Gestational weight gain prediction using privacy preserving federated learning

Chetanya Puri[1,3], Koustabh Dolui[2,3], Gerben Kooijman[3], Felipe Masculo[3], Shannon Van Sambeek[3], Sebastiaan Den Boer[3], Sam Michiels[2], Hans Hallez[2], Stijn Luca[4*] and Bart Vanrumste[1*]

***Abstract*— Gestational weight gain prediction in expecting women is associated with multiple risks. Manageable interventions can be devised if the weight gain can be predicted as early as possible. However, training the model to predict such weight gain requires access to centrally stored privacy sensitive weight data. Federated learning can help mitigate this problem by sending local copies of trained models instead of raw data and aggregate them at the central server. In this paper, we present a privacy preserving federated learning approach where the participating users collaboratively learn and update the global model. Furthermore, we show that this model updation can be done incrementally without having the need to store the local updates eternally. Our proposed model achieves a mean absolute error of 4.455 kgs whilst preserving privacy against 2.572 kgs achieved in a centralised approach.**

***Clinical relevance*— Privacy preserving training of machine learning algorithm for early gestational weight gain prediction with minor tradeoff to performance.**

I. INTRODUCTION

In pregnancy, inadequate or excessive weight gain remains a key health issue. Global estimates suggest that only around 30% of pregnant women end up being adequately weighed recommended by the Institute of medicine [5], [11]. There are several risks associated with such excessive or inadequate gestational weight gain, for example, excessive weight gain can lead to fetal macrosomia or post-partum maternal obesity putting the mothers at incresased risk of gestational diabetes [4]. Similarly, inadequate weight gain can lead to small-for-gestational-age infants [5].

Early prediction of gestational weight gain can help mitigate this problem by helping neonatal healthcare providers or expecting women in devising better management and interventions. Traditional approaches exist in which raw data from all the subjects is collected and sent to a central location. At this central location, the data is saved, processed and models to estimate gestational weight gain are trained. Even though, one can achieve high predictive performance in such a centralised approach, there are several privacy concerns associated with such a model building approach, especially in the light of the General Data Protection Regulation (GDPR) imposed by the European Union (EU) and increased awareness about privacy preservation among end-users. The centralized storage creates a large surface area for security and privacy attacks. It leaves the user lacking control of his/her own personal data. Finally, in applications where the data collected is large in size, especially larger than the model, handling sensitive data on the server side becomes cumbersome as well.

Google proposed federated learning [7] where many local devices collaboratively train a model in association with a central server, while keeping raw sensitive data distributed in the users' own devices. This is made possible by the ubiquity and the improved computational capabilities of the edge devices such as smart-phones. In federated learning, user devices only share model updates with the centralized server after training models iteratively on the local data available on-device. Federated learning is particularly applicable to use-cases where the data is collected from user devices and the data is sensitive in nature. In order to achieve this, we have designed and implemented a privacy-preserving federated approach for the prediction of gestational weight gain. The key contributions of this paper are (a) implementation of federated learning approach for prediction of gestational weight gain, (b) studying the effect of varying number of participants in collaborative learning, and (c) updating the global model incrementally such that the local updates are deleted once they are incorporated into a global model.

II. RELATED WORKS

Parametric methods such as maximum likelihood estimation or ARIMA [2] approaches have been used traditionally for time series prediction that utilise individual training data. Authors in [10] propose an improvement over these state-of-the-art techniques to predict an individual's end-of-pregnancy weight gain as early as day 140 with an average mean absolute error of around 2.572 kgs. The model is trained by learning an a-priori model based on data from other users stored at a data center and using this information in association with limited data from test individual to predict reliably. Such a centralised data storage implementation needs access to centrally stored data from a variety of users, in this case, pregnant women. This high performance is achieved at the expense of privacy sensitive information of

[1]eMedia Lab and STADIUS, Department of Electrical Engineering (ESAT), KU Leuven, Belgium. `<firstname>.<lastname>@kuleuven.be`

[2]imec-DistriNet, KU Leuven, Belgium `<firstname>.<lastname>@kuleuven.be`

[3]Philips Research, Eindhoven, The Netherlands `<firstname>.<lastname>@philips.com`

[4]Department of Data Analysis and Mathematical Modelling, Ghent University, Ghent, Belgium stijn.luca@ugent.be

*These authors contributed equally to this work.

users. Authors in [9] prove that such decentralised learning approach can help predict the gestational weight gain reliably with privacy preserved. However, the model aggregation in which local models are combined to form a single global model required storage of the local models eternally on the central server. This can lead to various forms of attacks on the models stored on the server or intercepted model updates including model inversion [3] and privacy leakage [8], [1]. In this work, we built upon our previous works and propose that such distributed learning can also be achieved by learning the global model incrementally without storing the local updates for infinite amount of time.

