Privacy-preserving pandemic monitoring

Win, Thu Yein; Tianfield, Hugo

Published in:
Data Science Advancements in Pandemic and Outbreak Management

DOI:
[10.4018/978-1-7998-6736-4.ch010](https://doi.org/10.4018/978-1-7998-6736-4.ch010)

Publication date:
2021

Document Version
Publisher's PDF, also known as Version of record

[Link to publication in ResearchOnline](#)

Citation for published version (Harvard):
Win, TY & Tianfield, H 2021, Privacy-preserving pandemic monitoring. in E Asimakopoulou & N Bessis (eds), *Data Science Advancements in Pandemic and Outbreak Management*. IGI Global, pp. 194-206.
<https://doi.org/10.4018/978-1-7998-6736-4.ch010>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please view our takedown policy at <https://edshare.gcu.ac.uk/id/eprint/5179> for details of how to contact us.

Privacy-Preserving Pandemic Monitoring

Thu Yein Win

University of Gloucestershire, UK

Hugo Tianfield

Glasgow Caledonian University, UK

ABSTRACT

The recent COVID-19 pandemic has posed a significant challenge for health organisations around the world in providing treatment and ensuring public health safety. While this has manifested the need for data sharing amongst health organisations, it has also highlighted the importance of ensuring patient data privacy in doing so. This chapter explores the different techniques which facilitate patient data privacy in pandemic monitoring. It also discusses the strengths as well as their limitations, along with possible areas for future research.

INTRODUCTION

The recent COVID-19 pandemic has posed a significant challenge for health organisations around the world in providing treatment and ensuring public health safety. While health professionals have risen up to this challenge, it has also highlighted the limitations amongst the health organisations in being able to detect it in a timely and accurate manner.

Due to the global nature of this pandemic, health organisations have obtained a growing amount of both unstructured and structured patient data which could potentially be leveraged to obtain insights to improve treatment as well as well control its spread. Due to increasing concerns over user data privacy, however, they are not allowed to be shared and stored in a centralised repository to ensure compliance with different data protection regulations (e.g., GDPR). This makes it an obstacle for traditional machine learning framework which requires the use of a centrally-stored data for both training and prediction. Privacy-preserving paradigms are essentially required.

APPROACHES IN EXISTING RESEARCH

The different approaches which have been developed to implement privacy-preserving pandemic monitoring can be categorised in terms of four techniques, which are namely:

1. Differential privacy
2. Federated Learning
3. Social media based approaches
4. Data sharing and access control

Differential privacy

The primary premise of differential privacy involves making sure a data subject is not affected (e.g., not harmed) by their entry or participation in a database, while maximizing utility/data accuracy (as opposed to random/empty outputs) for the queries.

Differential privacy guarantees that: (i) The raw data will not be viewed (and does not need to be modified); (ii) Maintaining the subject's privacy will be valued over mining insights from data; (iii) Resilience to post-processing; post-processing the output of a differentially private algorithm will not affect the differential privacy of the algorithm. In other words, a data analyst that does not have additional knowledge about the database cannot simply increase the privacy loss by thinking about the output of the differential privacy algorithm (Dwork & Roth, 2014).

Figure 1. Differential Privacy overview (Microsoft white paper (2012))

The workflow of differential privacy can be described as follows (as shown in Figure 1):

1. Analyst sends a query to an intermediate piece of software, the differential privacy guard.
2. The guard assesses the privacy impact of the query using a special algorithm.
3. The guard sends the query to the database, and gets back a clean answer based on data that has not been distorted in any way.
4. The guard then adds the appropriate amount of “noise,” scaled to the privacy impact, thus making the answer (hopefully slightly) imprecise in order to protect the confidentiality of the individuals whose information is in the database, and sends the modified response back to the analyst.

Aktay et al (2020) used differential privacy to obtain insights of user behaviour in response to the lockdown measures imposed in response to the pandemic. This involves using the data obtained from Google products, and applying ϵ -differential privacy prior to performing

analytics on them. This enables the researchers to obtain insights into user work patterns and the impact working from home has had on them.

Federated learning

Data science has received broadest attentions both in academia and industry and has generated profound impacts in a wide range of applications.

