Privacy-preserving Anonymization of FHIR healthcare data

Wenhui Yang[1], Lena Wiese[1,2]

[1]Fraunhofer Institute for Toxicology and Experimental Medicine, Nikolai-Fuchs-Str.1, Hannover, 30625, Germany
[2]Institute of Computer Science Goethe University Frankfurt, Robert-Mayer-Str.10, Frankfurt am Main, 60325, Germany

Abstract
Healthcare organizations are increasingly using Electronic Health Records (EHRs) to improve patient care and treatment efficacy. The introduction of the FHIR standard by HL7 aimed to handle varied and unstructured healthcare data formats using a resource-based approach. However, FHIR datasets contain Personally Identifiable Information (PII), posing privacy challenges under regulations like GDPR and HIPAA. This paper proposes a $k^{(n,t)}$-anonymity privacy model for FHIR datasets, utilizing a multidimensional anonymization method to safeguard patient privacy and maintain data structure integrity. This approach prevents the re-identification of individual records while handling FHIR's complex data structure.

Keywords
Privacy-preserving, Anonymization, FHIR, Healthcare data

1. Introduction

In recent years, healthcare organizations have increasingly relied on Electronic Health Records (EHRs) to collect patients' health data, aiming to enhance patient care, improve treatment efficacy, and ensure more accurate diagnoses [1]. To address the challenges of handling varied or unstructured data formats, the Health Level Seven International (HL7) healthcare standards organization introduced the Fast Healthcare Interoperability Resources (FHIR) standard in 2011. FHIR [1] uses a resource-based approach to data modeling, where every piece of medical data is defined as a resource. The standard defines numerous resources to represent real-world concepts in the healthcare system, such as Patient and Practitioner, as well as elements related to the healthcare process.

However, the presence of Personally Identifiable Information (PII) within FHIR poses significant challenges to data sharing due to data protection regulations. Numerous international and national laws and regulations, including the General Data Protection Regulation (GDPR) in Europe [2][3] and the Health Insurance Portability and Accountability Act (HIPAA) in the USA [4][5], require the anonymization or removal of personal or sensitive identifying information before any knowledge extraction tasks or queries are performed. Simply removing personal

LWDA'24: Lernen, Wissen, Daten, Analysen. October 23–25, 2024, Würzbzrg, Germany
wenhui.yang@item.fraunhofer.de (W. Yang); lwiese@cs.uni-frankfurt.de (L. Wiese)
0000-0003-3515-9209 (L. Wiese)

[1]http://www.hl7.org/fhir

identifiers does not adequately address privacy concerns, as quasi-identifiers can indirectly reference patients and cannot be removed due to their essential role. Additionally, removing identifiers can disrupt the data integrity due to the unique structure of FHIR datasets, where resources often reference each other using identifiers.

Several studies have proposed and analyzed methods for the anonymization and pseudonymization of FHIR resources [6][7][8]. [9] proposed a framework for anonymizing and standardizing FHIR datasets. This framework analyzes each attribute within the extracted data and applies rule-based methods to associate attributes with appropriate anonymization techniques. The framework in [10] uses a configuration file created by data experts to process FHIR data. It utilizes *FHIRPath* to locate specific elements and performs anonymization operations such as deletion, hashing, and substitution. While effective in practice, these methods have limitations. They focus solely on the Patient resource, addressing exposure issues related to its specific attributes. However, other FHIR resources should not be overlooked, as cross-references between resources can also compromise patient privacy.

In disease transmission networks or epidemiological graphs, health data is often presented in a graph-based format, similar to FHIR datasets. Various graph-based anonymization methods have been developed, including classic privacy models like k-degree [11], k-Automorphism [12], and k-Isomorphism [13]. [14] proposed sL-anonymity, which uses the Szemerédi regularity lemma to enforce k-anonymity by randomizing edges within sets of nodes to make them structurally indistinguishable. [15] introduced a method (k, l)-anonymity ensuring that even if an attacker knows up to l neighbors of a node, they cannot identify that node within a group of fewer than k nodes, by initially adding and then removing redundant edges. These models use graph structure and attribute information to prevent structural and node attribute disclosure attacks. Typically, nodes represent patients, and edges represent interactions between patients. However, this approach is unsuitable for FHIR datasets, where each node originates from different resources with distinct attribute structures, making it infeasible to directly apply general graph anonymization to FHIR data.