III. Data

We consider data from 80 women that were in their gestational week 5 or later recruited in Eindhoven, The Netherlands. The weight data was collected using a WiFi-connected weight scale, Withings WS30[1]. The participants were asked to log their weights weekly and the recorded weight data was sent to the cloud via a mobile application. Additional meta-data such as age, height and pre-pregnancy weight were also collected. The participants provided an informed consent pre-data collection and the study was approved by the Internal Ethics Committee for Biomedical Experiments of the involved organizations (ICBE Reference number 2015-0079). This sample dataset's distribution is close to that in [5], which is obtained from a large population of more than a million women, with almost half of the women gaining above the recommended guidelines [10].

TABLE I: Dataset description

Dataset Attribute	**Mean $\pm$ std**
Age (years)	31 ± 3.5
Height (meters)	1.69 ± 0.07
Pre-pregnancy weight (kgs)	69 ± 15
Pre-pregnancy BMI (kgs/m^2)	24 ± 4
Delivery (days)	277 ± 10
Weight Gained (kgs)	13.7 ± 4.7
Number of recorded weight gain samples	59.83 ± 41.02

IV. Methods

Given a population of N subjects that acquired N time series of gestational weight gain measurements as $\mathcal{X} = \{(\mathbf{x}^1, \mathbf{y}^1), \cdots, (\mathbf{x}^N, \mathbf{y}^N)\}$, where $\mathbf{x}^i = \{t_1^i, t_2^i, t_3^i, \cdots, t_{m_i}^i\}$ represents the input gestational days upto delivery day $t_{m_i}^i$ and $\mathbf{y}^i = \{y_1^i, y_2^i, y_3^i, \cdots, y_{m_i}^i)\}$ represents the output weight gain for i^{th} subject, where $y_k^i = y(t_k^i)$. It is important to note here that t_k^i *does not necessarily equal* $t_k^j, i, j \in \{1, 2, \cdots, N\}$. This is because the data is *self-reported* such that each subject acquires measurements at different times according to their personal preferences and adherence to data collection.

Furthermore, we are given individual weight measurements from test subject's initial t_d^+ days of pregnancy data, $\mathcal{D} = \{(t_1^+, y_1^+), (t_2^+, y_2^+), \cdots, (t_d^+, y_d^+)\}$.

[1]https://www.withings.com/

We try to learn function(s) f from $\mathcal{X}$ and $\mathcal{D}$, such that,

$$y^+ = f(t^+) + \epsilon \tag{1}$$

where $\epsilon \sim \mathcal{N}(0, \sigma^2)$ is independent and identically distributed (i.i.d) according to a Gaussian.

A. Centralised parametric approach

Traditionally used methods include parametric approach like fitting a p^{th}-order polynomial with $f = w_0 + w_1 t + w_2 t^2 + \cdots + w_p t^p$ in eq. (1) and estimating the coefficients $\mathbf{w} = [w_0, w_1, \cdots, w_p]^T$ by maximizing the likelihood ($\mathcal{L}$) over an individual's personal-training data $\mathcal{D}$, $\mathcal{L}(\mathbf{w}) = P(\mathcal{D}|\mathbf{w})$,

$$\hat{\mathbf{w}}_{MLE} = \underset{\mathbf{w}}{\operatorname{argmax}}\, P(\mathcal{D}|\mathbf{w}) = \prod_{i=1}^{d} p(y_i^+ | t_i^+; \mathbf{w}) \tag{2}$$

