Journal of
Data Acquisition and Processing

ISSN: 1004-9037
https://sjcjycl.cn/
DOI: 10.5281/zenodo.7698639

EVALUATION OF THE PRIVACY-PROTECTING EFFECTS OF LEARNING-BASED IOT ECOSYSTEM BEHAVIOR

[1] Srikanta Kolay [2]Dr Tryambak Hiwarkar
[1] Research Scholar [2]Professor
[1,2] Department of Computer Science& Engineering
[1,2] Sardar Patel University Balaghat

Abstract:

The Internet of Things has allowed for the development of numerous consumer-facing apps and services that enhance our knowledge of and ability to shape our built environments and the quality of our individual lives. These services couldn't exist without the persistent gathering and analysing of private and personal information about users. When it comes to protecting users from identification, profiling, localization and tracking, and information linking, smart heath care is one of the many IoT applications that requires privacy preservation strategies. Finding the right balance between privacy protection, data utility, and acceptable system performance in terms of accuracy, runtime, and resource consumption requires carefully selecting privacy preservation techniques (and solutions) based on the nature of data, system performance requirements, and resource constraints.

In this study, we evaluate the effects of introducing our preferred privacy preservation techniques on the functionality of various nodes in the IoT ecosystem, both in terms of data utility and overall system performance. Using both real-world and synthetic privacy-preserving smart health care datasets, we build, illustrate, and assess the results of our proposed methodologies. We begin with a comprehensive taxonomy and analysis of privacy preservation strategies and solutions that can be used as a starting point for making informed decisions about which methods to employ given the specifics of a given data set and the constraints of a given system. Furthermore, we discuss and implement a strategy for constructing realistic synthetic and privacy-preserving smart health care datasets utilising Generative Adversarial Networks and Differential Privacy to promote privacy-preserving data exchange. We utilise healthcare data as an example later on to describe and design a solution for private data analytics: the differential privacy library PyDPLib.

We present and implement a novel approach to reconfigurable data privacy in machine learning on resource-limited computing devices, complete with corresponding algorithms and an end-to-end system pipeline. This allows us to find appropriate trade-offs between providing the necessary privacy preservation, device resource consumption, and application accuracy.

Keywords: IoT Eco system, GAN

Introduction:

IoT Ecosystem : Kevin Ashton [1] is credited with creating the term "Internet of Things" in 1999. It refers to the system of networked sensors and devices that serve as the backbone of the data collection used by today's apps and services. Smart sensors can be defined as any type of sensor that can simultaneously gather data, process that data with on-board circuitry, and then transmit that processed data. The Internet of Things consists of all the devices that can gather data and send it to a central location via a network. The term "Internet of Things ecosystem" refers to the interconnected network of devices and the corresponding software and technology that enable the delivery of services based on large quantities of data. The software and hardware platforms, as well as the standards typically utilised for enabling such interconnection, may become a core of an IoT ecosystem, as pointed out by Mazhelis et al. [2]. In this thesis, we analyse the data, resource-limited edge devices, ML applications, and data analytics and visualisation that make up the Internet of Things ecosystem..

Privacy Preservation.: There is a great deal of flexibility in the methods that can be used to protect individual privacy. Information flow control [3], data or model obfuscation via noise addition [4], cryptography [5, 6], anonymization via generalisation and suppression of attributes [7-9], and the usage of private compute units [10] are all examples of procedures used by privacy preservation strategies. It is noted in [11] that all of these methods have some sort of performance overhead, be it in the form of greater resource use, decreased model accuracy, or longer processing times. So that the IoT ecosystem can meet the needs of volume, velocity, and variety, privacy preservation solutions that are lightweight, scalable, and efficient need to be applied to its many components. By doing so, we may determine the optimal compromises between the various goals of privacy protection, data utility retention, application performance, model accuracy, and device resource utilisation overhead.

Smart Health Care:

The field of Internet of Things (IoT) known as "smart health care" is expanding rapidly and has implications for both people and society as a whole. While the individual goal is to improve the quality of medical treatment while lower standards of living costs and providing better and personalised health care [12], the societal goal is to increase the ability to track physical activities and diet patterns via wearable trackers and diet logging applications.