In this paper, we propose a privacy model $k^{(n,t)}$-anonymity suitable for FHIR datasets and employ a multidimensional anonymization method to protect patient privacy in published datasets.

2. Problem Definition

2.1. Data Model

The FHIR standard includes more than a hundred different resources, making it impractical to detail each one. In our study, we abstract and discuss five key resources: Patient, Encounter, Observation, DiagnosticReport, and MedicationAdministration. Based on their inter-referential structures and inherent attributes, we can extract the graph structure as depicted in Figure 1. In this graph, circles represent the nodes, which are the resources in the FHIR dataset, while squares represent the attributes carried by each node. Here, we simplify the attributes to include only QIs (quasi-identifiers) and SAs (sensitive attributes) as the labels of the nodes.

Let G represent the entire collection of graphs, including the records of multiple patients. $G_i = (V(G_i), E(G_i))$ denotes the graph of patient i, which is their individual record. In our

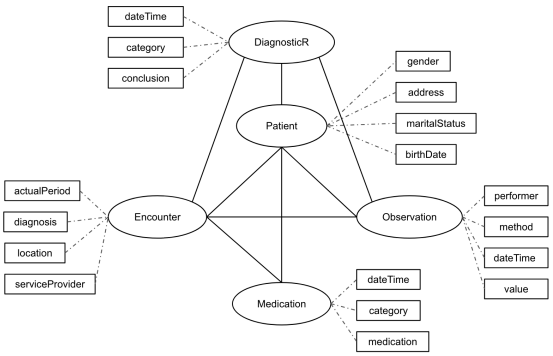

Figure 1: The data model extracted from the FHIR standard (using five types of resources as examples)

example, the nodes can be classified into five categories: P, E, O, M, and D, representing different types of resources. Each vertex $v \in V(G_i)$ is represented by a tuple of labels, denoted as $v = (l_1, l_2, \ldots, l_k)$, where l_i is the value of the i-th label. For a given node type T, the possible variants are denoted as $\{v_{T_1}, v_{T_2}, \ldots, v_{T_p}\}$, where v_{T_i} represents different variants of node type T.

2.2. Attack Model

We consider that the attacker possesses partial knowledge about an individual, specifically knowing some information present in their records. The attacker aims to use this partial knowledge to identify the complete record in G. They can use background knowledge of label values and structural relationships to filter records. If few matching records are found, there is a risk of privacy breach. We assume that the attacker knows n labels, which are distributed among different types of nodes, denoted as $L_{known} = \{l_1, l_2, \ldots, l_n\}$. Based on the known label values L_{known} the attacker can infer the corresponding node types T_{known} and their possible variants. The attacker then incrementally filters the graphs in the collection to find those containing these node variants, ultimately obtaining the candidate graph set $G_{matched}$. This process can be mathematically represented as follows:

- For each node type $T \in T_{known}$ and its set of inferred variants $\{v_{T_1}, v_{T_2}, \ldots, v_{T_p}\}$: $G_T = \{G_j \in G \mid \exists v \in V(G_j), v \in \{v_{T_1}, v_{T_2}, \ldots, v_{T_p}\}\}$
- $G_{matched} = \bigcap_{T \in T_{known}} G_T$

3. Privacy Guarantee

k^m-anonymity [16] extends k-anonymity [17] to scenarios where attackers may have knowledge of up to m elements of a record. It guarantees that even if an attacker knows up to m attributes (elements) of a record, will not be able to identify less than k records in the published data. We propose a new privacy guarantee $k^{(n,t)}$-anonymity which extends k^m-anonymity guarantee to handle graph datasets with different types of nodes. This ensures the protection of individual

identities associated with graph records from attackers with the aforementioned capabilities. We define $k^{(n,t)}$-anonymity as:

Definition 1. *A graph database G is considered $k^{(n,t)}$-anonymous if any attacker who has background knowledge of n labels and their distribution across t different types of nodes, is not able to use this knowledge to identify less than k records in G.*

In contrast to k-anonymity, which assumes that the set of QI is known, our assumption allows any node or any combination of nodes to be used by the attacker as QIs. This is due to the presence of holistic references to resources in the FHIR dataset, meaning that a node itself can serve as the QI for other nodes. Here we introduce a method to transform the original graph dataset G into a graph dataset G' that satisfies $k^{(n,t)}$-anonymity through multidimensional generalization.

