption techniques, differential privacy techniques, and secure data exchange methods.

These techniques may contribute to protecting personal data as well as data exchanges among different devices. However, such techniques can increase complexity and computational requirements for AI systems. An effective technique should maintain an appropriate balance between privacy concerns and factors like precision, efficiency, expenses, and ease-of-use.

It is important to emphasize that privacy-preserving approaches do not need to be fully opaque or prevent any information exchange. The key ideas behind privacy-preserving technologies involve collecting, processing, sharing, explaining, and controlling data exchanges.

6.3 Opportunities For The Future

Decentralized intelligence could benefit multiple domains. Health institutions might collaborate in healthcare solutions while maintaining patients' confidentiality. Banking institutions might detect fraud more efficiently by using partial user data instead of full datasets. Learning apps might deliver personalized education experiences while safeguarding learner data privacy.

Devices might assist farmers to track soil quality, climate change, and crop conditions. Cities with smart infrastructure might leverage device data to optimize transportation, energy consumption, and pollution monitoring. Vehicles and robots could learn from nearby machines and respond better to changing situations.

This technology may also be helpful in rural or low-connectivity areas. Devices that can process information locally may continue to provide useful services without depending completely on a constant connection to a central server.

6.4 Challenges That Still Need to Be Solved

Although decentralised AI has many possibilities, it will not completely replace centralised computing. Devices in the same network may have different hardware, storage space, battery levels, and internet speeds. However, some devices might go offline or fail to perform specific tasks.

Data gathered by each device will also differ from one another. Individuals will have varying routines; institutions cater to various demographics, and sensors might detect objects under diverse conditions. Due to this variability, developing one machine learning model that is effective for every device would prove challenging.

Sharing is yet another issue that developers face. While the initial dataset remains localised on each device, any new updates acquired during training processes will still require transmission across the network. It would consume energy and utilise limited network bandwidth. Moreover, it will become crucial to differentiate between faulty and potentially harmful updates from reliable learning updates. In addition, any privacy-preserving techniques might hinder efficiency and accuracy.
Decentralised artificial intelligence requires advancements in several areas, including security protocols, communication, reliability, effectiveness, and explainability.

6.5 Ethical Questions And Trust In Decentralized AI

Data protection is undoubtedly crucial; however, privacy alone cannot solve the issue of building an ethical AI system. Other aspects, such as fairness, transparency, accountability, and equitable access to services, need to be taken into account as well. When some populations contribute more data compared to others, the latter's experiences will weigh more heavily in the development process. Thus, a solution that works efficiently among certain individuals or organizations but fails to operate effectively among other parties may cause further discrimination or bias. There should be transparency about the types of data involved, the purpose of using such information, and how they are stored securely.

Additionally, there needs to be accountability whenever an AI system causes any damage or makes wrongful predictions. Decentralization cannot become a shield to absolve organizations from taking any responsibility for the AI application's performance. Trustworthiness relies equally on advanced technology and good governance. Security features, privacy regulations, audits, informed consents, and user education will all matter significantly.

7. CONCLUSION: INTELLIGENCE WITHOUT A CENTRAL BRAIN

Decentralized AI does not indicate that cloud computing becomes unnecessary. Cloud servers will always be needed for storing, training, and running intensive operations involving massive data sets. Future cloud technologies can operate alongside edge computing, federated learning, encrypted computing, and peer-to-peer communication methods. However, the difference lies in where the intelligence resides. Intelligence will not need to rely on any centralized infrastructure. Instead, various devices, organizations, vehicles, sensors, and groups of people can participate in collaborative learning while maintaining better control of their data.
In order for decentralized AI to become effective, considerations about privacy, fairness, security, and responsibility must be addressed at the onset.
As students and future technology professionals, we can help create AI systems that are not only intelligent, but also safe, inclusive, reliable, and worthy of people's trust.

Ambient Intelligence and Context-Aware Computing

Faculty Mentor:
Dr. Akansha Upadhyaya

Students Name:
Vinayak Verma (MCA-I)
Namo Narayan Shukla(MCA-I)

When every piece of Technology has its own Mind

WHAT IS AMBIENT INTELLIGENCE AND CONTEXT-AWARE COMPUTING ?

And the answer is, Ambient Intelligence(AmI) is a Concept of computing where technology invisibly blends into physical spaces to proactively assist user without requiring explicit commands And Context-Aware computing act as the core engine of this ecosystem enabling devices to collect, interpret, and act upon environmental variables such as - Location, Time, Identity, and Physical State
- to understand what a user needs in real time. Anind K dey(2001)
To understand more about the topic first let's start by the question: How is when you put your phone down and go to sleep in default mode and when you wake up you find the phone was in the sleep mode, or in technical way we say in hibernation mode, similarly when we charge our phone during the day it charges normally till 100% and when you charge it overnight just to look at that the phone is stopped charging at 80% to preserve the battery health, All those task was completed because now modern smartphones are capable of Ambient Intelligence and Context awareness computing.

Now the next question in mind should be this further leads us to another question: HOW CONTEXT-AWARENESS FUELS AMBIENT INTELLIGENCE ?

To find out the reasoning behind this we need to know how traditional computing works and function, As we know the traditional computing heavily relies on explicit inputs like Click keypress or voice commands to give user the desired output, Context-Aware system process raw situational data through a three stage pipeline: Anind K dey(2001)

1. Sensing Layer : Physical and digital sensors gather real-time data such as - Ambient light, ambient audio, device location and temperatures.
2. Context Interpretation : Machine learning models translate disparate sensor streams into meaningful situation (for example: recognizing the "low light +Quiet room + night time " means the user of the device is probably in sleep)
3. Adaptive execution : Actuators and background process autonomously adjust te environment without user intervention (In this case the phone make itself to go to the sleep mode or hibernation mode to preserve Energy)

The example I gave you and the current modern smart phone's current Ambient intelligence are capable of is just a small part of what it could do in the future. We will get into that later in this article, For now let's focus on what changes it's making on how the user utilizes the technology.

