Chapter 1
Privacy-Preserving Deep Learning for Healthcare Using Homomorphic Encryption

Parth Nagar
http://orcid.org/0009-0002-4027-3911
Sri Sathya Sai Institute of Higher Learning, India

Srinath M. S.
http://orcid.org/0000-0003-3275-6999
Sri Sathya Sai Institute of Higher Learning, India

ABSTRACT

Growing interest in cloud-based AI for clinical decision support raises critical patient privacy concerns, particularly under India's DPDPA. Homomorphic encryption (HE) offers a robust solution, enabling cloud servers to evaluate deep learning models on encrypted data and return encrypted results without accessing the plaintext. However, implementing HE requires navigating a strict privacy-efficiency-accuracy trilemma in deep learning architectures. This chapter provides a practical guide to designing neural network architectures that operate effectively under these cryptographic constraints. Per the authors, design choices such as polynomial activation function approximations, shallow-wide architectures, and quantization-aware training are discussed in detail. Through three healthcare case studies, the authors demonstrate that encrypted models can achieve accuracies within 1-3% of their plaintext counterparts. Finally, this chapter examines the integration of HE with federated learning and generative AI, concluding with a comprehensive implementation roadmap for healthcare teams.

DOI: 10.4018/979-8-3373-7426-0.ch001

1. INTRODUCTION

"Fully homomorphic encryption has long been regarded as cryptography's prized 'holy grail'-extremely useful yet rather elusive." - Vinod Vaikuntanathan, Computing Blindfolded, FOCS 2011 (Vaikuntanathan, 2011)

The increasing reliance of modern healthcare on artificial intelligence for diagnosis, risk assessment, and treatment decisions is in conflict with the strict guidelines imposed by privacy laws across the globe, such as the Digital Personal Data Protection Act (DPDP Act), 2023, of India, the GDPR of the European Union, and HIPAA in the United States of America, regarding the use of patient data by third-party service providers. This is because the best AI algorithms are run on cloud servers provided by technology companies, but the data cannot be allowed to leave the hospital premises without adequate privacy assurances. The following hypothetical scenario illustrates this tension concretely.

Dr. Priya Sharma runs a community hospital in rural Maharashtra. A patient named Rajesh was admitted with complaints of regular chest pain. Dr. Sharma prescribed a chest radiograph. The result revealed some alarming patterns - a suspicious opacity that might be suggestive of early malignancy. A technology firm in Bengaluru provides an AI-based diagnostic tool with expert-level performance on exactly this kind of image. However, Dr. Sharma does not have the freedom to upload Rajesh's image.

The Indian Digital Personal Data Protection Act (DPDP Act), 2023 (Ministry of Law and Justice, 2023) considers health data to be sensitive and requires explicit consent, purpose limitation, and institutional accountability for any processing that may be done by third parties (Gupta et al., 2025). Standard transport-layer encryption (TLS) protects Rajesh's scan as it is transmitted across the network - but as soon as it hits the server in Bengaluru, the GPU has to decrypt each pixel to perform inference. The scan is now visible in plaintext form on another person's computer. Removing metadata is insufficient to anonymize the image, as chest radiographs can be identified based on anatomical differences that are unique to each person (Yin et al., 2024). Dr. Sharma is in a bind. The technology is available, but the privacy hurdle is what keeps her from using it.

This is not just a hypothetical scenario. This is what healthcare systems around the world need to address - in district hospitals in India, radiology centers in Europe, and healthcare systems in America. It is advanced AI models that are what lie behind cloud APIs, and it is privacy regulations and ethics that dictate that the data cannot be shared in order to use them. This is not a matter of capability. It is a matter of needing to process private patient data without risking its exposure to third-party computation services.