Encryption breakthrough lays groundwork for privacy-preserving AI models

March 26 2025

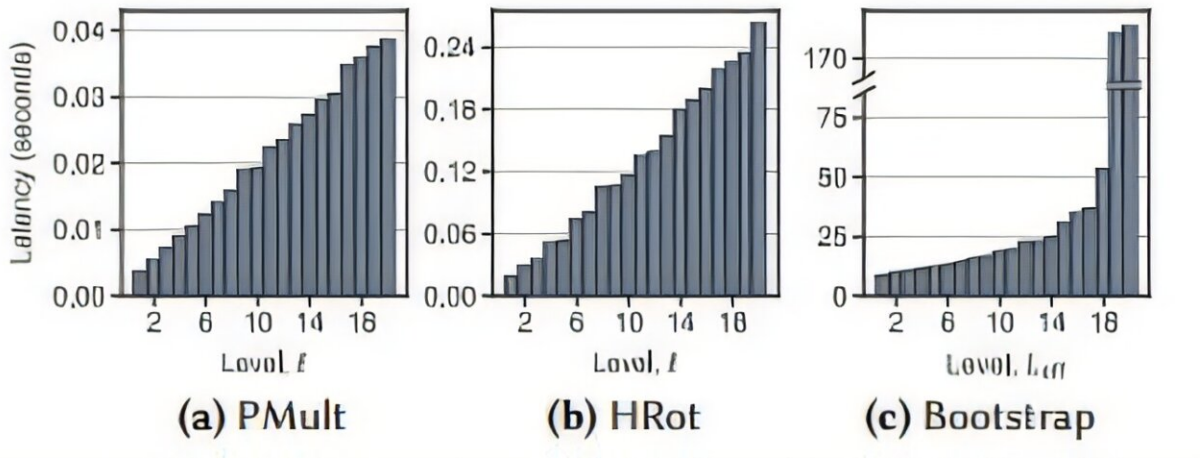

The latencies of key homomorphic operations as a function of ciphertext level (ℓ) for ring degree $N = 2^{16}$ and $\Delta \approx 2^{40}$. Setting ℓ_{eff} too low would require many low-latency bootstraps, while setting it too high would result in fewer but higher-latency bootstraps. We set $\ell_{\text{eff}} = 10$ in order to balance bootstrap latency with having a reasonable number of remaining levels for computation. Credit: *arXiv* (2023). DOI: 10.48550/arxiv.2311.03470

In an era where data privacy concerns loom large, a new approach in artificial intelligence (AI) could reshape how sensitive information is processed.

Researchers Austin Ebel and Karthik Garimella, Ph.D. students, and Assistant Professor of Electrical and Computer Engineering Brandon Reagen have introduced Orion, a novel framework that brings fully homomorphic encryption (FHE) to deep learning—allowing AI models to practically and efficiently operate directly on encrypted data without needing to decrypt it first.

The implications of this advancement, published on the *arXiv* preprint server and scheduled to be presented at the 2025 ACM International Conference on Architectural Support for Programming Languages and Operating Systems, are profound.

FHE has long been considered the holy grail of cryptography. Unlike traditional encryption, which protects data only when it is at rest or in transit, FHE allows computations to be performed on encrypted data without ever decrypting it. However, despite its promise, implementing deep learning models with FHE has been notoriously difficult due to the immense computational overhead and the technical hurdles in adapting neural networks to FHE's bespoke programming model.

"Whenever you use online services, there are machine learning models operating in the background, collecting both your inputs and outputs," says Garimella. "That compromises user privacy. Our goal is to bring FHE into the mainstream, and allow users to continue using the services they rely on every day without releasing their personal, private data."

Orion tackles these challenges head-on with an automated framework that seamlessly converts deep learning models written in PyTorch into efficient FHE programs. It does so by introducing a novel method to optimize how encrypted data is structured, significantly reducing computational overhead. The framework also streamlines encryption-related processes, making it easier to manage accumulated noise and execute deep learning computations efficiently.

By employing these techniques, Orion achieves a 2.38x speedup over existing state-of-the-art methods on ResNet-20, a common benchmark model used in FHE deep learning research that is comparatively small. But perhaps most impressively, Orion enables computations on much larger networks than previously possible. The researchers demonstrated the first-ever high-resolution FHE object detection using YOLO-v1—a deep learning model with 139 million parameters, roughly 500 times larger than ResNet-20—showcasing Orion's ability to handle real-world AI workloads.

The code the team produced is lightweight and could be used by anyone with a basic understanding of computer science. Not only would this help in increasing the efficiency of the computations; it also makes it easily deployable across industries.

"There has been an incredible barrier to entry for people who don't want to spend months to years learning the ins and outs," says Ebel. "With Orion, that barrier to entry is now almost non-existent."

The development of Orion marks a critical milestone in bridging the gap between FHE and practical deep learning applications. With this framework, industries reliant on privacy—such as health care, finance, and cybersecurity—could leverage AI without exposing sensitive user data.

"Take online advertising," says Reagen, who is also a member of the NYU Center for Cybersecurity. "If you want to process an individual's information in order to serve them targeted ads using neural networks, this allows service providers to analyze that data while keeping it totally confidential. For the marketers and the public, that's a win-win scenario."

While challenges remain in making FHE fully practical at scale, Orion

brings the technology closer to widespread adoption. The research team has open-sourced the project, making it accessible to developers and researchers worldwide.

As AI continues to integrate deeper into daily life, privacy-preserving techniques like Orion could redefine the balance between innovation and security—ensuring that smarter algorithms don't come at the cost of user privacy.

More information: Austin Ebel et al, Orion: A Fully Homomorphic Encryption Framework for Deep Learning, *arXiv* (2023). DOI: 10.48550/arxiv.2311.03470

Conference: www.asplos-conference.org/asplos2025/

Provided by NYU Tandon School of Engineering