This can be done on a local device using only the estimates of a single user following eq. (2) that refers to the model learnt from the individual's sparse limited observations upto given t_d days. Often, such a prediction is far from reliable as it uses only few points from personal data. Authors in [10] show that such a prediction can be improved by considering the public-training data. The public-training data ($\mathcal{X}$) can be exploited and the maximum likelihood point estimates (MLE) of $\hat{\mathbf{w}}^i$ for each individual time series in the public-training data following eq. (2) can be derived. If we assume gaussianity over the distribution of $\mathbf{w}$ such that $\mathbf{w} \sim \mathcal{N}(\mu_{\hat{\mathbf{w}}}, \Sigma_{\hat{\mathbf{w}}})$, we can find a closed-form solution of $\mathbf{w}_{MAP}$ analytically. Here, $\mu_{\hat{\mathbf{w}}}^N = mean([\hat{\mathbf{w}}^1, \hat{\mathbf{w}}^2, \cdots, \hat{\mathbf{w}}^N]^T)$, $\Sigma_{\hat{\mathbf{w}}}^N = cov([\hat{\mathbf{w}}^1, \hat{\mathbf{w}}^2, \cdots, \hat{\mathbf{w}}^N]^T)$ are mean and covariances of the polynomial coefficients $\hat{\mathbf{w}}^1, \hat{\mathbf{w}}^2, \cdots, \hat{\mathbf{w}}^N$ that are each obtained using the individual gestational weight gain data from each of the N subjects in the public-training data. This distribution over the MLE estimates of the coefficients, $p(\mathbf{w})$ is acquired from the N subjects in the public-training data as an *a-priori* estimate. The likelihood learnt from the individual's personal-training data ($\mathcal{D}$) and the *a-priori* distribution learnt from the population data are then combined using bayes theorem to calculate the maximum-a-posteriori (MAP) estimate of the coefficients $p(\mathbf{w}|\mathcal{D})$.

$$\hat{\mathbf{w}}_{MAP} = \underset{\mathbf{w}}{\operatorname{argmax}}\, p(\mathbf{w}|\mathcal{D}) = \underset{\mathbf{w}}{\operatorname{argmax}} \frac{P(\mathcal{D}|\mathbf{w}) p(\mathbf{w})}{P(\mathcal{D})} \tag{3}$$

The forecast at time t_m^+ is given by $\hat{\mathbf{w}}_{MAP}[t_m^+ \, {t_m^+}^2 \cdots {t_m^+}^p]^T$. This approach is called parametric because the choice of order of the polynomial p depends on the application of interest.

B. Federated approach with eternal updates (F_∞)

Federated learning is the process of storing only the model weights from individual subjects that are pushed to a central server. This preserves the privacy of a subject by only sending the model coefficients instead of complete raw data information as followed in the centralised approach. These small updates of local model coefficients ($\hat{\mathbf{w}}^i$) are sent to the central server where these updates are stored eternally, so that whenever a new model update arrives or a global update is needed all parties can participate and a global model can be aggregated as $\mu_{\hat{\mathbf{w}}} = mean([\hat{\mathbf{w}}^1, \hat{\mathbf{w}}^2, \cdots, \hat{\mathbf{w}}^N]^T)$, $\Sigma_{\hat{\mathbf{w}}} = cov([\hat{\mathbf{w}}^1, \hat{\mathbf{w}}^2, \cdots, \hat{\mathbf{w}}^N]^T)$. The federated learning process (Fig. 1) that we utilised is as follows:-

(1) the centralized server sends the meta-data, (for example, order p of the polynomial, current global model estimate) to the participating subjects, once all subjects agree upon it,
(2) the local subjects estimate model coefficients $\hat{\mathbf{w}}_{MLE}$ based on maximising the likelihood of the local data,
(3) these local model updates are then shared to the server
(4) the server aggregates the individual models and create an updated global model,
(5) the global model is shared with the participating subjects.

This process is repeated as new subjects participate or the already participating subjects gather more data to push updated local models to the server.