In a traditionally assumed paradigm, data is available in a centralised point, to which a desired model can thus be applied and trained. In such a paradigm, the effectiveness of the solution lies at the model. However, the assumption of data being available at a centralised point does not necessarily stand. In fact, in most of the real-world problems, data is intrinsically located in distributed sources. Moving data from distributed sources to a centralised point is practically impossible because it is either technically inviable, formidably costly, or would bring forth fundamental privacy risks.

Federated Learning aims to be technically viable, cost-effective and most importantly privacy-preserving data analysis framework over distributed datasets. The basic procedure of federated learning can be stated as follows:

1. Firstly, the distributed nodes each train their local models based on their local data;
2. Secondly, distributed nodes send their locally trained models to the aggregator node;
3. Thirdly, the aggregator node applies a proper mechanism to synchronise the local models so as to coordinate an ensemble model, which represents the global knowledge;
4. Fourthly, the latest ensemble model is deployed to all the distributed nodes to replace their local models;
5. Iterate the above cycle until the system converges.

Assuming that the datasets at the distributed nodes are statistically i.i.d. (independent and identically distributed), then federated learning would converge.

Federated learning is used together with deep learning in detecting COVID-19 infections from Chest X-ray images (Liu et al, 2020). It involves using four different COVID-19 specific deep learning frameworks, namely Covid-Net (Wang, et al., 2020), ResNeXt (Sharma & Muttoo, 2018), MobileNet-v2 (Sandler, et al., 2018) and ResNet (Ayyachamy, et al, 2019) on the publicly available COVIDx dataset containing pneumonia images. Implemented using PyTorch, it achieved an average accuracy of 90.34%.

Social media based approaches

Traditional means of pandemic monitoring involves the use of patient data in training a machine learning model. While the use of patient data in model training allows for a granular analysis of symptoms and thereby improved prediction accuracy, acquiring them in real-time tends to be a significant time lag between proper patient diagnosis and eventual storage on the hospital systems. This further exacerbated by the time taken to anonymise the patient data collected due to user data privacy concerns.

By contrast, online social media data provides a real-time view of the public concerns over the pandemic. In addition, the posts made online usually contain a geographical-related information which allows for pandemic monitoring and prediction in a given geographical location.

Two types of online data are usually used in current pandemic monitoring solutions, i.e.,

1. Online search data;
2. Twitter tweets.

2.1 Online search data

During a pandemic outbreak, there tends to be an increase in online search on pandemic as the public look for information on it as well as possible cure. This is also exacerbated by people who are suffering from it, as they look online for symptoms as well as measures for addressing them. These search queries usually are attached with timestamps, if properly anonymised, making them a useful data source to be used in real-time pandemic monitoring.

Lampos et al (2020) used patient surveys of COVID-19 symptoms together with online search data to track the outbreak of pandemic across eight different countries. This involves first using the Google Health Trends API to extract searches containing specific terms related to the pandemic in a given country and then filtering the relevant ones using the patient symptoms survey obtained from the National Healthcare Service (NHS). To further ensure that the search results obtained are indicative of the actual infections, the proportion of the pandemic coverage in the local media is used as a normalisation measurement. The search data obtained for one country is then used to build a time-series prediction model, before using the model to predict the pandemic status in other countries using transfer learning. Trained using the data obtained from Italy, it was able to forecast pandemic spread in eight countries with high pandemic occurrences over a two-week period.

Lu and Reis (2020) used a similar approach, and found that there is a correlation between an increase in COVID-related online searches and an increase in infection cases. It involves first collecting patient statistics from 32 different countries, before collecting anonymised search results. A correlation model is then developed using the data obtained. It is found that certain COVID-19-specific terms (e.g., “fever”) can be used to predict increases in pandemic cases up to 22.16 days in advance.

A similar approach is used by Yom-Tov et al (2020) to predict regional pandemic outbreaks based on Bing search results. This involves collecting searches on Influenza-Like Illnesses (ILI) for each Upper Tier Local Authority (UTLA). Time-series analysis is then carried out on each UTLA as well as other UTLAs within a 50 km distance to predict regional pandemic spread. Used together with the demographic statistics obtained from the UK Office of National Statistics (ONS), it achieved an Area Under Curve (AUC) accuracy of 63%.

To monitor outbreaks of influenza-like illnesses (ILI) in different countries, Zou et al (2019) built a machine learning model using regularised regression together with transfer learning. Using the influenza outbreak data in the United States, a “source” supervised regularised machine learning model is trained. The model is then transferred to other countries in which limited ground truth training data exists, to monitor outbreaks in those countries. Evaluated

against three different countries, the proposed approach is able to monitor ILI outbreaks with an accuracy of 91.6%.

