ITU Journal

Future and evolving technologies

FREE I FAST I FOR ALL

Volume 6, Issue 3, September 2025

ISSN: 2616-8375

ITU Journal
Future and evolving technologies

FREE I FAST I FOR ALL

The ITU Journal on Future and Evolving Technologies (ITU J-FET) is an international journal providing complete coverage of all communications and networking paradigms, free of charge for both readers and authors.

The ITU Journal considers yet-to-be-published papers addressing fundamental and applied research. It shares new techniques and concepts, analyses and tutorials, as well as learning from experiments and physical and simulated testbeds. It also discusses the implications of the latest research results for policy and regulation, legal frameworks, the economy and society. This publication builds bridges between disciplines, connects theory with application, and stimulates international dialogue. Its interdisciplinary approach reflects ITU's comprehensive field of interest and explores the convergence of ICT with other disciplines.

The ITU Journal welcomes submissions at any time, and on any topic within its scope.

Publication rights

SUGGESTED CITATION:

ITU Journal on Future and Evolving Technologies, Volume 6, Issue 3, September 2025

ADDITIONAL INFORMATION
Please visit the ITU J-FET website at:
https://www.itu.int/en/journal/j-fet/Pages/default.aspx.

Inquiries should be addressed to Alessia Magliarditi at: journal@itu.int.

Special issue on "Privacy and security challenges of generative AI"

Editorial

Generative Artificial Intelligence (GenAI) is ubiquitous across industries and societal domains. Its extraordinary ability to extract, process and expand data, information and knowledge makes it central to the digital economy. GenAI technologies hold the promise of addressing escalating demands in cost, power, capacity, coverage, latency, efficiency, flexibility, compatibility and quality of experience, factors that define the fabric of our digital lives.

Yet as GenAI application systems proliferate, privacy and security have assumed a pivotal role in their rapid development and massive deployment. Unauthorized use of data and model parameters, leakage of proprietary or classified information, and vulnerabilities during both training and inference phases represent critical risks. Private and secure generative AI is thus essential to ensuring compliance with legal frameworks, safeguarding individual rights and sustaining public trust.

This special issue was conceived to catalyse and steer advances in novel systems, architectures and safeguards that enable private and secure generative AI. It brings together contributions across algorithm design, cryptographic protection, system architectures, deployment strategies, governance frameworks and hardware implementations. The aim is to foster collaboration between scientists, engineers, manufacturers, software developers, policymakers and broader stakeholders in shaping the secure future of generative AI.

Reflecting the themes outlined in the call for papers, the seven papers included in this issue span three interconnected domains:

a. algorithms, architectures and applications

b. deployment, standardization and development

c. information and signal processing.

The first paper, *"**Exploring the benefits of differentially private pre-training and fine-tuning for table transformers**,"* situates itself at the intersection of private learning and parameter-efficient architectures. By combining table transformers with differential privacy and fine-tuning approaches such as adapters, LoRA and prompt tuning, it demonstrates how accuracy, efficiency and privacy can be jointly optimized, which is a critical advance in deploying GenAI securely on sensitive tabular datasets.

Addressing adversarial resilience, *"**The SkipSponge attack: Sponge weight poisoning of deep neural networks**"* uncovers a new form of sponge attack capable of stealthily increasing the energy consumption of neural networks, particularly GANs and autoencoders. Requiring minimal access to training data, SkipSponge exposes a systemic vulnerability that falls squarely under the security challenges this special issue seeks to highlight, urging the design of defences beyond traditional poisoning countermeasures.

Contributing to the advances in the role of encryption for securing GenAI, *"**A practical homomorphic encryption approach for GDPR-compliant machine learning full training protocol**"* introduces a selective encryption framework that strategically secures only privacy-critical layers. This approach delivers a thousand-fold improvement in training efficiency over fully encrypted models while maintaining compliance with privacy regulations, marking a breakthrough in practical homomorphic encryption deployment.

From a governance and policy perspective, *"**Strengthening AI governance: International policy frameworks, security challenges and ethical AI deployment**"* expands the lens to a global scale. Drawing on the Global Index for Responsible AI dataset, it reveals stark disparities in national

preparedness to handle GenAI's privacy and security risks. By emphasizing international cooperation and culturally-sensitive frameworks, this contribution resonates with the call for papers' emphasis on standardization and regulatory development.

Complementing the fourth paper above, *"**Challenges with handling keys for secure AI**"* addresses a crucial but underexplored issue: key management in encrypted AI systems. By analysing gaps in NIST KMS standards and proposing a hierarchical key management system, the paper illuminates technical bottlenecks and suggests solutions informed by IBM's He4Cloud design.

On the hardware front, *"**Private LLM technology: Security-layer definitions and optimal silicon solutions**"* explores the convergence of algorithmic design and silicon efficiency. The authors categorize private LLM requirements into security layers and introduce Cornami hardware as an alternative to GPUs and ASICs, achieving better trade-offs in latency, energy and cost. This work embodies the "co-design of algorithm and hardware" envisioned by this special issue.

The final contribution, *"**Adaptive hybrid convolutional neural network-autoencoder framework for backdoor detection in GenAI-driven semantic communication systems**,"* addresses adversarial robustness in semantic communication systems, where GenAI enables meaning-based transmission rather than raw data. By combining CNNs with adaptive autoencoders, the framework detects deeply embedded backdoors without altering model architectures or sacrificing inference quality, advancing both signal processing and security goals.

Together, these papers illustrate the multilayered nature of privacy and security in generative AI. They span differentially private algorithms, homomorphic encryption, adversarial resilience, governance frameworks, hardware optimization and semantic communication security. Each contribution not only advances the current state-of-the-art situation but they also respond directly to the research directions highlighted in the call for papers: encryption and decryption techniques, algorithm-hardware co-design, secure federated and differential learning, FHE development libraries and private LLM innovations.

As we conclude the introduction to this special issue, we would like to thank all authors for their valuable contributions, and we express our sincere gratitude to the reviewers for their timely and insightful comments on submitted papers. We hope that the content of this special issue is informative and useful across technology, standardization and implementation in addressing the privacy and security challenges of generative AI.

The Guest Editors

Fa-Long Luo, Cornami, USA (Lead Guest Editor)

Rosario Cammarota, Intel Labs, USA

Paul Master, Cornami, USA

Nir Drucker, IBM Europe, Israel

Donghoon Yoo, Desilo, Korea

Konstantinos Plataniotis, University of Toronto, Canada

EDITORIAL BOARD

TABLE OF CONTENTS