RESEARCH INSIGHTS REPORT · MAY 2026

Artificial Intelligence & Machine Learning

Research Insights 2025–2026

A comprehensive analysis of breakthrough research across Large Language Models, Generative AI, Deep Learning, AI Ethics, Healthcare AI, Federated Learning, and beyond.

60+	8	2026
Papers Reviewed	Research Domains	Latest Findings

Atine John Bosco

Chartered Information Scientist · MBCS · CISA

PhD Candidate in Computer Science

www.atine.info

Published May 2026 · Based on peer-reviewed literature from SciSpace, Google Scholar, ArXiv, and PubMed

Contents

1 Executive Summary **2**

2 Large Language Models and Generative AI **2**
2.1 The Scale Revolution . . . 3
2.2 Architectural Breadth and Generative Families . . . 3
2.3 Multimodal Emergence . . . 3

3 Deep Learning and Neural Network Innovations **4**
3.1 Self-Reflective Training Paradigms . . . 4
3.2 Efficient Scaling Strategies . . . 4
3.3 Hybrid Optimisation and Verifier Guidance . . . 5

4 AI Ethics, Governance, and Responsible AI **5**
4.1 Policy-Level Urgency . . . 5
4.2 Responsible Deployment Imperatives . . . 6
4.3 Alignment Bottlenecks . . . 6
4.4 Evaluation and Transparency . . . 6

5 AI in Healthcare and Medicine **6**
5.1 LLMs in Clinical Workflows . . . 6
5.2 Generative AI Opportunities and Caveats . . . 7
5.3 Federated Learning in Medical Imaging . . . 7

6 Federated Learning and Privacy-Preserving AI **8**
6.1 Decentralised Machine Learning Paradigms . . . 8
6.2 Technical Advances in 2025–2026 . . . 9
6.3 Remaining Challenges . . . 9

7 Reinforcement Learning Advances **9**
7.1 RL for LLM Alignment . . . 9
7.2 RL for Scientific Discovery . . . 10

8 AI for Scientific Discovery **10**
8.1 The AI Revolution in Chemistry . . . 10
8.2 Autonomous Drug Discovery . . . 11
8.3 AI in Materials Science . . . 11

9 AI in Cybersecurity **12**
9.1 Generative AI as a Dual-Use Technology . . . 12
9.2 AI-Powered Defence Mechanisms . . . 12

9.3 Privacy and Data Protection 13

10 Key Challenges and Future Directions **13**

10.1 Cross-Cutting Challenges 13

10.2 Promising Research Directions 14

10.3 Implications for ICT Professionals 14

11 Conclusion **15**

References **17**

About the Author **19**

1 Executive Summary

Key Takeaways

- **LLMs entered the trillion-parameter era**, with distributed training and multimodal integration defining the 2025–2026 frontier.
- **Generative AI** expanded beyond text into unified vision-language-audio models, with GPT-4o successors setting new interaction benchmarks.
- **Reinforcement Learning from Human Feedback (RLHF)** remained the dominant alignment paradigm, while novel self-reflective RL variants (GRPO) demonstrated measurable gains on multimodal tasks.
- **AI Ethics and Governance** emerged as a strategic priority, with policy bodies calling for safety-first deployment frameworks and transparent evaluation ecosystems.
- **Healthcare AI** advanced through federated learning, privacy-preserving clinical models, and LLM-driven documentation and decision support.
- **AI in Chemistry and Drug Discovery** saw an **8× increase in biochemistry patents** and autonomous LLM-driven discovery frameworks gaining traction.
- **Key challenges** persist: model robustness, interpretability, data governance, compute energy costs, and equitable AI access.

The years 2025–2026 represent a watershed period for artificial intelligence and machine learning. The confluence of unprecedented model scale, multimodal capability expansion, and growing societal awareness of AI's risks and benefits has reshaped both research priorities and industry practice. This report synthesises findings from over 60 peer-reviewed papers, preprints, and technical surveys drawn from SciSpace, Google Scholar, ArXiv, and PubMed, providing a structured overview of the most significant developments across eight research domains.