FROM ACTIVE COMMANDS TO PASSIVE SENSING :

Traditional digital systems are reactive, waiting for explicit human input before taking action. context -aware systems flip this dynamic by continuously interpreting situational variables such as - current

geographical location, time of the day, proximity to other devices, on screen content of the current device, and many more such data points, also rather than treating this data points in isolation , a context-aware framework process them through the three-stage intelligence loop that we have read above and By cycling continuously through Sense, Interpret, and Act, technology transits from an active tool requiring constant instruction into an ambient partner operating quietly in the background, this is how the new smart phones can find and display the current playing song in a room in lockscreen even without the user need to unlock it.

Now, we know what Ambient intelligence and Context-aware computing is , and how it works along with the subtle changes its making on how the user interacts with the technology as nowadays people are stop memorising things like password because of the auto fill feature which uses the Ambient intelligence and context-aware computed that realise that if the opened website needs login credentials and look for the ones saved for the same website and then input those credentials in the correct boxes and what is left for the user to do is to press the log in button, everything just happens in an instant without any manual input from the user.

Beyond smartphones:

The practical impact of the ambient intelligence extends far beyond automated reply suggestions on your smart phone, or automatically setting the sleep mode or autofilling your credential in a website, for instance take healthcare, Ambient Clinical intelligence listens to physician - patient interactions, automatically generating medical documentation while monitoring high-risk patients for fall hazards Yawen Guo (2026). In industrial settings, context-aware factory equipment continuously measures micro-vibrations and thermal changes, automatically predicting component failures before operational downtime occurs.

Smart Healthcare:

The integration of artificial intelligence (AI) with the internet of medical things (IoMT) to create an Ambient intelligence environment in the medical context further enriches these concepts within healthcare, infrastructure of which include both wearable and non wearable medical devices to realize the AmI vision in healthcare settings. The IoMT technologies can facilitate the utilization of interconnected medical devices and sensors, providing unparalleled opportunities for real-time monitoring data collection, and remote patient management (Biomedical IoT).

The integration of Artificial Intelligence algorithms in medical services enhances their capabilities by enabling advanced data analytics, predictive modeling, and decision support functionalities. The benefits of such integration can be listed as follows

1. Automated home monitoring services
2. Reduce hospitals and hospice occupancy
3. Reduce healthcare cost
4. Personalized healthcare services
5. Predictive analysis for early disease detection

Collectively, these AI- powered technologies can usher in a new era of intelligent healthcare that prioritizes patient-centered care while optimizing operational efficiency.(smart healthcare by mekhla sarkar)

WHAT WILL BE ITS FUTURE?

The future or we can say vision of ambient intelligence lies in the complete disappearance of traditional user interfaces, shifting computing from screen-based interaction to an invisible environment-wide intelligence layer. As multimodal models, low power edge

hardware, and non-invasive sensors converge, spaces will transition from rule-based automation to a predictive goal oriented environment. Lets understand this with an example from the future, Instead of opening multiple application searching, and toggling windows to book a ticket or updating the schedule and meetings for an upcoming vacation yourself spatial wearables and agentic AI act on high-level intent, where you just stare a goal like "plan and execute a trip to Vegas for the first week of the next month" and the AI will handle everything from the comparing flight prices, buying the actual tickets, reserving a hotel, making a list for the recommended places to visit, and rescheduling all the meeting for that week and update it to your calendar just by your single line command.

THE HIDDEN COST OF INVISIBILITY : THE RISKS OF AMBIENT INTELLIGENCE

Behind the promise of effortless, screenless convenience lies a profound shift in control, privacy, and human agency. As ambient intelligence and context-aware computing weave sensors, microphone and predictive algorithms into our physical environments they introduce significant technical, psychological and ethical vulnerabilities some examples as follows :

1. The death of privacy and the continuous panopticon : For Context-aware systems to anticipate user needs, they must collect continuous, granular data. Unlike traditional smartphones that record data when an application opens, ambient environments are always-on recording user's every move all the time. Ethical Implication of AmI by Jurgen Bohn

2. Contextual Misinterpretation and loss of human control: Algorithms rely on probabilistic models to infer human intent, but human behaviour is notoriously chaotic and context-dependent leading system misinterpreting a command lead to irreversible actions due to its automated response system

3. The 'frictionless' trap : By stripping away physical buttons, toggles and manual interfaces, ambient spaces often leave users with no immediate way to override a malfunctioning automated system which also can be misuse by corporate to make user purchase what they want as AI can be manipulated and user might buy something different from what would be best for them because the AI recommendation is the only option available.

CONCLUSION

Ambient intelligence and context-aware computing mark a fundamental evolution in human-computer interaction, shifting technology from a tool we actively operate into an invisible environment that quietly anticipates our needs. By removing the friction of manual inputs and screen-based navigation, this transition promises unprecedented efficiency across everyday life, healthcare and industry. However, an invisible interface must not mean invisible governance. The ultimate success of ambient computing will depend on our ability to engineer robust privacy protection, preserve manual overrides, and keep human agency at the center of autonomous decision-making. As computing fades into the background our design priority must remain clear : building intelligent spaces that serve human intent without quietly taking control over it.

Quantum Computing: The Next Leap in Computation

Faculty Mentor:
Dr. Deepti Khanna

Students Name:
Jagjit Singh (MCA-I)
Mann Vats (MCA-I)

INTRODUCTION

For more than half a century, computing has been built on a simple foundation: the bit, a switch that is either 0 or 1. This binary logic has powered everything from pocket calculators to global data centres. Quantum computing challenges that foundation. Instead of bits, it uses quantum bits, or qubits, which exploit two strange properties of quantum mechanics — superposition and entanglement
— to represent and process information in fundamentally new ways.

The idea is not merely academic. In 2019, Google's 53-qubit Sycamore processor completed a specific sampling task in about 200 seconds, a task the team estimated would take a classical supercomputer thousands of years, an event widely reported as the first demonstration of "quantum supremacy" (Arute et al., 2019). By late 2023, IBM had unveiled its 1,121-qubit Condor processor (IBM Quantum, 2023), and multiple companies were racing toward fault-tolerant, error-corrected machines. This article explores how quantum computers work, what makes them powerful, the hardware and software ecosystems being built around them, where India fits into the global race, and the obstacles that still stand between today's experimental devices and a practical quantum future.