In order to take use of machine learning-based services, the smart health care business is dependent on the availability of large-scale health information. Nonetheless, privacy rules like the EU's General Data Protection Regulation (GDPR) [13] to be followed in the collection and processing of this health data. Thus, there are unique difficulties in assuring patient confidentiality in the field of smart healthcare: There are a number of challenges when working with health care data: (1) there is a lack of open data to experiment with because the data is of a highly sensitive and private nature; (2) there are limitations on sharing data with confidential medical information because of the risk of misuse or re-identification; (3) smart health care data has highly diverse data types and formats, and bounded ranges; (4) health care data requires high utility and offers low tolerance towards data perturbance; and (5) there is low tolerance towards data quality variation. These factors combine to make privacy protection in the IoT sector of smart health care a formidable obstacle.

II. Literature Survey

During the past decade, the proliferation of the Internet of Things has led to the development of an incredible variety of smart devices (IoT). Cisco's Visual Networking Index (VNI) 2021 forecast [46] estimates that by the end of 2021, there will be 3.5 networked devices for every person on Earth. Any sensor with the ability to collect data, process it using integrated circuitry, and transmit it is considered a smart gadget. To better serve their users, these gadgets often upload collected data to the cloud for further analysis and storage. The term "edge computing" describes a more advanced variant of this method in which extra processing and analytical skills are offloaded to the devices.

Smith et al. [1] Privacy concerns are defined by as worries regarding the collection of personal information, worries regarding unauthorised secondary use (both internally and externally in organisations), worries regarding erroneous access to personal data, and worries regarding errors in collected personal information.

Before we can classify solutions for protecting users' privacy, we must first define the privacy problems that exist inside the IoT ecosystem and the architecture layers that are related to it. Following that, we will provide a broad overview of attacks on privacy as well as the risks that are associated with them.

Ziegeldorf et al. [2]. After that, we will give a quick explanation of the risks, and then we will discuss the many IoT architectural levels that are to blame for those risks. Take note that the threats commonly coexist in IoT systems, but this is mostly dependent on the sort of service that is being given.

Papernot et al. [3] conduct a research of the most advanced machine learning algorithms that safeguard users' personal information. In addition, differential privacy is often used in machine learning models as a defense mechanism against assaults that attempt to flip the model. is a fascinating illustration of a machine learning-as-a-service that protects users' privacy and was developed especially for cloud settings. Cloud environments are an essential part of the ecosystem for the Internet of Things (IoT). It does this by utilizing private computing units, which strengthens the promises of privacy (with SGX). In addition, there are implementations of k-anonymity mixed with ML algorithms and cryptographic approaches that use ML that are based on published works of the academic community. Federated learning is a method of data mining that is utilized by a few different implementations of recommender systems.

The technique of collaborative filtering is utilized rather frequently in recommender systems, which is an example of an implementation that protects privacy, combines k-anonymity with collaborative filtering. [5], which utilises misdirection; and [120], which helps make use of homomorphic encryption and discrepancy privacy to ensure that recommendations are kept private. [119]; [5]; and [120] are all examples of implementations that protect privacy. In addition, [89] suggests using a federated machine learning variation of collaborative filtering to provide more personalized suggestions.

III. Data Processing

In this part, we will detail the approach that we use to collect, impute, and alter data. In addition, we offer our methods for training models in a way that does not compromise users' privacy, which is then followed by the process for data inversion**.**