From trillion-parameter language models to federated healthcare systems, from autonomous drug-discovery pipelines to AI governance frameworks, the breadth of progress is extraordinary. Yet the field also confronts persistent challenges: reward hacking in alignment, interpretability gaps, energy consumption, and the widening divide between resource-rich and resource-limited AI actors. This report aims to equip ICT professionals, researchers, and decision-makers with the insights needed to navigate this rapidly evolving landscape.

2 Large Language Models and Generative AI

2.1 The Scale Revolution

Large Language Models entered the trillion-parameter era in 2025–2026, enabled by advances in distributed training and model-parallel architectures. Research confirms that scaling transformer-based models toward these regimes required overcoming fundamental compute and memory constraints through novel parallelism strategies [1].

Scale Milestones 2025–2026

- Trillion-parameter models became viable through pipeline and tensor parallelism
- Mixture-of-Experts (MoE) architectures enabled efficient sparse activation
- Context windows expanded to 1M+ tokens for long-document reasoning
- Inference optimisation (quantisation, speculative decoding) reduced deployment costs

2.2 Architectural Breadth and Generative Families

Transformer backbones continued to dominate text modelling while diffusion architectures gained prominence for high-fidelity image synthesis and multi-step generation tasks [2]. The period saw the emergence of three dominant generative model families, each with distinct strengths:

navymid

Transformers / LLMs	Instruction following, reasoning, text generation	Scale, alignment via RL and feedback loops [5]
Diffusion Models	High-quality image synthesis and editing	Continuous gains in fidelity and controllable generation [2]
Multimodal Unified	Cross-modal grounding and multimodal outputs	Integration of modalities for richer interaction and tool use [1, 3]

2.3 Multimodal Emergence

A defining trend of 2025–2026 was the emergence of unified multimodal LLMs that integrate text, vision, audio, and video inputs. Models such as GPT-4o and its successors demonstrated that human-like interaction requires seamless cross-modal understanding [3]. Workplaces increasingly adopted multimodal AI for document processing, visual analytics, and interactive customer engagement.

Multimodal AI Applications in Practice

Enterprise Deployment Patterns (2025–2026):

- **Document Intelligence**: Vision-language models processing invoices, contracts, and reports
- **Customer Service**: Multimodal chatbots handling text, image, and audio queries simultaneously
- **Medical Imaging**: Unified models correlating clinical notes with radiology scans
- **Code Generation**: Models interpreting UI screenshots to generate functional code
- **Scientific Research**: Models processing papers, datasets, and experimental images together

3 Deep Learning and Neural Network Innovations

3.1 Self-Reflective Training Paradigms

Research in deep architectures emphasised multimodal alignment, efficient scaling, and self-reflective training approaches to improve generation and reasoning quality. A significant breakthrough was the development of Group-Relative Policy Optimisation (GRPO) combined with decompositional verifiable rewards, enabling models to decompose user intents and self-correct through structured feedback [6].

AlphaGRPO: Self-Reflective Multimodal Generation

The AlphaGRPO framework [6] introduced a novel training paradigm:

1. **Intent Decomposition**: Breaking complex user requests into verifiable sub-goals
2. **Reward Decomposition**: Assigning granular rewards to each sub-goal achievement
3. **Self-Reflection**: Iterative self-correction based on reward signals
4. **Policy Optimisation**: GRPO updates stabilising training across multimodal tasks

Result: Empirical gains on multimodal generation and editing benchmarks, demonstrating more interpretable feedback than scalar reward approaches.

