DIFFERENTIAL LEARNING IN HEALTHCARE

Harshvardhan Aditya Harine M

ABSTRACT

In recent years, machine learning has opened up many opportunities in the healthcare industry. The ability of machine learning and neural network models to find subtle patterns in data has allowed scientists to make advances in finding, diagnosing and curing diseases. However, using patient data as-is comes with the risk of exposing confidential patient information.

Differential privacy is a standard method that enables privacy protection for such models. It requires noise to be added to the database in order to maintain privacy. In this project, we seek to find a differential privacy technique for training and testing neural networks in the context of patient data. The aim is to preserve model performance as much as possible while protecting sensitive information.

INTRODUCTION

Big strides have been made in the healthcare domain thanks to the introduction of machine learning. Computer vision has allowed faster analysis of complex medical images, providing secondary opinion to physicians and detecting abnormalities that may not be visible to the human eye. Even natural language concepts have found uses in the biomedical domain- this includes putting together huge chunks of patient data, combining them with patient history and providing insights on related or future risks. AI/ML tools and algorithms- especially LLMs- have played a big role in making genomic data usable, which has provided faster and more reliable ways of testing whether certain gene therapies would work, predictions of possible mutations of genes as well as new gene therapy methods.

Not only is patient data used in developing machine learning models, it’s also used in Internet of Medical Things (IoMT) objects. These devices provide real-time health feedback (akin to doctor’s advice) based on the wearer’s medical history. For the feedback to be accurate, the devices have to rely heavily on data which may be sensitive.

In such an era where health information- from patient history to genomic sequence- is actively being digitised and used, it has become necessary to safeguard sensitive patient data while preserving meaningful information. Even for datasets with millions of entries, a series of carefully-selected queries on publicly available statistics can expose a majority of the dataset.

The Fundamental Law of Information Recovery states that privacy cannot be protected without introducing some amount of noise. If too much focus is directed towards information gain, privacy is bound to be compromised whereas adding too much noise will affect model performance.

Differential privacy (DP) aims to provide a balance between individual privacy and information loss through adding noise to the dataset.

Extensive work has been done on maintaining this privacy-information balance by modifying how the noise is added. Recent works will show that changing the noise on the basis of some parameter (adaptive DP) shows more favourable results as compared to the traditional DP where noise is constant. The basis on which noise is generated is being experimented on, some of the more successful examples being entropy theory [1] and task-specificity [2][3].

TOOLS USED

Python, Tensorflow, diffprivlib

[1] Zhang, X., Yang, F., Guo, Y., Yu, H., Wang, Z., & Zhang, Q. (2023, January 8). Adaptive Differential Privacy Mechanism Based on Entropy Theory for Preserving Deep Neural Networks. Mathematics, 11(2), 330. https://doi.org/10.3390/math11020330

[2] Utaliyeva, A., Shin, J., & Choi, Y. H. (2023, February 10). Task-Specific Adaptive Differential Privacy Method for Structured Data. Sensors, 23(4), 1980. https://doi.org/10.3390/s23041980

[3] Ziller, A., Usynin, D., Braren, R., Makowski, M., Rueckert, D., & Kaissis, G. (2021, June 29). Medical imaging deep learning with differential privacy. Scientific Reports, 11(1). https://doi.org/10.1038/s41598-021-93030-0