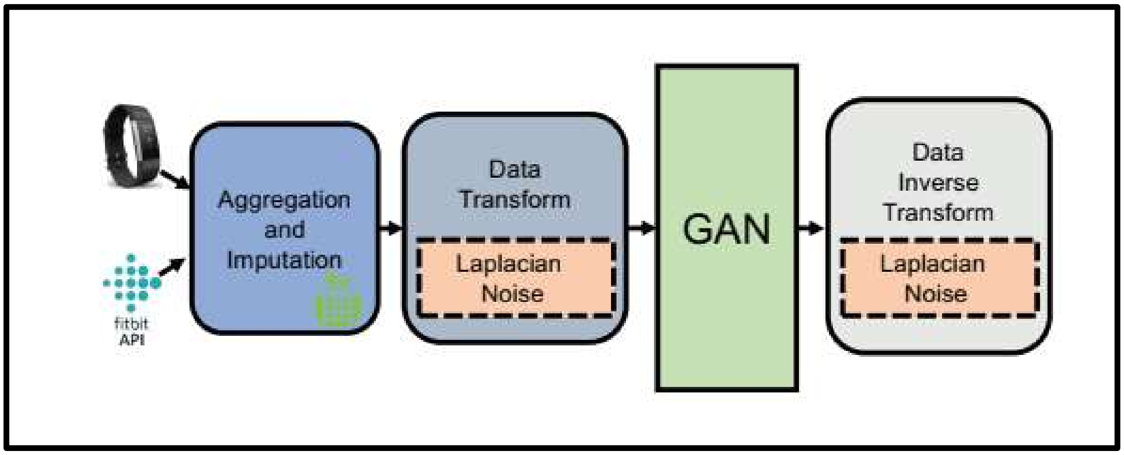

Figure: 3.1 Data Preprocessing Process

3.1 The collection of information and imputing of values

For the purpose of this investigation, we merged the automatic data collection capabilities of the Fitbit Charge 2 HR wearables with the manually logged meal capabilities of the Fitbit app. Throughout the course of this inquiry, a total of 25 people from Belgium and Sweden were under observation at various points. Data were acquired with the use of 12 different devices from two continuing participants (a man and a woman) as well as ten active users. The users were given the instruction to record their observations for a period of at least one month. In order to provide an accurate picture of the eating patterns and overall health of the local populations, we included people from six different widely categorized ethnic groups in our participant pool.

More than 17 million measurements were taken of various aspects of the users' meal records, calorie intake, heart rate, number of calories burned, number of steps done, daily activity pattern, and sleep. The website not only collects numerical data, but it also collects demographic information about users, such as their age, gender, height, and weight, in addition to collecting number data. Due to the fact that the users were not provided with intelligent scales, the weight measurement is recorded manually. After that, the Fitbit platform was utilised to export each and every one of these data and analytics. The data that was collected revealed a great number of inconsistencies and user errors, such as the following: 1) users forgetting to wear the watch on some days; 2) incorrectly recording very large portion sizes of meals; 3) manually recording the meals without a caloric breakdown; and 4) incorrectly wearing the watch, which led to inconsistencies between recorded activities; Also, there were some people who reported their data in languages other than English, such as French or Italian.

3.1.1 Meal logs imputation

For the purpose of this investigation, we merged the automatic data collection capabilities of the Fitbit Charge 2 HR wearables with the manually logged meal capabilities of the Fitbit app. Throughout the course of this inquiry, a total of 25 people from Belgium and Sweden were under observation at various points. Data were acquired with the use of 12 different devices from two continuing participants (a man and a woman) as well as ten active users. The users were given the instruction to record their observations for a period of at least one month. In order to provide an accurate picture of the eating patterns and overall health of the local populations, we included people from six different widely categorised ethnic groups in our participant pool.

More than 17 million measurements were taken of various aspects of the users' meal records, calorie intake, heart rate, number of calories burned, number of steps done, daily activity pattern, and sleep. The website not only collects numerical data, but it also collects demographic information about users, such as their age, gender, height, and weight, in addition to collecting number data. Due to the fact that the users were not provided with intelligent scales, the weight measurement is recorded manually. After that, the Fitbit platform was utilised to export each and every one of these data and analytics. The data that was collected revealed a great number of inconsistencies and user errors, such as the following: 1) users forgetting to wear the watch on some days; 2) incorrectly recording very large portion sizes of meals; 3) manually recording the meals without a caloric breakdown; and 4) incorrectly wearing the watch, which led to inconsistencies between recorded activities; Also, there were some people who reported their data in languages other than English, such as French or Italian.