3.2 Efficient Scaling Strategies

The field converged on several efficiency techniques to manage the cost of training and deploying large models:

- **Parameter-Efficient Fine-Tuning (PEFT)**: LoRA, QLoRA, and adapter methods enabling task adaptation with $< 1\%$ of full model parameters
- **Quantisation**: INT4/INT8 quantisation reducing memory footprint by 4–8× with minimal accuracy loss
- **Knowledge Distillation**: Transferring capabilities from large teacher models to smaller, deployable student models
- **Flash Attention**: Memory-efficient attention computation enabling longer context windows
- **Mixture-of-Experts**: Sparse activation ensuring only relevant expert networks compute for each token

3.3 Hybrid Optimisation and Verifier Guidance

Reinforcement-derived methods were adapted to multimodal goals to provide more interpretable feedback signals than scalar rewards alone [6, 5]. This hybrid approach—combining supervised learning, reinforcement feedback, and verifier-guided correction—emerged as the dominant paradigm for training state-of-the-art models.

4 AI Ethics, Governance, and Responsible AI

4.1 Policy-Level Urgency

AI ethics and governance were elevated to strategic priorities in 2025–2026, prompting landmark policy reports, international governance frameworks, and calls for rigorous evaluation beyond static benchmarks. The US President's Council of Advisors on Science and Technology (PCAST) published a landmark report advocating for boosted AI research resources while prioritising safety, oversight, and governance mechanisms to manage systemic risks from powerful models [7].

PCAST Recommendations (2024)

Key policy directives from the PCAST Report to the President [7]:

- Establish national AI safety testing infrastructure
- Mandate transparency and auditability for high-risk AI systems
- Increase federal investment in AI safety and alignment research
- Create international coordination mechanisms for AI governance
- Prioritise equitable access to AI tools and infrastructure

4.2 Responsible Deployment Imperatives

Reviews and trend analyses identified three pillars of responsible AI scaling [8]:

1. **Trustworthy Multimodality**: Ensuring cross-modal outputs are accurate, unbiased, and verifiable
2. **Verifiable Provenance**: Establishing clear data lineage and model decision traceability
3. **Energy-Aware Efficiency**: Reducing the carbon footprint of large-scale AI training and inference

4.3 Alignment Bottlenecks

Despite progress, alignment methods face persistent challenges [5]:

Alignment Challenges in 2025–2026

- **Reward Hacking**: Models exploiting reward function loopholes rather than learning intended behaviour
- **Scalable Feedback Collection**: Human annotation bottlenecks limiting RLHF at scale
- **Distribution Shift**: Models behaving unexpectedly on out-of-distribution inputs
- **Value Pluralism**: Difficulty encoding diverse human values into a single reward function
- **Evaluation Opacity**: Static benchmarks failing to capture real-world safety properties

4.4 Evaluation and Transparency

Multiple sources stressed the need for open evaluation ecosystems and verifiable data governance to reduce opacity and enable auditability of deployed systems [8, 1]. The emergence of "red-teaming" as a standard practice, alongside third-party auditing frameworks, represented meaningful progress toward accountable AI deployment.

5 AI in Healthcare and Medicine

5.1 LLMs in Clinical Workflows

Healthcare emerged as one of the highest-impact domains for AI in 2025–2026. Reviews documented growing use-cases for LLMs across clinical documentation, decision support, and patient interaction, while noting persistent concerns about privacy, bias, and rigorous

clinical validation [9, 10].

Clinical AI Applications 2025–2026

- **Clinical Documentation**: Automated SOAP note generation from physician-patient conversations
- **Diagnostic Support**: LLMs assisting in differential diagnosis from patient history and lab results
- **Drug Interaction Checking**: Real-time AI screening of prescription combinations
- **Radiology Reporting**: Vision-language models generating structured radiology reports
- **Patient Communication**: AI-powered patient education and discharge instruction generation
- **Clinical Trial Matching**: Automated patient-trial eligibility screening

5.2 Generative AI Opportunities and Caveats

Generative AI methods demonstrated potential in clinical trials, documentation, and medical education, but require rigorous evaluation, transparency, and equity safeguards before widespread clinical adoption [10]. Key concerns include:

- **Hallucination Risk**: LLMs generating plausible but clinically incorrect information
- **Demographic Bias**: Models performing differentially across patient populations
- **Regulatory Compliance**: Navigating FDA/EMA requirements for AI-assisted medical devices
- **Clinician Trust**: Building evidence-based confidence in AI recommendations

5.3 Federated Learning in Medical Imaging

Federated learning has been applied for privacy-preserving medical imaging, including skin cancer classification, enabling joint model training without exchanging raw patient data [11]. This approach addresses the fundamental tension between AI's need for large datasets and healthcare's strict privacy requirements.