Input Given a threshold k and n labels, denoted as $L_{known} = \{l_1, l_2, \dots, l_n\}$, which belong to t different types. Each type's labels can be used to infer the variants of that type. This can be represented as follows:

Let T_i represent the i-th type, where $i = 1, 2, \dots, t$. Let L_i denote the set of labels in type T_i. It can be defined as:

$$L_i = \{l_{i1}, l_{i2}, \dots, l_{i|L_i|}\}$$

where l_{ij} represents the j-th label in the i-th type, and $|L_i|$ is the number of labels in type T_i. Then $L_{known} = L_1 \cup L_2 \cup \cdots \cup L_t$. Furthermore, each type T_i can have a set of variants V_i inferred from its labels:

$$V_i = \{v_{i1}, v_{i2}, \dots, v_{im_i}\}$$

where v_{ij} represents the j-th variant of type T_i, and m_i is the number of variants for type T_i.

Node Anonymization As previously mentioned, a node can be represented by a tuple of its labels, meaning that a type of node can be understood as a relational data table. The first step is to ensure that each known individual type within the dataset achieves k-anonymity. There are many algorithms for achieving k-anonymity in relational data tables, and here we use the well-known Incognito algorithm [18], which can find minimal full-domain generalizations. Therefore, for each type node in T_{known}, we apply Incognito for initial anonymization.

Combinations Anonymization After applying the Incognito algorithm to each type of node, we obtain a set of generalization rules for the labels within each type. These generalization rules correspond to the different variants of each type. We then construct a frequency tree to represent the frequency of these variant combinations in the graph dataset. The structure is illustrated in Figure 2.

We randomly start with type T_1 among the known types, traverse all the graph records in G, and determine the frequency of each variant. Next, we move to the variants in Type T_2 and find

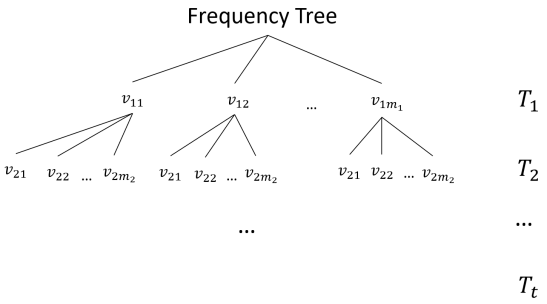

Figure 2: The frequency tree of variants and their combinations in the Graph dataset after initial anonymization

the frequency of their combinations, specifically the frequency of graphs that simultaneously has both T_1 variants and T_2 variants. This process continues in this manner until reaching T_t. Thus, this frequency tree can track graph records containing all points along the same path. Additionally, we maintain a list I to track the changes in frequency during the merging process of variants. Initially, I stores each individual variant in ascending order of their frequency of occurrence in the graph dataset. When two variants are merged, the resulting merged variant replaces the two original variants in I.

The anonymization process starts at the T_t level. We examine the frequency of all vertices at this level; if the frequency is greater than k , it indicates that the combination on that path meets the anonymization requirements. For vertices with a frequency less than k, similar variants need to be merged to increase the frequency. At this point, we query list I to identify the variant with the smallest frequency on that path. We believe that this variant affects the frequency of the combination on that path. Thus, we merge this variant with the next smallest variant of the same type. Merging means that the generalization levels of their labels will take the union. By merging variants of the same type in this manner, the frequency of the combination on each path will eventually exceed k , thus meeting the anonymization requirements.