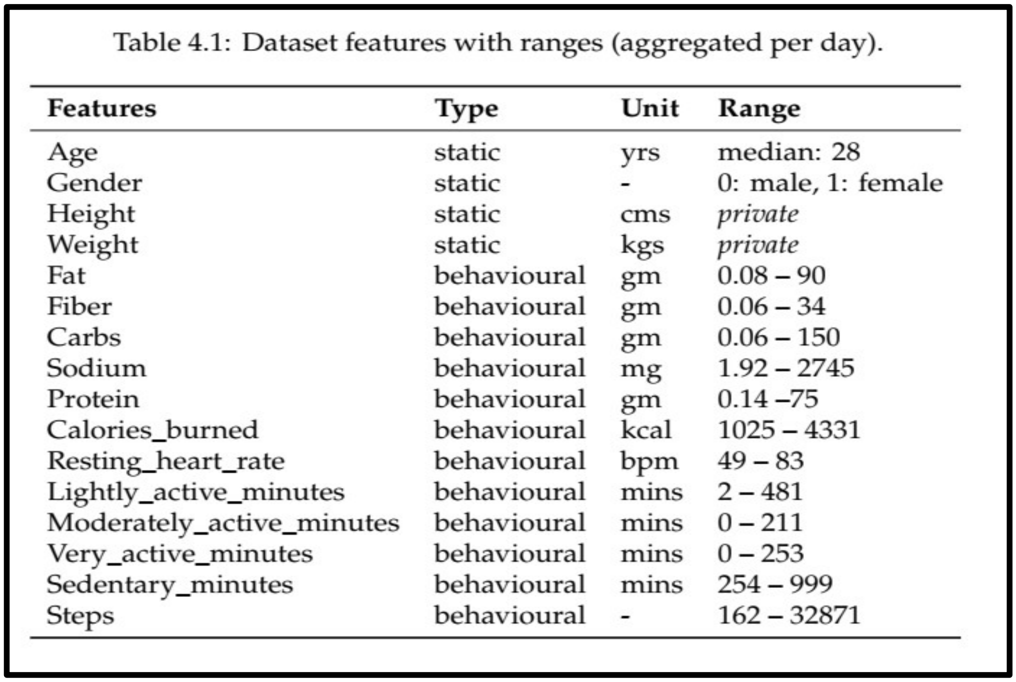

Table 4.1: Dataset features with ranges (aggregated per day).

Features	Type	Unit	Range
Age	static	yrs	median: 28
Gender	static	-	0: male, 1: female
Height	static	cms	*private*
Weight	static	kgs	*private*
Fat	behavioural	gm	0.08 – 90
Fiber	behavioural	gm	0.06 – 34
Carbs	behavioural	gm	0.06 – 150
Sodium	behavioural	mg	1.92 – 2745
Protein	behavioural	gm	0.14 –75
Calories_burned	behavioural	kcal	1025 – 4331
Resting_heart_rate	behavioural	bpm	49 – 83
Lightly_active_minutes	behavioural	mins	2 – 481
Moderately_active_minutes	behavioural	mins	0 – 211
Very_active_minutes	behavioural	mins	0 – 253
Sedentary_minutes	behavioural	mins	254 – 999
Steps	behavioural	-	162 – 32871

3.1.2 Transformation of Data

In order to get the data ready for training, we start by stripping it of any date and gender information that could be in there. After this, the remaining characteristics are normalised, and the result is given back into the model so that it may be trained.

We are able to send DP-input data to the GAN, and this, in conjunction with the privacy setting (noisy input), enables the GAN to produce DP-synthetic samples in accordance with the requirements of the post-processing theorem. With that, we include some Laplacian noise "because the data comprises categorical or static properties, which call for more stringent privacy settings." "= 0:2 to assure high noise addition and, as a result, more restrictive privacy settings."

On the other hand, the probability of re-identification is significantly reduced by using behavioural traits. In order to do this, we mix the Laplacian noise with "= 0:5 to guarantee enough addition of noise without compromising the usefulness of the data

IV. Model Training

In order to generate synthetic data samples, we make use of BGAN. To train the BGAN, we chose participants at random, taking into account their gender and geographic area. We trained the model in each of these three distinct privacy configurations (non-DP, noisy input, and noisy output)

V. Results:

our pipeline offers a multitude of sites for the insertion of Laplacian noise to provide differential privacy. This makes it possible to build up three separate experiments. The GAN network is able to acquire knowledge of the distribution and produce examples that are convincing for each scenario.