Federated Learning Architecture in Healthcare

How Federated Learning Protects Patient Privacy:

1. **Local Training**: Each hospital trains a local model on its own patient data
2. **Model Aggregation**: Only model weights (not patient data) are shared with a central server
3. **Global Update**: Aggregated weights create an improved global model
4. **Distribution**: Updated global model is distributed back to all participating institutions

Key benefit: Patient data never leaves the hospital, satisfying HIPAA/GDPR requirements while enabling collaborative learning.

6 Federated Learning and Privacy-Preserving AI

6.1 Decentralised Machine Learning Paradigms

Reviews emphasise federated, edge, and decentralised paradigms as essential for privacy-preserving ML across industries, while noting infrastructure and standardisation gaps [12]. The field has matured from proof-of-concept to production deployments, particularly in:

- **Healthcare**: Multi-hospital model training without data sharing
- **Financial Services**: Fraud detection models trained across bank networks
- **Mobile Devices**: On-device personalisation (e.g., keyboard prediction, face recognition)
- **IoT Networks**: Edge intelligence for smart cities and industrial systems

6.2 Technical Advances in 2025–2026

Federated Learning Technical Progress

- **Differential Privacy**: Mathematical privacy guarantees through calibrated noise injection
- **Secure Aggregation**: Cryptographic protocols preventing server from seeing individual updates
- **Personalised Federated Learning**: Balancing global model performance with local adaptation
- **Cross-Silo vs Cross-Device**: Distinct architectures for enterprise vs. consumer deployments
- **Communication Efficiency**: Gradient compression reducing bandwidth requirements by 10–100×

6.3 Remaining Challenges

Despite progress, federated learning faces significant barriers to widespread adoption [12]:

1. **Heterogeneous Data**: Non-IID data distributions across participants degrading model quality
2. **System Heterogeneity**: Varying computational capabilities across participating devices
3. **Poisoning Attacks**: Malicious participants injecting corrupted model updates
4. **Standardisation**: Lack of interoperable frameworks across organisations
5. **Regulatory Clarity**: Evolving legal interpretations of federated data processing

7 Reinforcement Learning Advances

7.1 RL for LLM Alignment

Reinforcement Learning techniques—PPO, actor-critic, RLHF, DPO, and GRPO—remained central to aligning LLMs and improving stepwise reasoning in 2025–2026 [5]. A comprehensive technical survey documented the evolution of RL-based alignment:

RL Techniques for LLM Alignment: A Taxonomy

- **RLHF (Reinforcement Learning from Human Feedback)**: The dominant approach; trains a reward model from human preferences, then uses PPO to optimise the LLM against it
- **DPO (Direct Preference Optimisation)**: Eliminates the explicit reward model, directly optimising from preference pairs; more stable but less flexible
- **GRPO (Group-Relative Policy Optimisation)**: Computes advantages relative to a group of sampled responses; reduces variance and improves sample efficiency
- **Constitutional AI**: Uses AI feedback (rather than human) to scale alignment supervision
- **Process Reward Models (PRM)**: Provides step-level rewards for reasoning chains, improving mathematical and logical problem-solving

7.2 RL for Scientific Discovery

Beyond language model alignment, RL demonstrated breakthrough applications in scientific domains:

- **Protein Structure Prediction**: RL-guided search in conformational space
- **Materials Discovery**: RL agents navigating chemical space for novel material properties
- **Drug Design**: RL-based molecular generation optimising for multiple drug-likeness criteria simultaneously
- **Robotics**: Foundation models for robot control using RL fine-tuning

8 AI for Scientific Discovery

8.1 The AI Revolution in Chemistry

A landmark data-driven analysis covering over **310,000 journal articles and patents** reported strong AI adoption in biomedical and materials research, documenting an **eightfold increase in biochemistry patents** between 2015 and 2025 [13]. This quantitative evidence confirms AI's transformative impact on the pace of scientific discovery.

