European Heart Journal - Digital Health (2026) 7 (Suppl 1)

e-Cardiology/Digital Health – Artificial Intelligence, Data Synthesis

CardioSynth: privacy-preserving generation and evaluation of multimodal synthetic data for cardiovascular research

K. Kechagias[1], M. Miranda[2], M. Garofalakis[1], E. Morley-Fletcher[3], Y. Ioannidis[1], S. Bratieres[2], M.J. Boonstra[4], C. Capelli[3]

[1]Athena Research and Innovation Center, Athens, Greece
[2]Translated, Rome, Italy
[3]Lynkeus, Rome, Italy
[4]Amsterdam University Medical Centre (AUMC), Amsterdam, Netherlands (The)

Funding Acknowledgements: Type of funding sources: Public grant(s) – EU funding. Main funding source(s): DataTools4Heart is funded by the European Union's Horizon Europe Framework Under Grant Agreement No. 101057849

Background: Secondary use of cardiology data is vital for trend analysis and data-driven healthcare innovation, but privacy and regulations limit access. Synthetic data offers a solution, yet cardiology lacks a standardized, privacy-preserving framework for generating high-quality, multimodal datasets.
Purpose: This study presents a framework for generating multimodal, differentially private (DP) synthetic data to enable collaborative research under diverse privacy constraints. The data will form part of CardioSynth, a European dataset for privacy-preserving AI in cardiology.
Methods: A secure SD generation framework for cardiology was deployed in eight European sites using real-world multimodal data—structured, images, and unstructured—from both local hospital systems and the publicly available MIMIC-IV. For structured data, we evaluated Probabilistic Graphical Models and Deep Generative Models; the latter were also applied to imaging data. Both model families were trained using publicly available data and clinical data from participating hospitals. Centralized and federated configurations were tested. Fidelity, utility, and formal privacy were assessed under inferential privacy budgets $\epsilon \in \{0.1, 1, 5, 10, 40\}$, enabling private cardiovascular risk modeling. For unstructured data, the DP-In Context Learning (ICL) approach was adopted. Private datasets were partitioned into multiple subsets, generating ICL prompts for an LLM, whose outputs were aggregated using Embedding Space Aggregation to preserve both privacy and semantic meaning. Model fine-tuning was not required. The entire pipeline is available as a web app. Generated notes were evaluated by 19 experts who were presented with a subsample of 5 generated samples. Evaluation focused on data coherence, usability, language, and style to assess utility. For all these features, the reviewers were asked to rate the texts on a scale from 1 to 5.
Results: PrivBayes preserved marginals of structured SD (<5%) but had limited fidelity (AUROC = 0.71). DP-GAN improved structure (correlation error < 0.06; AUROC = 0.82; MIA ≤ 1.2%). Federated DP-GAN scaled to 9 hospitals, converged (10–27 rounds), and retained utility (AUROC = 0.79 at $\epsilon = 2$). As expected, higher ϵ resulted in better fidelity but lower levels of privacy. DP-GAN enabled imaging synthesis under $\epsilon \leq 5$.
The outputs for unstructured generation were successfully generated using Llama3.1. The clinical evaluation showed good coherence and generally actionable generations. The average scores for the main text features across all samples and voters are in the range between 3.2 and 3.92. Reviewers found the generated texts realistic.
Conclusions: A framework was established to generate multimodal cardiology SD from clinically available sources. This work lays the groundwork for creating CardioSynth, a public comprehensive SD repository that supports collaborative research while accommodating diverse privacy requirements.

Downloaded from https://academic.oup.com/ehjdh/article/7/Supplement_1/ztaf143.026/8423000 by guest on 13 May 2026