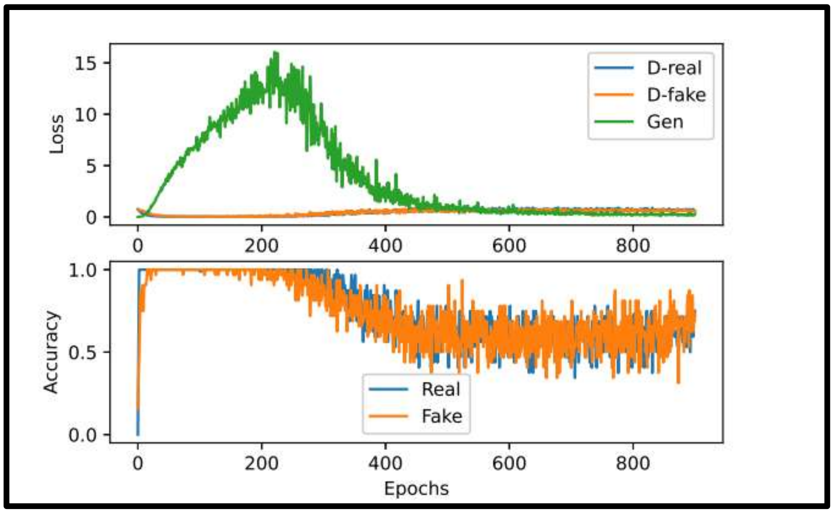

Fig: 5 Line Plots of loss and Accuracy for stable GAN

Table **5.1** Example data samples from Belgium population. Age and Gender are hidden.

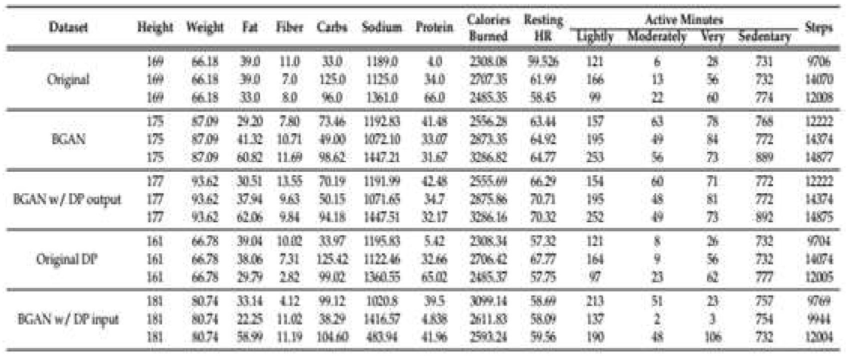

Dataset	Height	Weight	Fat	Fiber	Carbs	Sodium	Protein	Calories Burned	Resting HR	Active Minutes				Steps
										Lightly	Moderately	Very	Sedentary	
Original	169	66.18	39.0	11.0	33.0	1189.0	4.0	2308.08	59.526	121	6	28	731	9706
	169	66.18	39.0	7.0	125.0	1125.0	34.0	2707.35	61.99	166	13	56	732	14070
	169	66.18	33.0	8.0	96.0	1361.0	66.0	2485.35	58.45	99	22	60	774	12008
BGAN	175	87.09	29.20	7.80	73.46	1192.83	41.48	2556.28	63.44	157	63	78	768	12222
	175	87.09	41.32	10.71	49.00	1072.10	33.07	2873.35	64.92	195	49	84	772	14374
	175	87.09	60.82	11.69	98.62	1447.21	31.67	3286.82	64.77	253	56	73	889	14877
BGAN w/ DP output	177	93.62	30.51	13.55	70.19	1191.99	42.48	2555.69	66.29	154	60	71	772	12222
	177	93.62	37.94	9.63	50.15	1071.65	34.7	2875.86	70.71	195	48	81	772	14374
	177	93.62	62.06	9.84	94.18	1447.51	32.17	3286.16	70.32	252	49	73	892	14875
Original DP	161	66.78	39.04	10.02	33.97	1195.83	5.42	2308.34	57.32	121	8	26	732	9704
	161	66.78	38.06	7.31	125.42	1122.46	32.66	2706.42	67.77	164	9	56	732	14074
	161	66.78	29.79	2.82	99.02	1360.55	65.02	2485.37	57.75	97	23	62	777	12005
BGAN w/ DP input	181	80.74	33.14	4.12	99.12	1020.8	39.5	3099.14	58.69	213	51	23	757	9769
	181	80.74	22.25	11.02	38.29	1416.57	4.838	2611.83	58.09	137	2	3	754	9944
	181	80.74	58.99	11.19	104.60	483.94	41.96	2593.24	59.56	190	48	106	732	12004

Discussion:

Our proposed technique generates private and artificially intelligent smart health care data through the utilisation of BGAN and DP. The GAN network has the potential to provide outcomes that are consistent and credible. The stability of the proposed GAN can be seen in the top subplot which presents line graphs for the discriminator loss for genuine data (blue), the discriminator loss for manufactured fake samples (orange), and the generator loss for generated fake samples (green). As can be seen, the three losses first exhibit a significant amount of instability until reaching a point of stability between epochs 420 and 600. Beyond that point, losses continue to be constant, which demonstrates the GAN's reliable behaviour.