The theoretical seeds were sown decades before any working device existed. In 1982, physicist Richard Feynman argued that simulating quantum systems on classical computers was inherently inefficient and proposed building computers that themselves obeyed quantum rules. David Deutsch formalised the idea of a universal quantum computer in 1985, and by the mid-1990s, Peter Shor's and Lov Grover's algorithms gave the field its first concrete reasons to believe quantum machines could outperform classical ones on practical problems, not just simulate physics.

FUNDAMENTALS OF QUANTUM COMPUTING

Quantum computers do not simply run classical programs faster; they process information in a way that has no classical equivalent. Understanding this begins with three core ideas: superposition, entanglement, and quantum gates.

1. Qubits and Superposition

A classical bit is always definitely 0 or definitely 1. A qubit, by contrast, can exist in a superposition, a state written as $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where α and β are complex numbers whose squared magnitudes ($|\alpha|Z$ and $|\beta|Z$) give the probability of measuring 0 or 1, with $|\alpha|Z + |\beta|Z = 1$. Only when the qubit is measured does it "collapse" into one definite classical value. Because of superposition, n qubits can jointly describe 2^n possible states at once — but this larger state space does not, by itself, make a quantum computer faster. A genuine speed advantage only appears when an algorithm uses interference to amplify correct answers and cancel out wrong ones, which is why only certain problems, such as factoring or unstructured search, see a proven quantum speed-up.

2. Entanglement and Quantum Gates

Entanglement links two or more qubits so tightly that measuring one instantly reveals information about the state of the other, no matter the distance separating them — a phenomenon Einstein famously called "spooky

action at a distance." This does not allow faster-than-light communication: extracting any useful information from entanglement still requires comparing measurement results over an ordinary, light-speed-limited channel, so entanglement does not violate relativity. Entangled qubits nonetheless allow quantum computers to coordinate information in ways that classical bits never can, forming the backbone of quantum communication and many quantum algorithms.

2.1 Quantum Logic Gates

Just as classical circuits use AND, OR, and NOT gates, quantum circuits use quantum gates such as the Hadamard gate (which creates superposition), the Pauli-X gate (a quantum NOT), and the CNOT gate (which creates entanglement between two qubits). Unlike classical gates, quantum gates are reversible mathematical operations, represented as unitary matrices acting on qubit states.

2.2 Quantum Circuits and Coherence

A quantum circuit is a sequence of gates applied to a set of qubits, ending in measurement. The circuit must complete before "decoherence" — the loss of a qubit's fragile quantum state due to heat, vibration, or electromagnetic noise — destroys the information. Coherence times vary widely by platform and device: IBM's superconducting processors, for example, have reported qubit coherence times on the order of 100–300 microseconds (IBM Quantum, 2023), while trapped-ion qubits can retain coherence for seconds or longer. This is why superconducting quantum processors specifically are cooled to near absolute zero, colder than deep space, inside dilution refrigerators.

3. Quantum Algorithms

Quantum advantage depends on algorithms specifically designed to exploit superposition and entanglement. Shor's algorithm (Shor, 1997) can factor large integers in polynomial time on an ideal, fault-tolerant quantum computer — an exponential speed-up over the best known classical methods — which is why it threatens widely used encryption schemes such as RSA once sufficiently large, error-corrected quantum hardware exists. Grover's algorithm (Grover, 1996) offers a proven quadratic speed-up for searching unsorted databases. Variational algorithms such as the Quantum Approximate Optimization Algorithm (QAOA) are being explored for near-term, noisy devices where full error correction is not yet available.

QUANTUM HARDWARE PLATFORMS

Building a qubit that is stable enough to control yet isolated enough to avoid unwanted interference is one of the hardest engineering problems in modern physics. No single technology has won outright, and several competing platforms are being developed in parallel.

1. Superconducting Qubits

Used by IBM, Google, and Rigetti, superconducting qubits are tiny circuits made from superconducting metal that behave like artificial atoms when cooled to roughly 10–20 millikelvin (about 15 millikelvin for typical IBM and Google systems) — colder than outer space — inside a dilution refrigerator (IBM Quantum, 2023). This extreme cooling requirement is specific to superconducting-qubit hardware and does not apply to platforms such as trapped ions or photonics. Superconducting qubits rely on Josephson junctions to create discrete, controllable energy levels, offer fast gate operations, and benefit from mature semiconductor fabrication techniques, but they demand extreme cooling infrastructure and currently suffer from relatively short coherence times.

2. Trapped-Ion Qubits

Companies such as IonQ and Quantinuum trap individual charged atoms using electromagnetic fields and manipulate their internal energy states with precisely tuned lasers. Trapped-ion qubits typically achieve longer coherence times and higher gate

fidelities than superconducting qubits, since every ion of a given species is naturally identical. The trade-off is speed: laser-based gate operations are generally slower, which can limit how many operations fit within a qubit's usable lifetime.

2.1 Photonic Qubits

Photonic approaches, pursued by Xanadu and PsiQuantum, encode information in properties of light such as polarisation or path. Because photons barely interact with their environment, photonic qubits can operate at room temperature and travel easily through optical fibre, which is attractive for quantum networking. Photon loss during transmission and the difficulty of making photons interact strongly enough to form logic gates remain significant hurdles.

2.2 Topological Qubits

Microsoft's long-term bet is on topological qubits, which would encode information in the braided paths of exotic quasiparticles called non-abelian anyons. Because information here is stored in a global topological property rather than a fragile local state, these qubits are theoretically far more resistant to noise. The approach remains largely experimental, with the underlying physics still being validated in research settings.

3. Comparing the Platforms

No platform currently dominates on every metric. Superconducting qubits lead on qubit count and gate speed (typical two-qubit gate times of tens of nanoseconds) but have coherence times of roughly 100–300 microseconds; trapped ions lead on gate fidelity (often above 99.9%) and connectivity but have slower, microsecond-to-millisecond gate speeds; photonics leads on room-temperature operation and networking but faces higher photon-loss rates; topological qubits promise built-in error resistance but remain unproven at scale (Nielsen & Chuang, 2010; Preskill, 2018). Many researchers expect the eventual winner, if one emerges, may combine ideas across platforms rather than following a single approach exclusively.