CardioSynth: Cardiology Synth Data

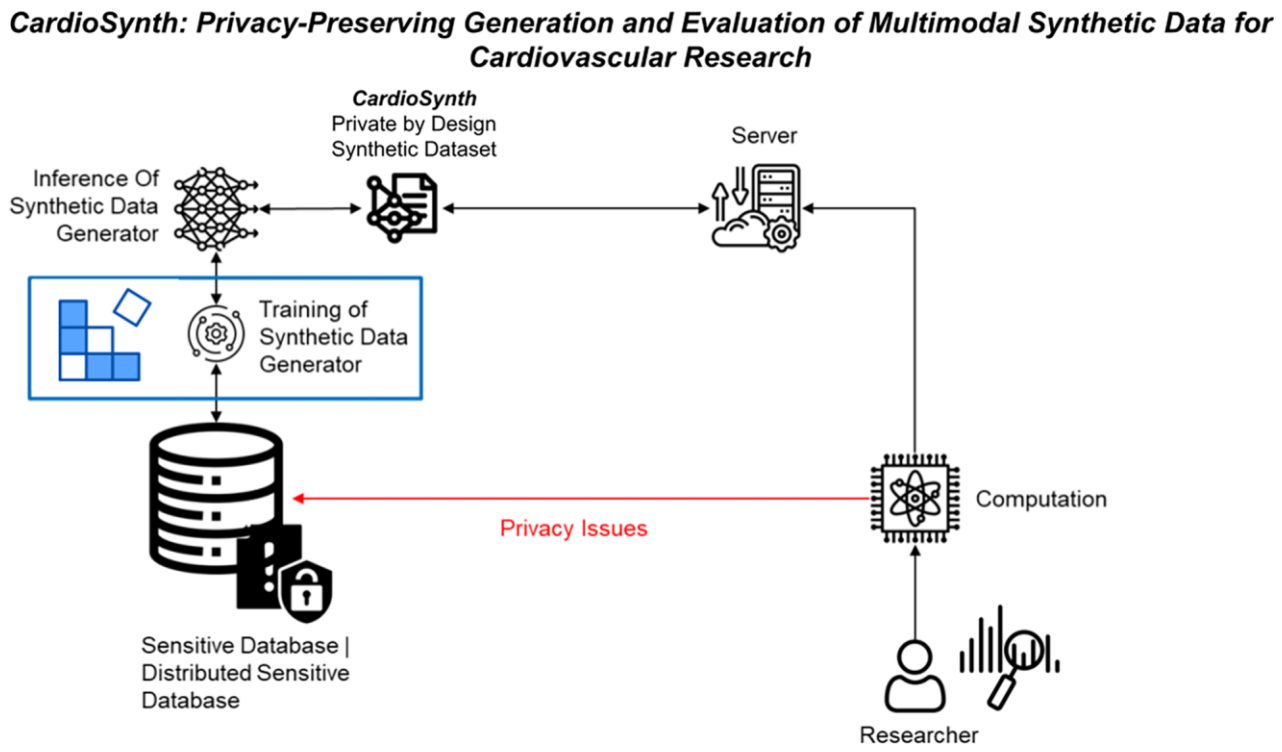

CardioSynth: Synthetic Data Evaluation

Comparison of one model from each family—PrivBayes (PGMs) and DP-GAN (DL) under varying ε.

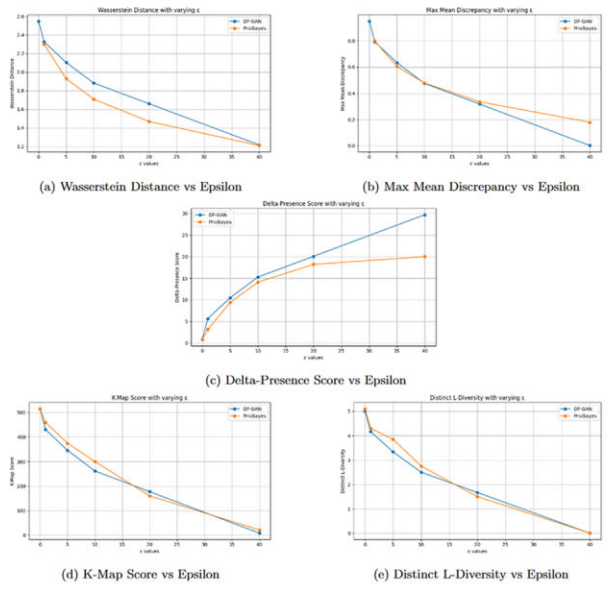

(a) Wasserstein Distance vs Epsilon (b) Max Mean Discrepancy vs Epsilon

(c) Delta-Presence Score vs Epsilon

(d) K-Map Score vs Epsilon (e) Distinct L-Diversity vs Epsilon

Ratings for unstructured generation samples, averaged across all voters and 5 samples.

Internal Coherence	3.34
Coherence with labels	3.2
Language used	3.50
Timeline	3.48
Clinical measurements	3.92
Medications and dosages	3.89

Downloaded from https://academic.oup.com/ehjdh/article/7/Supplement_1/ztaf143.026/8423000 by guest on 13 May 2026