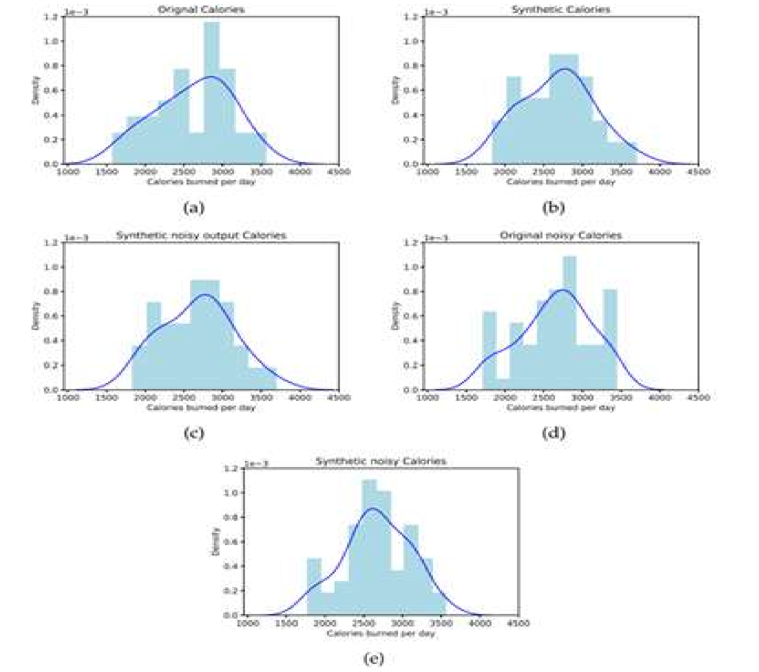

Figure 6 Histogram distributions for calories burned per day (kcal). Samples: Belgium males with RHR=70-75bpm.

despite the fact that the variety grows. The discriminator experiences a loss of around 0:5 regardless of whether it is presented with genuine samples or false samples, but the loss experienced by the generator is substantially greater between the ratios of 0:5 and 1:0. It is anticipated that the model will provide data that is plausible between epochs 420 and 600. In the bottom subplot, you can observe a line plot depicting the discriminator accuracy during training on genuine (blue) and false (orange) samples. A behaviour that is analogous to this can be seen in the subplot of loss, which shows that the accuracy initially varies considerably between the two sample types, stabilises between epochs 420 and 600 at a value that is roughly 60–70%, and then continues to stabilise after that, albeit with a greater degree of variation.

Table 5.1 displays an example row from the primary dataset, whereas the synthetic rows are the result of trained GANs being applied to the dataset. In this context, Original and GAN both stand for samples taken from datasets before the introduction of noise (non-DP). The produced data samples that have DP-noise added to them are displayed in the GAN along with the DP output (noisy output). In a manner analogous to that of Original DP, GAN with DP input represents, respectively, the synthetic samples that were constructed and the original DP-input (noisy input). As can be seen, the generated instances all give the impression of being realistic given the privacy settings that were selected.

Results for Fitbit Dataset

To begin, we begin by calculating the cost measures for the non-private version of the Fitbit dataset on the edge device. These cost measures include memory usage, bandwidth needs, and the processor instructions. The amount of RAM that was taken up by the non-private data in its entirety was 247 KB. Because of the relatively modest size of the dataset, the needed bandwidth for this dataset over a 4G network with 100 Mbps was just 0.0198 Mbps. Due to the fact that this dataset is not private, there were no additional processor instructions that were necessary for it to be processed.