QUANTUM SOFTWARE ECOSYSTEM

1. Programming Frameworks

Writing quantum programs no longer requires specifying hardware pulses by hand. Open-source frameworks such as IBM's Qiskit, Google's Cirq, Microsoft's Q#, and Xanadu's PennyLane let developers describe circuits at a high level, in Python or a dedicated quantum language, and compile them down to the gate set of a specific quantum processor. These toolkits also include simulators that mimic quantum behaviour on ordinary classical computers, which is how most students and researchers first learn to write and debug quantum code before ever touching real hardware.

For example, the following short Qiskit program builds a two-qubit Bell state (an entangled pair) and measures it — often the first circuit students run:

```
from qiskit import QuantumCircuit
qc = QuantumCircuit(2, 2)
qc.h(0)      # put qubit 0 into superposition
qc.cx(0, 1) # entangle qubit 0 with qubit 1
qc.measure([0, 1], [0, 1]) # measure both qubits
```

2. Cloud Access To Real Quantum Hardware

Perhaps the most important democratising development in the field is cloud access. Platforms such as IBM Quantum Experience, Amazon Braket, and Microsoft Azure Quantum allow anyone with an internet connection to submit circuits to real quantum processors and retrieve results, often for free at small scale — though exact queue times (ranging from minutes to hours depending on demand), free-tier limits, and available hardware vary by provider and change over time. This has turned quantum computing from a resource available only to a handful of

national laboratories into something a university student or hobbyist can experiment with directly, dramatically accelerating education and grassroots research worldwide.

INDIA'S QUANTUM AMBITIONS

India has entered the global quantum race with the National Quantum Mission, approved by the Union Cabinet in April 2023 with an outlay of about ₹6,003 crore through 2030–31 (Department of Science and Technology, Government of India, 2023). The mission, coordinated by the Department of Science and Technology, aims to develop intermediate-scale quantum computers with 50 to 1,000 physical qubits within eight years, alongside advances in quantum communication, quantum sensing, and quantum materials. Four Thematic Hubs have been set up at leading institutions, including the Indian Institute of Science, Bengaluru, and IIT Madras, IIT Bombay, and IIT Delhi, to anchor research across these domains. For MCA and engineering students in India, this national push is opening new academic and research pathways into a field that, until recently, was pursued mainly abroad.

APPLICATIONS AND CHALLENGES

1. Where Quantum Computing Could Matter Most Interest in quantum computing is not evenly spread across industries. A handful of problem types keep recurring in research papers and corporate pilot projects because they map naturally onto what quantum hardware does well: modelling probabilistic, interacting systems rather than executing long sequences of deterministic steps.

1.1 Cryptography and Cybersecurity

Shor's algorithm means a sufficiently large, fault-tolerant quantum computer could, in the future, break the RSA and elliptic-curve encryption that currently secures banking, government, and internet traffic — a capability that does not yet exist on today's noisy, error-prone hardware. In anticipation of this, there has been a global push toward "post-quantum cryptography" — classical encryption schemes believed to resist quantum attacks — with the U.S. National Institute of Standards and Technology finalising its first post-quantum standards in 2024 (NIST, 2024).

1.2 Drug Discovery and Chemistry

Molecules are quantum systems, so simulating them accurately is naturally suited to quantum hardware. Pharmaceutical firms such as Roche and Merck, along with quantum computing companies, are already running early, small-scale proof-of-concept experiments to model simple chemical reactions and protein interactions; simulating pharmaceutically relevant molecules at full accuracy remains a future goal that will likely require fault-tolerant hardware (Preskill, 2018).

1.3 Optimization and Logistics

Airlines, banks, and logistics firms are piloting quantum and quantum-inspired algorithms for problems such as flight-crew scheduling, portfolio risk balancing, and delivery-route optimization — problems where the number of possible combinations grows so quickly that classical computers can only search a small fraction of them. These remain pilot studies to date; a demonstrated, reproducible quantum advantage over the best classical optimization heuristics has not yet been established for practical business problems (Preskill, 2018).

1.4 Artificial Intelligence

Hybrid quantum-classical machine learning, where a quantum circuit handles a specific sub-routine within an otherwise classical model, is an active area of academic research. Claims of quantum advantage in AI remain contested and largely unproven at practical scale, and most experts view this as a longer-term, less certain application than cryptography or chemistry (Preskill, 2018).

2. Table: Classical vs Quantum Computing

Feature	Classical Computing	Quantum Computing
Basic Unit	**Bit (0 or 1)**	**Qubit (can exist in superposition)**
State	**Definite(0 or 1)**	**Probabilistic until measured**
Scaling	**Linear resource growth**	**State space grows as 2^n (not itself a speed-up)**
Hardware	**Room-temperature transistors**	**Near-absolute-zero qubits**
Error Rate	**Extremely low**	**High; needs error correction**
Best Suited For	**General-purpose tasks**	**Simulation, factoring, optimization**

Table 1: Classical vs Quantum Computing

3. KEY CHALLENGES

Despite rapid progress, several obstacles stand between today's devices and widespread practical use. Decoherence and noise mean physical qubits make errors far more often than classical transistors, so quantum error correction is essential — but current schemes, such as the surface code, are estimated to require on the order of a thousand noisy physical qubits to form a single reliable "logical" qubit, with the exact ratio depending on the physical error rate, architecture, and error-correction method used (Preskill, 2018); this implies a useful fault-tolerant computer could need anywhere from hundreds of thousands to millions of physical qubits. Scalability is a related engineering challenge: wiring, control electronics, and cooling systems all become harder to manage as qubit counts grow. There is also a global shortage of engineers and physicists trained in quantum information science, and the specialised cryogenic and photonic equipment involved makes research infrastructure expensive to build and maintain.