4. Summary

This paper discusses the anonymization of FHIR datasets to protect individual records from re-identification. We first model the inherent data structure of the FHIR dataset, and then introduce a novel privacy model $k^{(n,t)}$-anonymity, which aims to enhance the protection of individual records by considering the number of labels n and different types of resources t. To meet the privacy requirements specified by the $k^{(n,t)}$ model, we propose a multidimensional anonymization method, providing a practical and effective solution for FHIR dataset anonymization. However, it also has some potential limitations, such as the need to consider data utility during the combinations anonymization process. Future work plans to evaluate the practicality and impact of this approach.

References

[1] A. Aminifar, Y. Lamo, K. I. Pun, F. Rabbi, A practical methodology for anonymization of structured health data (2019).

[2] P. Regulation, Regulation (eu) 2016/679 of the european parliament and of the council, Regulation (eu) 679 (2016) 2016.

[3] J. P. Albrecht, How the gdpr will change the world, Eur. Data Prot. L. Rev. 2 (2016) 287.

[4] A. Act, Health insurance portability and accountability act of 1996, Public law 104 (1996) 191.

[5] O. for Civil Rights, Summary of the hipaa privacy rule, 2022. URL: https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html, accessed: 2024-07-04.

[6] S. Neto, F. S. Ferraz, C. A. G. Ferraz, Towards identity management in healthcare systems, in: Proceedings on the International Conference on Internet Computing (ICOMP), The Steering Committee of The World Congress in Computer Science, Computer ..., 2016, p. 157.

[7] D.-Y. Kim, S.-h. Hwang, M.-G. Kim, J.-H. Song, S.-W. Lee, I. K. Kim, Development of parkinson patient generated data collection platform using fhir and iot devices, in: MEDINFO 2017: Precision Healthcare through Informatics, IOS Press, 2017, pp. 141–145.

[8] S. Dimopoulou, C. Symvoulidis, K. Koutsoukos, A. Kiourtis, A. Mavrogiorgou, D. Kyriazis, Mobile anonymization and pseudonymization of structured health data for research, in: 2022 Seventh International Conference On Mobile And Secure Services (MobiSecServ), IEEE, 2022, pp. 1–6.

[9] M. Ciampi, M. Sicuranza, S. Silvestri, A privacy-preserving and standard-based architecture for secondary use of clinical data, Information 13 (2022) 87.

[10] E. Raso, P. Loreti, M. Ravaziol, L. Bracciale, Anonymisation and pseudonymisation of fhir resources for secondary use of healthcare data, IEEE Access (2024).

[11] K. Liu, E. Terzi, Towards identity anonymization on graphs, in: Proceedings of the 2008 ACM SIGMOD international conference on Management of data, 2008, pp. 93–106.

[12] L. Zou, L. Chen, M. T. Özsu, K-automorphism: A general framework for privacy preserving network publication, Proceedings of the VLDB Endowment 2 (2009) 946–957.

[13] J. Cheng, A. W.-c. Fu, J. Liu, K-isomorphism: privacy preserving network publication against structural attacks, in: Proceedings of the 2010 ACM SIGMOD International Conference on Management of data, 2010, pp. 459–470.

[14] D. Foffano, L. Rossi, A. Torsello, You can't see me: Anonymizing graphs using the szemerédi regularity lemma, Frontiers in big Data 2 (2019) 7.

[15] R. Mortazavi, S. Erfani, Gram: An efficient (k, l) graph anonymization method, Expert Systems with Applications 153 (2020) 113454.

[16] M. Terrovitis, N. Mamoulis, P. Kalnis, Privacy-preserving anonymization of set-valued data, Proceedings of the VLDB Endowment 1 (2008) 115–125.

[17] L. Sweeney, k-anonymity: A model for protecting privacy, International journal of uncertainty, fuzziness and knowledge-based systems 10 (2002) 557–570.

[18] K. LeFevre, D. J. DeWitt, R. Ramakrishnan, Incognito: Efficient full-domain k-anonymity, in: Proceedings of the 2005 ACM SIGMOD international conference on Management of data, 2005, pp. 49–60.