Following that, we used the privacy encoder injective functions that were built for each feature to turn every feature in this dataset into a private feature. Using the use of the appropriate injective privacy functions, we determined the amount of memory that would be needed to process each feature. The total amount of RAM that was utilised throughout this procedure for the protection of personal information was 533.378 KB. After the transformation of the features into privacy encoded features, the total memory of the private dataframe was calculated to be 1,750.369 KB. If we were to make all of the characteristics of the Fitbit dataset private, the total amount of additional RAM that would be required would thus be 2283 747 KB. It would take 0.1827 Mbps of bandwidth to send this data over a local area network that has a speed of 100 Mbps. The Raspberry pi RPi 1 Model A required a total of 252 processor instructions in order to turn the non-private features into the private features.

There are four aspects that, due to the sensitive nature of our data, we have decided to make private in order to safeguard our users. These features are represented by the letter EF. These essential factors are one's age, gender, height, and weight, all of which must be kept confidential regardless of the availability of resources. After then, we determined the amount of resources used by the dataset that included these four EF in the private category. In this particular instance, we only transformed these four Elements by utilising the injective privacy encoder methods that were built into them; the values for all of the other features were left unchanged.

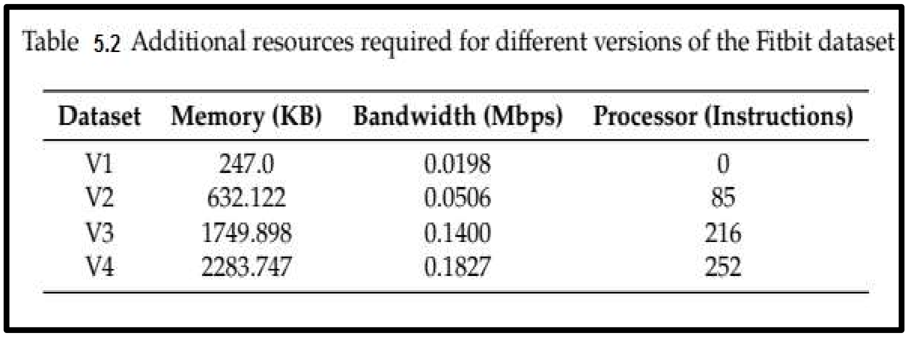

Table 5.2 Additional resources required for different versions of the Fitbit dataset

Dataset	Memory (KB)	Bandwidth (Mbps)	Processor (Instructions)
V1	247.0	0.0198	0
V2	632.122	0.0506	85
V3	1749.898	0.1400	216
V4	2283.747	0.1827	252

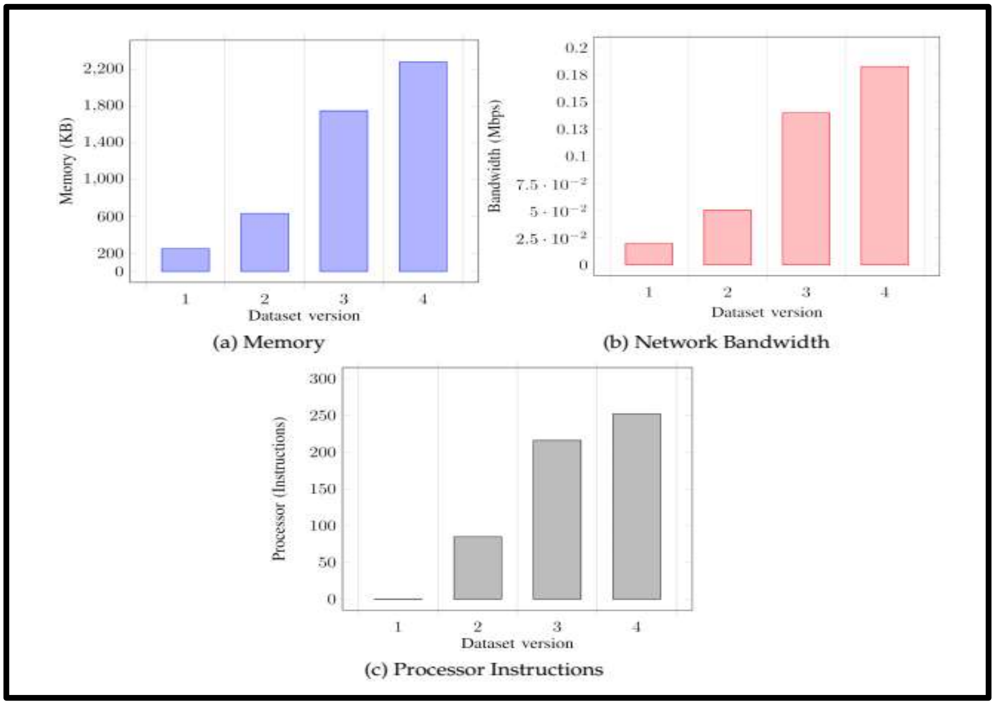

(a) Memory

(b) Network Bandwidth

(c) Processor Instructions

Figure 5.1: Additional resource consumption for increasingly private versions of Fitbit dataset.