4. THE ROAD AHEAD: NISQ AND BEYOND

Today's devices belong to what physicist John Preskill termed the "Noisy Intermediate-Scale Quantum" (NISQ) era (Preskill,2018): powerful enough to explore, too error-prone for guaranteed practical advantage on most real-world problems. The next milestone is fault-tolerant quantum computing, which uses many physical qubits to form a smaller number of highly reliable "logical" qubits through quantum error correction. IBM's 2023 quantum development roadmap update targets early error-corrected systems later this decade (IBM Quantum, 2023) — a timeline that, like all technology roadmaps, is subject to revision. Google, IonQ, and other players pursue competing hardware approaches — superconducting circuits, trapped ions, and photonics — each with distinct trade-offs in coherence time, gate speed, and scalability.

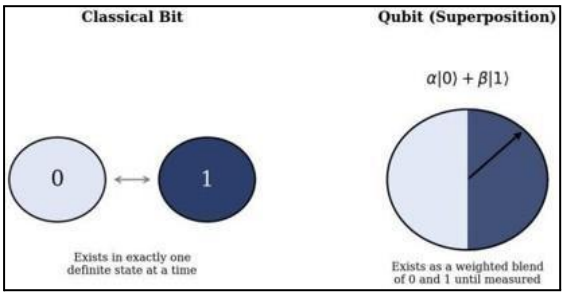


Fig 1: Classical bit vs. qubit in superposition

Figure 1 contrasts a classical bit, which is always definitely 0 or 1, with a qubit's Bloch-sphere representation, in which the qubit's state can point anywhere on the sphere's surface until measurement forces it to collapse to one of the two poles (0 or 1).

5. Quantum Computing Timeline: Key Milestones

Year	Milestone
1982	**Richard Feynman proposes computers based on quantum mechanics for simulating quantum systems.**
1985	**David Deutsch formalises the concept of a universal quantum computer.**
1994	**Peter Shor publishes a quantum algorithm for fast integer factorization.**
1996	**Lov Grover publishes a quantum algorithm for fast unstructured search.**
2019	**Google's 53-qubit Sycamore processor claims the first demonstration of quantum supremacy.**
2023	**IBM unveils the 1,121-qubit Condor processor; India launches its National Quantum Mission.**
2024	**U.S. NIST finalises its first post-quantum cryptography standards (FIPS 203).**

Table 2: Quantum Computing Timeline

CONCLUSION

Quantum computing will not replace classical computers; it will complement them, taking on the narrow class of problems — simulating molecules, breaking certain codes, exploring vast optimization landscapes — where quantum mechanics itself offers the shortcut. The field has moved from theoretical curiosity to global race in less than three decades, with billions of dollars invested by IBM, Google, Microsoft, Amazon, and national governments, including India's own National Quantum Mission. Open-source software frameworks and cloud-accessible hardware mean that, for the first time, students anywhere in the world can write and run real quantum circuits rather than only reading about them. For MCA students specifically, a practical starting point is learning Python alongside a framework such as Qiskit or Cirq, brushing up on the linear algebra behind quantum states, and running small circuits on free cloud simulators and hardware before moving on to research-level topics such as error correction or variational algorithms.

Whether the first genuinely useful, fault-tolerant quantum computer arrives in five years or fifteen, the foundational work being done today in labs, start-ups, and universities

— including student projects exploring these same fundamentals

— is laying the groundwork for what may become the defining computational technology of the twenty-first century.

REFERENCES

1. Nielsen, M. A., & Chuang, I. L. (2010). Quantum computation and quantum information (10th anniversary ed.). Cambridge University Press.

2. Preskill, J. (2018). Quantum computing in the NISQ era and beyond. Quantum, 2, 79. https://doi.org/10.22331/q-2018-08-06-79

3. Arute, F., Arya, K., Babbush, R., et al. (2019). Quantum supremacy using a programmable superconducting processor. Nature, 574(7779), 505–510. https://doi.org/10.1038/s41586-019-1666-5

4. Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM Journal on Computing, 26(5), 1484–1509.

5. Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. Proceedings of the 28th Annual ACM Symposium on Theory of Computing, 212–219.

Neuromorphic Computing: Designing Machines That Think Like the Brain

Faculty Mentor:
Dr. Deepti Sharma

Students Name:
Gurpreet Kaur (MCA-I)
Megha (MCA-I)

INTRODUCTION

For nearly eight decades, digital computers have used the same basic blueprint: a processor that fetches instructions and data from separate memory, executes them step at a time, and writes the results back. This design, first formalised by John von Neumann in the 1940s, has scaled remarkably well, but it fundamentally mismatches the demands of modern artificial intelligence. Training and running large neural networks requires moving enormous volumes of data between memory and processor, and this movement, not the arithmetic itself, now consumes most of the energy and time in AI systems.

The human brain offers a strikingly different template. It contains on the order of eighty-six billion neurons and hundreds of trillions of synaptic connections, yet it performs perception, reasoning, and motor control on a power budget of roughly twenty watts – far less than a single modern graphics processor. It achieves this efficiency by storing and processing information in the same physical location, communicating through brief, sparse electrical events called spikes, and remaining largely silent except when something meaningful happens. Neuromorphic computing is the discipline that tries to translate these organising principles – distributed memory, event-driven communication, and local learning – into silicon. The field traces its origins to Carver Mead's work on analogue very-large-scale-integration circuits in the late 1980s, and it has since grown into a broad research area spanning device physics, circuit design, algorithms, and cognitive science. This article introduces the motivation for neuromorphic computing, explains how biological neurons inspire artificial spiking models, surveys hardware platforms that have emerged from this research, and discusses the challenges that still stand between today's prototypes and truly brain-like machines.

1. REDEFINING COMPUTER ARCHITECTURE THROUGH NEUROSCIENCE

Before going into how neuromorphic systems are built, it is useful to understand precisely which restriction of conventional computing they are trying to overcome, and which biological mechanisms they borrow to do so.