2.2 Twitter tweets

In addition to anonymised online search queries, another key source for information for pandemic monitoring is social media data, more specifically, micro-blogging services such as Twitter. Due to its ability to post condensed messages and its ability to reach a wide audience, social media enables people to provide latest updates on the status of pandemic spread within a given region. Tweets usually are accompanied with timestamps and geographical locations as well, making them a valuable resource in temporal-based regional pandemic monitoring.

Gencoglu and Gruber (2020) modelled the causality between the user sentiment on Twitter and the characteristics of the pandemic as a means of pandemic monitoring. This involves first extracting daily tweets containing keywords related to the pandemic and then representing them as feature vectors. They are then passed along with different pandemic characteristics (e.g., mortality statistics, infection statistics, etc.) into a Bayesian Network which learns their conditional probabilities for pandemic monitoring. Implemented using publicly available pandemic datasets, the proposed approach is able to achieve an average accuracy of 83.3%.

Dewhurst et al (2020) also used the Twitter dataset to conduct an exploratory analysis on pandemic outbreaks, but combined clustering with time-series model in doing so. Using one percent of the COVID19-related tweets across multiple languages, the proposed approach creates two clusters, one of which monitors discussions on treatments and the other monitoring the collective attention towards the pandemic.

Dewhurst et al (2020) also used the Twitter dataset to conduct an exploratory analysis on pandemic outbreaks, but combined clustering with time-series model in doing so. Using one percent of the COVID19-related tweets across multiple languages, the proposed approach creates two clusters one of which monitors discussions on treatments and the other monitoring the collective attention towards the pandemic.

Zhang et al (2016) used pervasive social network (PSN) together with blockchain to securely share health data in wireless body area networks (WBAN). It first uses an improved version of the IEEE 802.15.6 protocol to establish secure connections between the sensor nodes. It then uses a separate data sharing protocol to exchange healthcare and adding their encrypted versions to the dedicated healthcare blockchain. While the experiment results indicate the proposed approach is able to exchange data securely between the nodes, it did not provide the implementation details of the data sharing protocol.

Data Access and Sharing

Multi-authority ABE (MA-ABE)

Originally proposed by Chase (2007), multi-authority attribute based encryption enables data to be encrypted using a specific set of attributes for each authority to which data is to be sent. It involves the use of a central authority that is responsible for the key management as well as

regulating access control amongst different parties which need access to an object. Figure 2 illustrates how Multi-Authority Attribute Based Encryption works.

Figure 2. Multi-authority Attribute Based Encryption overview

When a Data Owner (**DO**) wishes to make data available for multiple authorities, he/she first registers with them using his/her Global ID (**GID**) such as NI number. Upon receiving the public-private key pair which are specific to each authority, he/she then uses it along with the required access control policy attributes to encrypt data before storing it in a shared data server. When a Data User (**DU**) wishes to access the encrypted data, he/she first sends his/her GID along with the access policy attributes which are encrypted using the public key of the Data Owner. Each of the participating authorities then verifies the policy attributes based on their pre-defined policies. If the access policy contains t out of n required attributes, the decryption key is then sent to the Data User who can then access the encrypted data.

Yu et al (2020) incorporated multi-authority ABE together with blockchain in implementing a tamper-resistant verification scheme. Designed to support access revocation, the proposed approach features the use of two separate blockchains, one of which is used to maintain data and the other used for regulating access control.

Qian et al (2015) used multi-authority ABE in protecting the privacy of personal health records (PHR) in cloud environments using a protocol which generates secret user keys anonymously. It features the use of an access tree for the encrypting user to define fine-grained access policies for different authorities which he/she wants to share data with. To prevent the authorities from collusion, it features the use of an anonymous secret key generation algorithm by running a 2-party secure computation protocol between the user and each participating authority.

While the aforementioned techniques involve the in-device encryption of ciphertext, Policy-Hidden Outsourced ABE (PHOABE) involves outsourcing it to an external party (Belguith et al, 2018). Designed to be used in Internet of Things (IoT) environments, it features the use of a Semi-Trusted Computing Server (STCS) to offload the intensive computation workload. To protect the privacy of user information, it hides the values of the attributes which are used to encrypt the ciphertext together with the access control structure. It is tested on five resource-constrained devices, and the computational costs remain the same with increasing number of encryption attributes.