Conclusion:

We designed, implemented, and evaluated a solution for generating realistic synthetic private smart health care datasets from sensitive non-private datasets in order to enable privacy-preserving data sharing. This was done in order to generate realistic synthetic private smart health care datasets. For the purpose of producing realistic and confidential smart health care datasets, we suggested using a generative adversarial network model in conjunction with differential privacy safeguards. Our solution was tailored to meet the particular problems posed by smart health care data, which included volume, velocity, and diversity in the form of a wide range of data kinds and distributions. In addition to enriching and augmenting the input data samples, our proposed solution was also able to generate realistic synthetic data samples and differentially private data samples under a variety of conditions, including learning from a noisy distribution and then noisifying the distribution that was learned. We put our suggested method through its paces by testing and evaluating it using data taken from an actual Fitbit device. According to the findings of our research, our method is able to produce synthetic differentially private datasets of a high quality that are able to maintain the statistical features of the dataset that was originally collected.

Reference:

1. H. J. Smith, S. J. Milberg, and S. J. Burke, “Information privacy: measuring individuals’ concerns about organizational practices,” *MIS quarterly*, pp. 167–196, 1996

2. J. H. Ziegeldorf, O. G. Morchon, and K. Wehrle, “Privacy in the internet of things: threats and challenges,” *Security and Communication Networks*, vol. 7, no. 12, pp. 2728–2742, 2014.

3. N. Papernot, P. McDaniel, A. Sinha, and M. P. Wellman, “SoK: Security and privacy in machine learning,” in *IEEE EuroS&P*. IEEE, 2018, pp. 399–414.

4. K. Zhao and L. Ge, “A survey on the internet of things security,” in *9th International Conference on Computational Intelligence and Security*, Dec 2013, pp. 63–667.

5. -K. Chen, “Challenges and opportunities of internet of things,” in *17th Asia and South Pacific design automation conference*. IEEE, 2012, pp. 383–388.

6. N. Papernot, M. Abadi, Ú. Erlingsson, I. Goodfellow, and K. Talwar, “Semisupervised knowledge transfer for deep learning from private training data,” *arXiv preprint arXiv:1610.05755*, 2016.

7. N. Carlini, C. Liu, J. Kos, Ú. Erlingsson, and D. Song, “The secret sharer: Measuring unintended neural network memorization & extracting secrets,” *arXiv preprint arXiv:1802.08232*, 2018.

8. F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, “Stealing machine learning models via prediction apis,” in *25th* f*USENIX*g *Security Symposium (*f*USENIX*g *Security 16)*, 2016, pp. 601–618.

9. B. Wang and N. Z. Gong, “Stealing hyperparameters in machine learning,” in *Symposium on Security and Privacy (IEEE S&P)*. IEEE, 2018, pp. 36–52.

10. M. Juuti, S. Szyller, A. Dmitrenko, S. Marchal, and N. Asokan, “Prada: protecting against DNN model stealing attacks,” arXiv preprint arXiv:1805.02628, 2018.

11. Z. Ji, Z. C. Lipton, and C. Elkan, “Differential privacy and machine learning: a survey and review,” arXiv preprint arXiv:1412.7584, 2014.

12. M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, “Deep learning with differential privacy,” in Proceedings of SIGSAC Conference on Computer and Communications Security. ACM, 2016, pp. 308–318.

13. C. Dwork and V. Feldman, “Privacy-preserving prediction,” arXiv preprint arXiv:1803.10266, 2018.

14. T. Hunt, C. Song, R. Shokri, V. Shmatikov, and E. Witchel, “Chiron: Privacypreserving machine learning as a service,” arXiv preprint arXiv:1803.05961, 2018.

15. S. Nijssen and E. Fromont, “Optimal constraint-based decision tree induction from itemset lattices,” Data Mining and Knowledge Discovery, vol. 21, no. 1, pp. 9–51, 2010.