1.1 The Von Neumann Bottleneck

In a conventional processor, every operation – however small requires data to travel across a bus connecting the arithmetic units to memory. As models have grown larger, this constant back-and-forth has become the dominant cost of computation, a phenomenon widely referred to as the von Neumann bottleneck. Because memory access consumes considerably more energy than the arithmetic itself, systems built on this architecture spend an increasing share of their power budget simply shuttling numbers back and forth rather than computing with them. Figure 1 contrasts this arrangement with a neuromorphic layout, in which small, self-contained cores keep their own memory locally and exchange only brief messages when necessary.

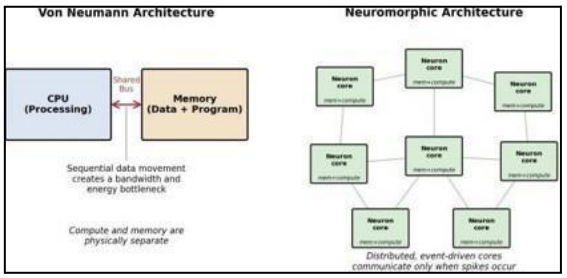


Fig 1: Centralised von Neumann architecture compared with a distributed, memory-local neuromorphic architecture

1.2 Learning from Biology

Neuromorphic engineering does not attempt to copy the brain exactly; rather, it borrows a small number of organising principles that appear to explain much of the brain's efficiency, and reproduces them in electronic form.

1.2.1 Neurons and Synapses as Computational Primitives

A biological neuron collects incoming electrical signals through its dendrites, integrates them at the cell body, and, once the accumulated potential crosses a threshold, fires a brief electrical pulse along its axon to other neurons through junctions called synapses. Synapses are not fixed; their strength changes with experience, which is the physical basis of learning and memory. Artificial spiking neurons reproduce this behaviour in a simplified mathematical form, most commonly the leaky integrate-and-fire model shown in Figure 2. Incoming spikes add to a virtual membrane potential that slowly decays, or 'leaks', over time; when the potential exceeds a set threshold, the artificial neuron emits its own spike and resets.

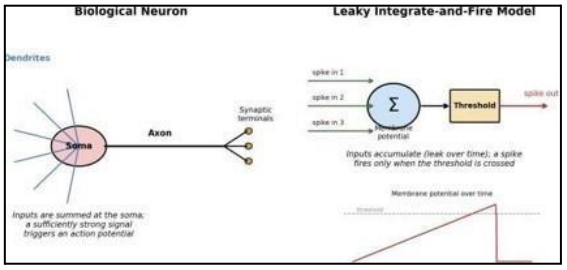


Fig 2: A biological neuron and its simplified artificial counterpart, the leaky integrate-and-fire model

1.2.2 Spiking Neural Networks and Temporal Coding

Networks built from these artificial neurons are known as spiking neural networks. Unlike conventional artificial neural networks, which exchange continuous numeric values at every layer during every forward pass, spiking networks exchange only discrete, timed events. Information can therefore be encoded not just in how many spikes occur, but in exactly when they occur relative to one another, a strategy known as temporal coding. Because most neurons in a biological or neuromorphic network are silent at any given instant, computation becomes inherently sparse. Figure 3 illustrates this contrast: a conventional system samples every unit on every clock cycle regardless of activity, whereas an event-driven system performs work only where and when a spike demands it, which is the principal source of the large energy savings reported in neuromorphic hardware studies (Roy, Jaiswal, & Panda, 2019).

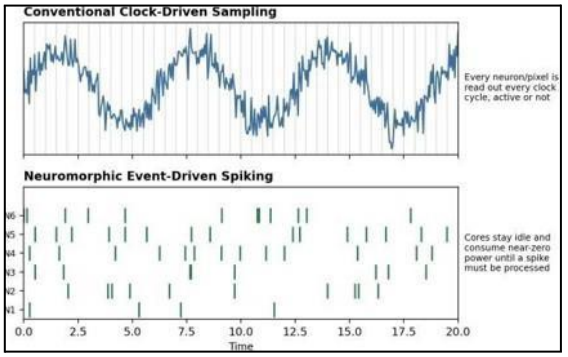


Fig 3: Clock-driven sampling of every unit compared with sparse, asynchronous, event-driven spiking

1.3 From Concept to Silicon: Neuromorphic Hardware Platforms

Translating spiking neuron models into working chips has produced several notable research platforms over the past decade, each exploring a different point in the design space between biological realism, programmability, and manufacturing scale. Table 1 summarises four influential examples.

Chip	Developer	Neurons (approx.)	Key Distinguishing Feature
TrueNorth	IBM Research	1 million	65 mW fully digital neurosynaptic chip with 256 million programmable synapses (Merolla et al., 2014).
Loihi	Intel Labs	130,000	On-chip spike-timing-dependent learning with 128 asynchronous neuromorphic cores (Davies et al., 2018).
SpiNNaker	University of Manchester	Up to 1 billion (full machine)	Massively parallel ARM-core mesh built to simulate spiking networks in biological real time (Furber et al., 2014).
Tianjic	Tsinghua University	40,000	Hybrid architecture supporting both spiking and artificial-neural-network computation on one chip (Pei et al., 2019).

Table 1: A comparison of representative neuromorphic hardware platforms

IBM's TrueNorth demonstrated that a fully digital chip could emulate a million neurons while consuming only tens of milliwatts, by keeping every core event-driven and clock-free. Intel's Loihi went further by adding on-chip learning, allowing synaptic weights to adapt locally through rules inspired by spike-timing-dependent plasticity rather than relying solely on off-chip training. The University of Manchester's SpiNNaker took a different route, using large numbers of conventional low-power processor cores connected by a custom communication fabric to simulate spiking networks at, or faster than, biological real time. More recently, Tsinghua University's Tianjic chip showed that spiking and conventional artificial-neural-network computation could coexist on a single hybrid architecture, a step that proved influential for combining the efficiency of spiking hardware with the mature tool-chains built around deep learning.