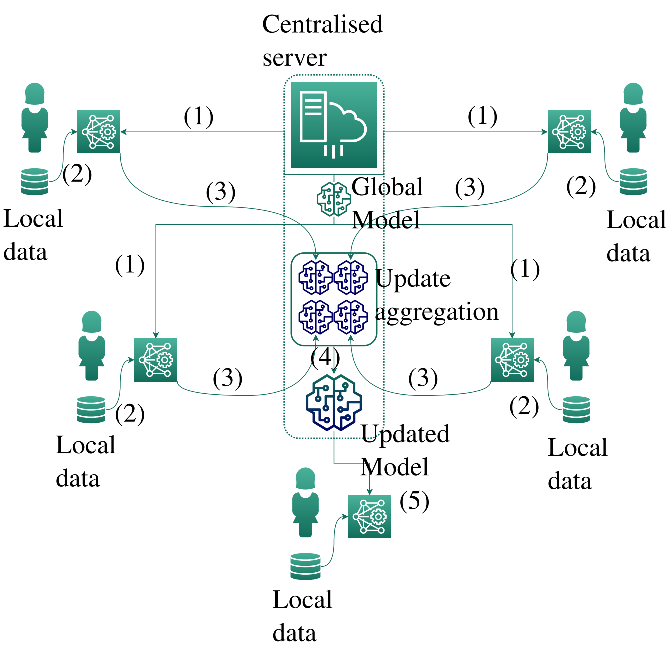

Fig. 1: Federated learning ensures local data remains on-device and only model weights are shared at the central server.

The local updates from participating subjects are stored eternally at the central server for secure aggregation to accurately estimate the global update. Hence, this method is also denoted as F_{∞} as the updates are stored for infinite time.

C. Federated approach with ephemeral updates ($F_{\cancel{\infty}}$)

Although federated learning with eternal updates gives better privacy guarantees than sharing user data and learning on a central server, it still leaves the system vulnerable to attacks from older model updates or models themselves. The reason why stored model updates over time can still reveal sensitive information is because they are derived from the sensitive data of the user and is a representation of high level statistical distribution of the data [6]. We propose a scenario where only incremental updates from participating users are shared and are deleted from the central server once the global model is updated. Assuming a multivariate normal distribution, the global model μ^N, $\Sigma^N_{\hat{\mathbf{w}}}$ can be updated using the past global model μ^{N-1}, $\Sigma^{N-1}_{\hat{\mathbf{w}}}$ and the new shared local model ($\hat{\mathbf{w}}^N$) as follows:-

$$\begin{aligned} \mu^N_{\hat{\mathbf{w}}} &= \frac{(N-1)\mu^{N-1}_{\hat{\mathbf{w}}} + \hat{\mathbf{w}}^N}{N} \\ &= \mu^{N-1}_{\hat{\mathbf{w}}} + \frac{\hat{\mathbf{w}}^N - \mu^{N-1}_{\hat{\mathbf{w}}}}{N} \end{aligned} \tag{4}$$

Similarly, covariance for N^{th} update can be estimated as [2],

$$\Sigma^N_{\hat{\mathbf{w}}} = \Sigma^{N-1}_{\hat{\mathbf{w}}} + \frac{\hat{\mathbf{w}}^N \hat{\mathbf{w}}^{N^\top}}{N-1} - \frac{N \cdot \mu^N_{\hat{\mathbf{w}}} \mu^{N^\top}_{\hat{\mathbf{w}}}}{N-1} + \mu^{N-1}_{\hat{\mathbf{w}}} \mu^{N-1^\top}_{\hat{\mathbf{w}}} \tag{5}$$

V. Experiments

We perform *leave-one-out* cross validation to evaluate and compare performance of our approaches, where training dataset in each iteration consists of weight gain data from $n \leq N$ public-training subjects and self-training data from the test subject. Here, n denotes the number of participants that had already participated in the federated learning pregnancy and an updated global model exists based on these n number of participants. We experiment with different values of n to show the effect of number of initiating users on the regression performance. We subtract the pre-pregnancy weight from the absolute data to get weight-gain data to ensure further local model security. The performance of regression was computed using Mean Absolute Error (MAE), $MAE = \frac{1}{N} \sum_N |y(t^i_m) - y_{ref}(t^i_m)|$. We experiment with first, second, third, fourth and fifth order polynomial based approach to fit our weight-*gain* data. The weight-gain data is normalised to pass through origin, so intercept term can be omitted. We chose third-order polynomial as it obtains minimum prediction error.