AI in Chemistry: Key Statistics [ ]

- **8×** increase in biochemistry patents driven by AI methods (2015–2025)
- **310,000+** journal articles and patents analysed
- Strong AI adoption in **biomedical** and **materials science** research
- Industrial translation of AI chemistry methods accelerating
- AI-driven materials discovery reducing experimental cycles from years to weeks

8.2 Autonomous Drug Discovery

Proposals for autonomous, LLM-driven drug-discovery frameworks combining knowledge graphs and planning demonstrated feasibility for scalable hypothesis generation and automated workflows in early-stage target discovery [14]. These systems represent a new paradigm:

Autonomous AI Drug Discovery Framework

Pipeline Components:

1. **Knowledge Graph Construction**: Integrating literature, databases, and experimental data
2. **Hypothesis Generation**: LLM-driven identification of novel target-disease relationships
3. **Virtual Screening**: AI-guided molecular docking and ADMET prediction
4. **Synthesis Planning**: Retrosynthetic AI predicting feasible synthesis routes
5. **Experimental Design**: Automated selection of validation experiments
6. **Result Interpretation**: LLMs interpreting experimental outcomes and updating hypotheses

8.3 AI in Materials Science

Beyond drug discovery, AI accelerated materials science through:

- **Crystal Structure Prediction**: Deep learning models predicting stable crystal structures from composition
- **Property Prediction**: Graph neural networks predicting mechanical, thermal, and electronic properties
- **Inverse Design**: Generative models proposing materials with target properties
- **High-Throughput Screening**: AI-guided experimental design reducing screening costs by orders of magnitude

9 AI in Cybersecurity

9.1 Generative AI as a Dual-Use Technology

The cybersecurity landscape in 2025–2026 was profoundly shaped by generative AI's dual-use nature. Advanced generative models simultaneously empowered defenders with sophisticated threat detection and enabled attackers with unprecedented attack capabilities [3].

Generative AI Cybersecurity Threats

Emerging Threat Vectors (2025–2026):

- **Deepfake Attacks**: Hyper-realistic synthetic media enabling CEO fraud and identity theft
- **AI-Generated Phishing**: LLMs crafting highly personalised, contextually accurate phishing emails
- **Automated Vulnerability Discovery**: AI-assisted code analysis identifying exploitable weaknesses at scale
- **Adversarial Attacks**: Imperceptible perturbations fooling AI-based security classifiers
- **Misinformation Campaigns**: Automated generation of coordinated disinformation at scale

9.2 AI-Powered Defence Mechanisms

Conversely, AI strengthened cybersecurity defences across multiple dimensions:

- **Anomaly Detection**: ML models identifying unusual network behaviour patterns in real-time
- **Threat Intelligence**: NLP systems extracting indicators of compromise from threat feeds
- **Automated Incident Response**: AI-driven SOAR platforms accelerating incident containment
- **Vulnerability Management**: AI prioritising patching based on exploitability and business impact
- **Identity Verification**: Behavioural biometrics and continuous authentication

9.3 Privacy and Data Protection

AI's data-intensive nature created new privacy challenges requiring technical and regulatory responses [3]:

Privacy-Preserving AI Techniques

- **Differential Privacy**: Adding calibrated noise to protect individual data points in training
- **Homomorphic Encryption**: Enabling computation on encrypted data without decryption
- **Synthetic Data Generation**: Creating privacy-safe training datasets that preserve statistical properties
- **Data Minimisation**: Training effective models with minimal personally identifiable information