2. APPLICATIONS, CHALLENGES, AND THE PATH FORWARD

2.1 Where Neuromorphic Computing Is Making an Impact

The combination of low power consumption and event-driven processing makes neuromorphic systems particularly attractive wherever a device must sense and react continuously on a limited energy budget. Event-based cameras, which report only pixel-level brightness changes rather than full image frames, pair naturally with spiking processors for high-speed, low-power vision in robotics and autonomous vehicles. Wearable and implantable health-monitoring devices benefit from the same efficiency when processing continuous physiological signals such as electrocardiograms. Neuromorphic chips have also been explored for solving certain optimisation and constraint-satisfaction problems, and for on-device keyword spotting and gesture recognition, where always-on sensing must not drain a battery.

2.2 Open Challenges

Despite these promising demonstrations, neuromorphic computing has not yet become a mainstream computing platform, and several practical obstacles explain why.

2.2.1 Programming Models and Algorithm-Hardware Co-design

Decades of tooling, compilers, and libraries have been built around conventional processors and, more recently, around dense matrix operations for deep learning accelerators. Spiking hardware has no equivalent mature ecosystem: programming a neuromorphic chip typically means reasoning directly about spike timing and neuron parameters, a very different skill from writing ordinary software. Closing this gap requires close co-design between algorithms and hardware, so that spiking networks can be specified at a high level and compiled efficiently onto diverse neuromorphic chips (Roy et al., 2019).

2.2.2 Learning Rules and On-Chip Training

Training deep neural networks relies on the back-propagation algorithm, which assumes smooth, differentiable activation functions. The discrete, all-or-nothing nature of a spike makes standard back-propagation difficult to apply directly, and biologically plausible alternatives such as spike-timing-dependent plasticity are still less effective on complex tasks. Much current research therefore focuses on approximate or surrogate-gradient methods that make spiking networks trainable with conventional tools while preserving their event-driven behaviour once deployed.

2.3 The Road Ahead

Neuromorphic computing is unlikely to replace conventional processors outright; rather, it is expected to complement them, taking over the always-on, sensory, and pattern-recognition workloads for which its energy profile is best suited, while conventional and graphics processors continue to handle general-purpose and large-scale training tasks. Progress in emerging memory devices, such as memristive synapses that combine storage and computation in a single nanoscale element, may eventually allow neuromorphic chips to approach the density and adaptability of real neural tissue. As algorithms, compilers, and fabrication techniques mature together, machines that compute the way brains do may move from specialised laboratory demonstrations into everyday low-power intelligent devices.

REFERENCES

1. Davies, M., Srinivasa, N., Lin, T.-H., Chinya, G., Cao, Y.,

Choday, S. H., Dimou, G., Joshi, P., Imam, N., Jain, S.,

Liao, Y., Lin, C.-K., Lines, A., Liu, R., Mathaikutty, D., McCoy, S., Paul, A., Tse, J., Venkataramanan, G., … Wang, H. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. IEEE Micro, 38(1), 82–99.

2. Furber, S. B., Galluppi, F., Temple, S., & Plana, L. A. (2014). The SpiNNaker project. Proceedings of the IEEE, 102(5), 652–665.

3. Maass, W. (1997). Networks of spiking neurons: The third generation of neural network models. Neural Networks, 10(9), 1659–1671.

4. Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy,

A. S., Sawada, J., Akopyan, F., Jackson, B. L., Imam, N.,

Guo, C., Nakamura, Y., … Modha, D. S. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. Science, 345(6197), 668–673.

5. Pei, J., Deng, L., Song, S., Zhao, M., Zhang, Y., Wu, S.,

Wang, G., Zou, Z., Wu, Z., He, W., … Shi, L. (2019).

Towards artificial general intelligence with hybrid Tianjic chip architecture. Nature, 572(7767), 106–111.

6. Roy, K., Jaiswal, A., & Panda, P. (2019). Towards spike-based machine intelligence with neuromorphic computing. Nature, 575(7784), 607–617.

Brain – Computer Interfaces: Computing Beyond Screens

Faculty Mentor:
Dr. Sanjive Saxena

Students Name:
Anvesha Gupta (MCA-I)
Puneet Kaur (MCA-I)

INTRODUCTION

Brain-computer interfaces (BCI) are communication channels that enable interaction between the brain and an external device by decoding electrical signals, bypassing the need for manual control of a computer or robotic arm using hands or legs. Their basic purpose is to "restore, enhance, or otherwise replace" a human function, and they take advantage of the brain's ability to form new connections through neuroplasticity to adapt to the presence of a prosthesis. Depending on how invasively they interact with the brain, different BCI's can have different risks and benefits. BCIs interact with the brain in several ways, including non-invasive techniques like EEG, MEG, or MRI; partially invasive techniques like ECoG or endovascular approaches; and fully invasive microelectrode arrays.

THE JOURNEY OF THOUGHT

The development of brain–computer interfaces (BCIs) spans nearly a century of neuroscientific and technological advancements:

· *1924 (EEG Discovery):* The first recording of the activity of the human brain using an EEG (electroencephalography) method was made by Hans Berger who observed prominent oscillations of alpha waves using a Siemens double-coil galvanometer.

· *1965 (Early BMI Application):* Music for Solo Performer was the first BMI application created by an American composer, Alvin Lucier, and used an EEG with analog circuits to transform alpha waves into kinetic energy to play various percussion instruments.

· *1973-1977 (Invention of BCI):* The term "brain–computer interface (BCI)" was introduced by Jacques Vidal in 1973 research and the 'BCI challenge' was defined and in 1977, the first non-invasive EEG trial was also conducted.

· The subject used visual evoked potentials (VEP) to move a cursor on a computer screen around a simple maze.

· *1988 (Robotic Control):* Scientists demonstrated for the first time non-invasive control of a robotic device using EEG.

· The robot was directed to passively follow a line drawn on the floor and instructed to move from start to finish, 1990 (Closed-Loop BCI): The first closed-loop (bidirectional) adaptive BCI was built.

· A contingent negative variation (CNV) brain potential named Electroexpectogram (EXG) was used to control the computer buzzer in a feedback loop, 2010s (Neural Restoration and Passive BCI): Restoration of function was achieved through the stimulation of certain molecules in the system with neural stimulation.

· In 2011, Thorsten O. Zander introduced a passive BCI and DARPA was beginning to fund the BRAIN initiative in 2013.