VI. Results

Initially, we assume that $n = 10$ random users have already participated in the model building process and we perform leave-one-out cross validation on the rest of 70 subjects by sending a global model learnt based on $n = 10$ subjects as an initial estimate. Fig. 2(a) and 2(b) show the worst and best performing subjects respectively in terms of estimating end-of-pregnancy weight gain based on such a federated learning scheme. Note that personal weight gain data until day 120 is used which is shown in black in Fig. 2 and the further values to be predicted are plotted in green. Fig. 2(c) shows that when a global model initiated by $n = 70$ users is distributed, the regression performance improves. The end-of-pregnancy weight prediction also improves with only $n = 10$ participating users if the personal-data availability increases (Fig. 2(d)).

[2] See Appendix for proof.

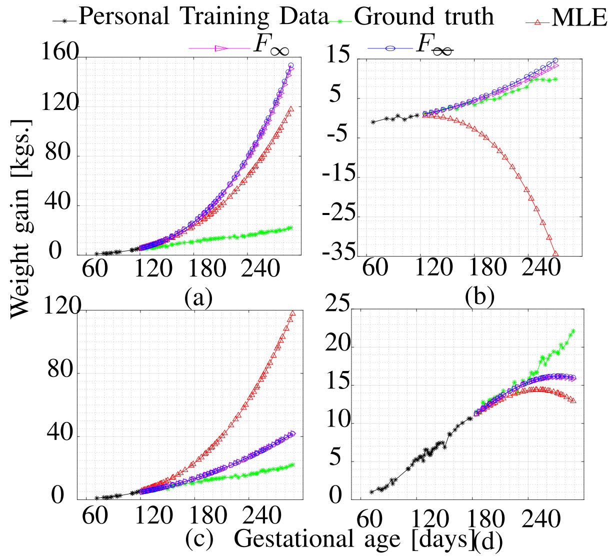

Fig. 2: Federated learning generates (a) worst (subject id #14) and (b) best result (subject id #47) with limited personal data upto 120 days when only 10 users have participated initially. Performance for the subject id #14 can be seen improving when (c) 70 users participated in federated learning or when (d) the availability of personal-data increased (upto 180 days).

Next, we present the prediction results averaged over $N - n$ subjects where n is varied as 10, 40, and 70 and $N = 80$. Fig. 3 shows that performance improves (MAE decreases) as self-training data availability increases or when the initial number of users participating in federated learning increase.

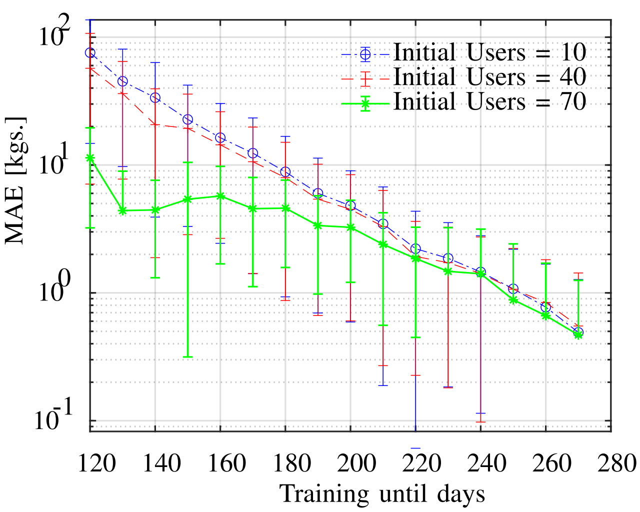

Fig. 3: Average mean absolute error decreases as personal training data increases or number of initial users increase.

The centralised approach with 80 subjects produces the minimum absolute error in prediction with around 2.57 kgs error in predicting end-of-pregnancy weight gain. The federated approach with ephemeral updates (F_{∞}) performs worse by about 1.89 kgs than centralised MAP approach with around 4.46 kgs mean absolute error. Fig. 4 shows that the federated learning out-performs the rest of the state-of-the-arts in predicting gestational weight gain in the presence of limited personal data (upto 200 days).