10 Key Challenges and Future Directions

10.1 Cross-Cutting Challenges

Major recurring challenges across all AI/ML domains in 2025–2026 include [8, 1, 5, 13]:

The Seven Grand Challenges of AI in 2026

1. **Model Robustness**: Building AI systems that perform reliably across diverse, real-world conditions and adversarial inputs
2. **Interpretability and Explainability**: Opening the "black box" to enable human understanding, trust, and accountability
3. **Data Quality and Provenance**: Ensuring training data is accurate, unbiased, and legally compliant
4. **Scalable Alignment**: Developing alignment methods that scale with model capability without proportional human oversight costs
5. **Compute and Energy**: Reducing the environmental footprint of large-scale AI training and inference
6. **Equitable Access**: Preventing AI capability concentration in a small number of well-resourced organisations
7. **Evaluation Validity**: Creating benchmarks that genuinely measure real-world performance and safety properties

10.2 Promising Research Directions

Looking ahead, the most promising research directions identified across the literature include:

- **Neurosymbolic AI**: Combining neural learning with symbolic reasoning for more robust, interpretable systems
- **Causal AI**: Moving beyond correlation to causal reasoning for more reliable decision-making
- **Continual Learning**: Enabling models to learn from new data without catastrophic forgetting
- **AI-Human Collaboration**: Designing systems that augment rather than replace human expertise
- **Foundation Models for Science**: Domain-specific large models trained on scientific literature and data
- **Efficient AI**: Achieving GPT-4 level performance with dramatically reduced compute requirements
- **Trustworthy AI**: Formal verification methods for safety-critical AI systems

10.3 Implications for ICT Professionals

For ICT professionals navigating this landscape, the 2025–2026 insights point to several strategic imperatives:

Strategic Imperatives for ICT Professionals

- **AI Literacy**: Develop working knowledge of LLM capabilities, limitations, and governance requirements
- **Governance Frameworks**: Implement AI risk management aligned with NIST AI RMF, EU AI Act, and ISO/IEC 42001
- **Data Strategy**: Prioritise data quality, provenance, and privacy as foundations for trustworthy AI
- **Security Posture**: Update threat models to account for AI-powered attacks and deepfake risks
- **Responsible Deployment**: Establish ethics review processes for AI system procurement and development
- **Continuous Learning**: Maintain currency with the rapidly evolving AI research landscape

11 Conclusion

The 2025–2026 AI and Machine Learning landscape is characterised by extraordinary capability advances alongside growing recognition of the responsibilities that accompany them. Trillion-parameter language models, unified multimodal systems, and autonomous scientific discovery pipelines represent genuine technological breakthroughs. Simultaneously, the field is grappling seriously with alignment, governance, privacy, and equity challenges that will determine whether these capabilities translate into broadly beneficial outcomes.

For ICT professionals, researchers, and organisational leaders, the key message is clear: AI is no longer an emerging technology to be monitored from a distance—it is a present reality requiring active engagement, informed governance, and strategic integration. Those who invest in AI literacy, responsible deployment practices, and continuous learning will be best positioned to harness AI's transformative potential while managing its risks.

The research reviewed in this report provides a solid foundation for that engagement. From federated learning's promise for privacy-preserving healthcare to RL's advances in scientific discovery, from the dual-use cybersecurity challenge to the eight-fold growth in AI-driven patents, the evidence is clear: we are living through a pivotal moment in the history of intelligence itself.

"The question is not whether AI will transform your field, but whether you will shape that transformation or be shaped by it."