WHEN THE BODY ISN'T THE INTERFACE

When restricted or absent, typical forms of manual movement are ineffective for interacting with a computer. The brain-computer interface bypasses the neuromuscular system by interpreting electrical signals in the brain as commands.

·**Communication:** Non-invasive EEG or invasive microelectrode recordings can be

decoded into information on a digital device. People who are physically incapacitated or have restricted motor function (e.g., due to a condition such as Amyotrophic Lateral Sclerosis) can type or navigate a computer interface using only their brain.

·**Prosthetic Limbs**: The BCIs allow for two-way communication with advanced robotic limbs.The movement of the motor cortex sends signals to the robotic limbs, and the feeling of touch can be sent back to the sensory cortex, allowing a person to move and "feel" a robotic arm as an organic part of their body.

· **Wheelchair Control**: The movement of wheelchairs can be controlled by thinking about how you want to move your left and right hands.

· Some systems use a combination of thought and vision to avoid obstacles and allow for natural movement in different terrains.

·**Rehabilitation:** Using a BCI in neurorehabilitation can help the brain create new connections by linking the intention of movement in the motor cortex with the movement of a robotic limb or the use of functional electrical stimulation to stimulate the muscles. Assistive Technology: Using neural signals, a person can control their environment independently and conveniently from controlling smart devices to robotic limbs, BCI can be used as a universal remote for a person's independent lifestyle.

BEYOND HEALTHCARE: COULD THINKING BECOME AN INPUT DEVICE?

Although brain–computer interfaces are primarily developed within the scope of neuroprosthetics and the rehabilitation of people with disabilities, their evolution is destined to go beyond medicine and enter the realm of regular human–computer interactions. The shift from therapeutic implants to consumer devices represents the fundamental change in how people engage with technology.

·Gaming

Now: Can track simple focus or binary-choice decisions

In Future: Could allow us to make decisions and move around virtual environments using our brains

·Virtual and Augmented Reality

Now: Virtual and augmented reality headsets use video cameras and eye-tracking sensors

In Future: Think about moving and manipulating items in a virtual space

·Robotics

Now: Robots are largely confined to highly structured environments such as factories

In Future: We might be able to control robotic systems through our thoughts

· Smart Homes

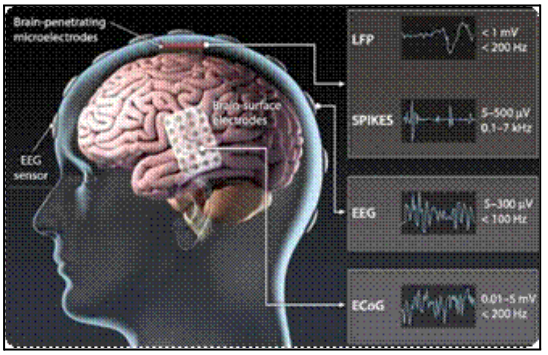
Now: Smart homes are controlled by voice or manual input

In Future: Wearable technology could change aspects of the home environment based on thought

·Everyday Computing

Now: Requires touchscreens, mice, and keyboards

In Future: Typing on our thoughts would replace current means of computing



SCREENLESS COMPUTING: THE NEXT INTERFACE?

Transitioning toward screenless computing doesn't mean they'll disappear overnight.

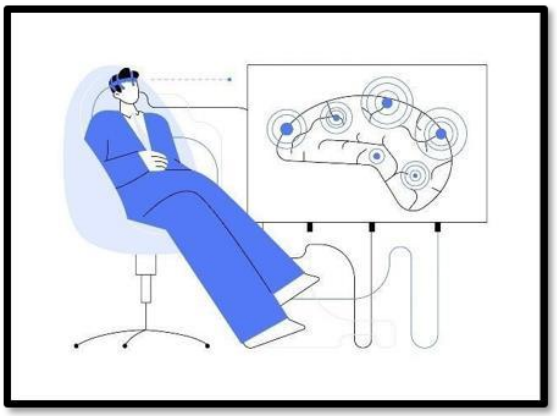
Monitors, displays, and screens are the easiest medium for conveying high levels of visual information such as viewing complex documents, watching video, or inspecting intricate graphics.

The BCI is also used as a supplementary layer of control to allow for immediate access to navigation or control, without requiring tactile or verbal feedback, and to free the visual display for more important information

AUGMENTED COGNITION & COGNITIVE LOAD BALANCING

Beyond the control of various devices and removal of screens, we can imagine a future in which brain implants provide a passive neural monitoring capability for the purpose of dynamic real-time adaptation.

Rather than using the device to translate our commands (i.e., 'click here' or 'type this'), we can think of the 'next generation' devices as ways to best interpret the information needed.
A few possible applications might include:
Dynamic Cognitive Offloading: Imagine a future in which non-invasive neural sensors can detect when I'm fatigued, overloaded, or focused.
In such a world, my smartphone might automatically reduce the amount of new information it tries to push to me, filter out unimportant notifications, or simplify the user interface when it detects I'm experiencing cognitive stress.
Accelerated Adaptive Learning: In a similar vein, imagine an educational or training software that could detect when you are beginning to forget what it was just teaching you.
Using a combination of BCI and biofeedback, such a program could be used to help determine what methods, media, or pacing are best suited for your learning.
Emotional and Affective Computing: With enough information about your frustration, focus, or anxiety levels, BCI could be used to help dynamically alter your work environment. For example, perhaps your smartwatch could recognize when you become frustrated and recommend a break or reduce it's visual output if it detects you're becoming distracted.
...and so on.



OUR VISION

Academic Excellence With Purpose To be an Institute of Academic Excellence with total commitment to quality education and research in Management and Information Technology with a holistic concern for better life, environment and society.

OUR MISSION

Transforming Education Into Impact To serve society and improve quality of life through quality education, training, research, consultancy and knowledge dissemination.



JAGAN INSTITUTE OF MANAGEMENT STUDIES

3, Institutional Area, Sector-5, Rohini, Delhi

techbyte@jimsindia.org

www.jimsindia.org