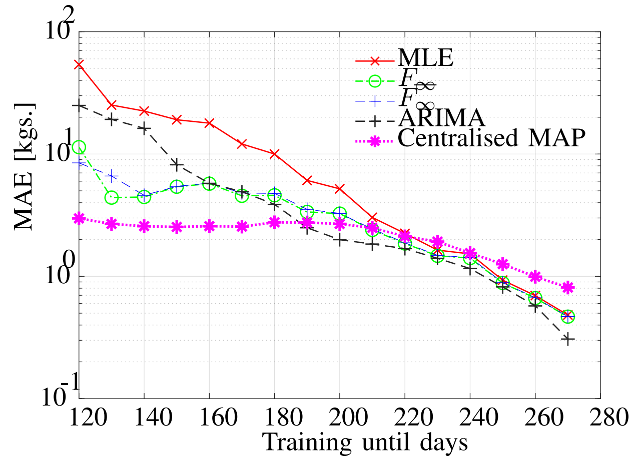

Fig. 4: Performance of federated learning as compared to state-of-the-art approaches.

VII. Discussion

In this paper, we propose the implementation of federated learning in time series related to healthcare. Apart from centralised learning where the raw data is shared and stored at the server for model building, we discuss two different federated scenarios that differ in how long the updates are stored at the central server.

Fig. 3 shows that the performance of the federated learning approach is different when the initial number of users participating in the training process varies. As the number of users that are involved in initial model building increases, the performance improves. This can be attributed to the fact that the global model becomes more generalised when the number of users have increased. Similarly, a decreasing trend is observed in mean absolute error from Fig. 3 and Fig. 4 with respect to the number of training days available. It is intuitive that as more and more training data becomes available, the individual model starts estimating the end-of-the-pregnancy weight more accurately. But, it is desirable to predict the weight gain as early as possible for necessary intervention.

Fig. 4 shows that the performance of the two federated learning approaches with different local model storage strategies have identical performance as the availability of the training data increases. It can be observed that the federated approaches (F_{∞} (green) and F_{∞}(blue)) performance in early prediction of the weight gain is much better than the state-of-the-arts and is very close to the centralised approach, thus guaranteeing a good trade-off in performance and privacy preservation. As more and more training data for an individual pregnancy is available the performance of centralised approach, MLE and the federated learning approaches is close to each other as the global a model a-priori has less influence on local model.

VIII. Conclusion

In this paper, we try and propose a federated learning strategy that enables the preservation of privacy of a user while attaining state-of-the-art performance. We try and predict the gestational weight gain at the end of pregnancy as early as possible. The proposed approach achieves around 4.455 kgs of mean absolute error as early as 140 days into the pregnancy. In the future, we would like to improve upon the privacy of the shared model updates by making them differentially private (adding a noise to local weights) and establishing formal privacy guarantees.

APPENDIX

Proof of Federated Covariance estimation with ephemeral updates:

$$\begin{aligned}
\Sigma_{\hat{\mathbf{w}}}^{N} &= \frac{1}{N-1}\sum_{i=1}^{N}\left(\hat{\mathbf{w}}^{i}-\mu_{\hat{\mathbf{w}}}^{N}\right)\left(\hat{\mathbf{w}}^{i}-\mu_{\hat{\mathbf{w}}}^{N}\right)^{\top} \\
&= \frac{1}{N-1}\sum_{i=1}^{N}\left[\hat{\mathbf{w}}^{i}\hat{\mathbf{w}}^{i^{\top}}-\hat{\mathbf{w}}^{i}\mu_{\hat{\mathbf{w}}}^{N^{\top}}-\mu_{\hat{\mathbf{w}}}^{N}\hat{\mathbf{w}}^{i^{\top}}+\mu_{\hat{\mathbf{w}}}^{N}\mu_{\hat{\mathbf{w}}}^{N^{\top}}\right] \\
&= \frac{1}{N-1}\sum_{i=1}^{N}\hat{\mathbf{w}}^{i}\hat{\mathbf{w}}^{i^{\top}}-2\left(\sum_{i=1}^{N}\hat{\mathbf{w}}^{i}\right)\mu_{\hat{\mathbf{w}}}^{N^{\top}}+\sum_{i=1}^{N}\mu_{\hat{\mathbf{w}}}^{N}\mu_{\hat{\mathbf{w}}}^{N^{\top}} \\
&= \frac{1}{N-1}\sum_{i=1}^{N}\hat{\mathbf{w}}^{i}\hat{\mathbf{w}}^{i^{\top}}-2N\mu_{\hat{\mathbf{w}}}^{N}\mu_{\hat{\mathbf{w}}}^{N^{\top}}+N\mu_{\hat{\mathbf{w}}}^{N}\mu_{\hat{\mathbf{w}}}^{N^{\top}} \\
&= \frac{1}{N-1}\sum_{i=1}^{N}\hat{\mathbf{w}}^{i}\hat{\mathbf{w}}^{i^{\top}}-N\mu_{\hat{\mathbf{w}}}^{N}\mu_{\hat{\mathbf{w}}}^{N^{\top}}
\end{aligned} \tag{6}$$