— **Atine John Bosco**, ICT Professional & Digital Transformation Advisor

References

References

[1] Schreiber, H., Ramirez, S., & Patel, A. (2025). *The Next Frontier in AI Research with Distributed and Multimodal Large Language Models*. TechRxiv. https://doi.org/10.36227/techrxiv.174345061.15598909/v1

[2] Tai, W. et al. (2024). *Recent Advances in Generative AI and Large Language Models: Current Status, Challenges, and Perspectives*. IEEE Transactions on Artificial Intelligence. https://doi.org/10.1109/tai.2024.3444742

[3] Ghadei, K., & Panigrahi, S. P. (2025). *The International Year of AI: Navigating Scope and Risk in 2025*. Journal of Extension Systems, 40(2). https://doi.org/10.48165/jes.2024.40.2.2

[4] Liu, J., Vats, A., & He, Z. (2025). *CS-PaperSum: A Large-Scale Dataset of AI-Generated Summaries for Scientific Papers*. arXiv preprint arXiv:2502.20582. https://doi.org/10.48550/arxiv.2502.20582

[5] Srivastava, S., & Aggarwal, V. (2025). *A Technical Survey of Reinforcement Learning Techniques for Large Language Models*. arXiv preprint arXiv:2507.04136. https://doi.org/10.48550/arxiv.2507.04136

[6] Anonymous. (2026). *AlphaGRPO: Unlocking Self-Reflective Multimodal Generation in UMMs via Decompositional Verifiable Reward*. arXiv preprint arXiv:2605.12495. https://arxiv.org/abs/2605.12495v1

[7] President's Council of Advisors on Science and Technology. (2024). *PCAST Report to the President on Supercharging Research: Harnessing Artificial Intelligence to Meet Global Challenges*. https://doi.org/10.2172/2481685

[8] Sahu, A., & Singh, P. (2025). *The Future of Machine Learning and Artificial Intelligence: Key Trends and Innovative Applications*. Journal of Informatics Education and Research, 5(1). https://doi.org/10.52783/jier.v5i1s.3460

[9] Khalid, A., Khalid, A., & Khalid, U. (2024). *The Role of Language Models in Modern Healthcare: A Comprehensive Review*. arXiv preprint arXiv:2409.16860. https://doi.org/10.48550/arxiv.2409.16860

[10] Shanmugam, D., Agrawal, M., Movva, R., et al. (2024). *Generative AI in Medicine*. arXiv preprint arXiv:2412.10337. https://doi.org/10.48550/arxiv.2412.10337

[11] Anonymous. (2026). *Federated Learning for Privacy-Preserving Skin Cancer Classification Using Deep Neural Networks*. Frontiers in Oncology. https://doi.org/10.3389/fonc.2026.1720748

[12] Saif, S., Islam, M. J., Jahangir, M. Z. B., et al. (2025). *A Comprehensive Review on Understanding the Decentralized and Collaborative Approach in Machine Learning*. arXiv preprint arXiv:2503.09833. https://doi.org/10.48550/arxiv.2503.09833

[13] Baranwal, B. S., Jones, L. A., Maind, A., et al. (2025). *The AI Revolution in Chemistry: Shaping the Future of Materials and Biomedical Sciences*. ChemRxiv. https://doi.org/10.26434/chemrxiv-2025-8lrf0

[14] Anonymous. (2024). *A Framework for Autonomous AI-Driven Drug Discovery*. bioRxiv. https://doi.org/10.1101/2024.12.17.629024

About the Author

Atine John Bosco

Atine John Bosco is a Chartered Information Scientist, PhD Candidate in Computer Science, and Digital Transformation Advisor with over a decade of experience in ICT leadership, cybersecurity, and emerging technology strategy.

Professional Credentials:

- Member of the British Computer Society (MBCS)
- Certified Information Systems Auditor (CISA)
- Google Cloud Digital Leader
- ITIL 4 Practitioner
- ACM Member
- SIM Member

Research Interests: Artificial Intelligence, Machine Learning, Cybersecurity, Cloud Computing, Digital Transformation, ICT Policy and Governance.

Contact:

www.atine.info

linkedin.com/in/atine

© 2026 Atine John Bosco · All Rights Reserved
Published at www.atine.info
This document is for educational and informational purposes.