

Arenberg Doctoral School of Science, Engineering & Technology
Faculty of Engineering Science

Federated Learning for Privacy-Preserving Healthcare

Thomas Peeters

Supervisors:

Prof. dr. ir. Elke Martens

Prof. dr. Ruben Dewaele (co-
supervisor)

Dissertation presented in partial fulfil-
ment of the requirements for the degree
of Doctor of Engineering Science (PhD):
Computer Science

June 2026

Federated Learning for Privacy-Preserving Healthcare

Training clinical models across hospitals without sharing patient data

Thomas Peeters

Examination committee:

Prof. dr. ir. Bart Lemmens, chair

Prof. dr. ir. Elke Martens, supervisor

Prof. dr. Ruben Dewaele, co-supervisor

Prof. dr. Hanne Goossens

Prof. dr. Sofia Marques, Vertex University
Hospital, Portugal

Dissertation presented in partial fulfilment of the requirements for the degree of Doctor of Engineering Science (PhD): Computer Science

June 2026

Unofficial template, not affiliated with or endorsed by KU Leuven. The author, committee and results in this sample are fictional. Delete this line before you submit.

© 2026 KU Leuven, Faculty of Engineering Science

Uitgegeven in eigen beheer, Thomas Peeters, Kasteelpark Arenberg, B-3001 Leuven (Belgium)

Alle rechten voorbehouden. Niets uit deze uitgave mag worden vermenigvuldigd en/of openbaar gemaakt worden door middel van druk, fotokopie, microfilm, elektronisch of op welke andere wijze ook zonder voorafgaande schriftelijke toestemming van de uitgever.

All rights reserved. No part of the publication may be reproduced in any form by print, photoprint, microfilm, electronic or any other means without written permission from the publisher.

Preface

I thank my supervisors for their guidance, the members of my examination committee for their careful reading, and my colleagues in the research group for their support and friendship.

Abstract

Clinical prediction models improve with data from many hospitals, but privacy law and patient trust make it difficult to pool records. Federated learning trains a shared model by exchanging parameter updates instead of data. This dissertation studies federated learning for tabular clinical data, where hospitals differ in case mix, coding practice and size. It proposes a reweighting of client updates that accounts for differences in outcome prevalence, and combines it with differentially private aggregation. Across eleven simulated hospitals, the method recovers 96% of the discrimination of a model trained on pooled data while guaranteeing $(\epsilon, \delta) = (3, 10^{-5})$ differential privacy for every patient.

Popular Abstract

Computer models that predict, for example, which patients risk complications work better when they learn from many hospitals. Hospitals cannot simply share patient records. This thesis shows how models can learn from many hospitals while the records stay where they are: each hospital trains the model locally and only shares mathematical updates, which are further protected so that no individual patient can be recognized. The resulting models are almost as good as models that see all the data at once.

Beknopte samenvatting

Klinische voorspellingsmodellen worden beter met gegevens uit veel ziekenhuizen, maar privacywetgeving maakt het samenvoegen van dossiers moeilijk. Dit proefschrift bestudeert gefedereerd leren voor tabellarische klinische gegevens en stelt een herweging van de bijdragen van ziekenhuizen voor, gecombineerd met differentieel private aggregatie. Over elf gesimuleerde ziekenhuizen haalt de methode 96% van de discriminatie van een model dat op alle gegevens samen is getraind.

Populaire samenvatting

Computersmodellen die voorspellen welke patiënten risico lopen, werken beter als ze leren uit veel ziekenhuizen. Dit proefschrift toont hoe dat kan zonder dat patiëntgegevens het ziekenhuis verlaten: elk ziekenhuis traint het model lokaal en deelt alleen beschermde wiskundige aanpassingen.

List of Abbreviations

AUROC	area under the receiver operating characteristic curve
DP	differential privacy
FL	federated learning

List of Symbols

ϵ, δ	privacy parameters
K	number of hospitals
w_k	weight of hospital k

Contents

Preface	i
Abstract	iii
Popular Abstract	v
Beknopte samenvatting	vii
Populaire samenvatting	ix
List of Abbreviations	xi
List of Symbols	xiii
1 Introduction	1
2 Background	3
3 Prevalence-Aware Private Aggregation	5
4 Results	7
5 Conclusion	9
A Privacy Accounting	11

List of Figures

3.1 One round of private, prevalence-aware federated training. 5

List of Tables

4.1 Discrimination (AUROC) on held-out patients. 7

Chapter 1

Introduction

Hospitals collect detailed records that could train models to predict complications, readmissions and adverse drug events. A model trained in one hospital often fails in another, because patients, practice and coding differ. Pooling data would help, but patient records are among the most sensitive data there are. Federated learning offers a middle way [2]: the model travels, the data stay put.

Chapter 2

Background

In federated averaging, K clients each train on local data and the server combines their parameters,

$$\theta^{t+1} = \sum_{k=1}^K w_k \theta_k^{t+1}, \quad w_k = \frac{n_k}{\sum_j n_j}, \quad (2.1)$$

where n_k is the number of examples at client k [2]. Differential privacy bounds how much any single patient can change the output [1].

Chapter 3

Prevalence-Aware Private Aggregation

Weighting clients by size alone lets large hospitals with low outcome prevalence dominate the model. We instead weight each client by its effective number of positive and negative cases, and add calibrated Gaussian noise to the clipped updates before aggregation so that the whole training run satisfies differential privacy.




Figure 3.1: One round of private, prevalence-aware federated training.

Chapter 4

Results

Table 4.1 compares discrimination on held-out patients from all eleven hospitals.

Table 4.1: Discrimination (AUROC) on held-out patients.		
Training	AUROC	Relative to pooled
Pooled data (not private)	0.842	1.00
Local models only	0.731	0.87
Federated averaging, private	0.779	0.93
Prevalence-aware, private	0.808	0.96

Chapter 5

Conclusion

Federated learning can bring clinical models close to what pooled data would give, with a formal privacy guarantee, provided that aggregation accounts for how hospitals differ.

Appendix A

Privacy Accounting

Privacy loss was tracked with the moments accountant over 200 rounds with a sampling rate of 0.1 and a noise multiplier of 1.1, giving $\epsilon = 3$ at $\delta = 10^{-5}$.

Bibliography

- [1] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4): 211–407, 2014.
- [2] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pages 1273–1282, 2017.