In order to calculate the update, we use the $\Delta\Sigma = \Sigma_{\hat{\mathbf{w}}}^{N} - \Sigma_{\hat{\mathbf{w}}}^{N-1}$. Substituting eq. 6 to calculate $\Delta\Sigma$, we get

$$\begin{aligned}
\Delta\Sigma &= \Sigma_{\hat{\mathbf{w}}}^{N} - \Sigma_{\hat{\mathbf{w}}}^{N-1} \\
&= \frac{\hat{\mathbf{w}}^{N}\hat{\mathbf{w}}^{N^{T}}}{N-1} - \frac{N\cdot\mu_{\hat{\mathbf{w}}}^{N}\mu_{\hat{\mathbf{w}}}^{N^{T}}}{N-1} + \mu_{\hat{\mathbf{w}}}^{N-1}\mu_{\hat{\mathbf{w}}}^{N-1^{T}}
\end{aligned} \tag{7}$$

ACKNOWLEDGMENT

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No 766139. This publication reflects only the authors' view and the REA is not responsible for any use that may be made of the information it contains.

References

[1] A. Bhowmick, J. Duchi, J. Freudiger, G. Kapoor, and R. Rogers, "Protection against reconstruction and its applications in private federated learning," arXiv preprint arXiv:1812.00984, 2018.

[2] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, Time series analysis: forecasting and control. John Wiley & Sons, 2015.

[3] M. Fredrikson, S. Jha, and T. Ristenpart, "Model inversion attacks that exploit confidence information and basic countermeasures," in Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, 2015, pp. 1322–1333.

[4] R. Gaillard, B. Durmuş, A. Hofman, J. P. Mackenbach, E. A. Steegers, and V. W. Jaddoe, "Risk factors and outcomes of maternal obesity and excessive weight gain during pregnancy," Obesity, vol. 21, no. 5, pp. 1046–1055, 2013.

[5] R. F. Goldstein, S. K. Abell, S. Ranasinha, M. Misso, J. A. Boyle, M. H. Black, N. Li, G. Hu, F. Corrado, L. Rode, et al., "Association of gestational weight gain with maternal and infant outcomes: a systematic review and meta-analysis," Jama, vol. 317, no. 21, pp. 2207–2225, 2017.

[6] L. Lyu, "Privacy-preserving machine learning and data aggregation for internet of things," Ph.D. dissertation, 2018.

[7] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, et al., "Communication-efficient learning of deep networks from decentralized data," arXiv preprint arXiv:1602.05629, 2016.

[8] L. Melis, C. Song, E. De Cristofaro, and V. Shmatikov, "Exploiting unintended feature leakage in collaborative learning," in 2019 IEEE Symposium on Security and Privacy (SP). IEEE, 2019, pp. 691–706.

[9] C. Puri, K. Dolui, G. Kooijman, F. Masculo, S. Van Sambeek, S. D. Boer, S. Michiels, H. Hallez, S. Luca, and B. Vanrumste, "Privacy preserving pregnancy weight gain management: demo abstract," in Proceedings of the 17th Conference on Embedded Networked Sensor Systems, 2019, pp. 398–399.

[10] C. Puri, G. Kooijman, F. Masculo, S. Van Sambeek, S. Den Boer, S. Luca, and B. Vanrumste, "Pregdict: Early prediction of gestational weight gain for pregnancy care," in 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). IEEE, 2019, pp. 4274–4278.

[11] K. M. Rasmussen, P. M. Catalano, and A. L. Yaktine, "New guidelines for weight gain during pregnancy: what obstetrician/gynecologists should know," Current opinion in obstetrics & gynecology, vol. 21, no. 6, p. 521, 2009.