A similar approach is proposed by Tian et al (2019) to protect the privacy and integrity in auditing user data in fog-to-cloud computing environments. Designed for Internet of Things (IoT) environments, it uses homomorphic encryption to generate data tags at each mobile sink which then sends them to the fog nodes. The fog nodes then generates the corresponding fog tags before sending the encrypted data to the cloud for processing. The third-party auditor (TPA) at the cloud end then performs auditing on the encrypted tags. The experimental results show that it is able to provide privacy-preserving data auditing with minimal performance overhead.

The use of distributed authorities for data encryption and data access helps to address the issues associated with one centralised authority. However, it is possible for participating authorities as well as parties to collude in obtaining the decryption keys as Meamari et al (2020) discussed.

Smart contracts and blockchain

Azaria *et al* (2016) proposed the use of blockchain for secure electronic medical record sharing with access control in their implementation of *MedRec*. Designed to be used by patients as well as researchers, it involves the use of three smart contracts using Ethereum to associate patient health records with different healthcare providers. Upon registration the registrator contracts associate the patient identities with their Ethereum addresses. Throughout their time within the system, the patient-provider contract is responsible for the access management of patient data with the summary contract maintaining a record of access by other parties within the system.

To provide an added protection layer for the secret keys, Guo et al (2019) proposed the use of smart contracts along with multi-authority attribute based access control. This involves the data owner first encrypting the data using his secretly key together with AES (Advanced Encryption System) encryption. Smart contracts are used in this set-up to obtain the different attributes which are provided to the user by different authorities. They are then used to create an access control structure and the user is granted access if the number of correct attributes is beyond a pre-defined threshold.

Blockchain is also used together with ABE to provide anonymous user authentication and multi-party healthcare data access (Guo et al, 2020). It first uses an attribute based multi-signature (ABMS) scheme to encrypt the patient attributes using the keys provided by the attribute authority, before storing the resulting signature on a blockchain. The patient data is encrypted in a similar manner, however, it features the use of a user-defined access control policy in doing so. The healthcare providers in the network then use the aforementioned key to access the patient attributes before using them to access the encrypted data. Implemented

on the Hyperledger Fabric and Hyperledger Ursa platform for the multi-party signature and multi-party ABE respectively, the performance of the approach is found to increase with the increase in patient attributes.

DISCUSSION

In response to the challenges posed by the COVID-19 pandemic, a number of different privacy-preserving data sharing techniques have been proposed to facilitate the sharing and analysis of patient health data. While they have proven to be effective in addressing different aspects of data sharing, a number of issues need further exploration to facilitate widespread adoption.

One of the key aspects of existing approaches is the performance overhead associated with their operations. Federated learning- based approaches reduce the need for a centralised data server by allowing model training to take place at each local node, before sharing the model parameters to create a shared global model. While this allows the use of local data at each node, the communication overhead tends to increase with increase in the number of participating nodes. By the same token, blockchain based approaches also suffer from the same issue in generating the Proof-of-Work (PoW) amongst the participating nodes. Given the heterogeneous nature of participating nodes (e.g., mobile and IoT devices) as well as its reliance on the underlying network infrastructure, the need for a low-latency and efficient network communication protocol is an open research problem which is worth exploring.

The second aspect of existing approaches is the guarantee provided by existing techniques in ensuring user data privacy. MA-ABE based approaches allow for a means of exchanging patient data amongst different participating authorities in a privacy-preserving manner. Similar to federated learning based approaches, this reduces the need for a central authority which regulates access control by distributing the responsibility amongst the authorities. However, the possibility of collusion attacks exists in which either participating authorities or parties can collude to obtain the access control attributes needed to decrypt patient data.

CONCLUSION

The global nature of the current pandemic has increased the need for international collaboration in their efforts to addressing it. To ensure compliance with different international regulations on patient and user data privacy (e.g., HIPAA, GDPR), a number of techniques have been proposed and used in addressing the pandemic. This chapter presents these techniques, along with their overall implementations. It also discusses the strengths as well as their limitations, along with possible areas for future research.

REFERENCES

- Ayyachamy, S., Alex, V., Khened, M. and Krishnamurthi, G., 2019. *Medical image retrieval using Resnet-18*. s.l., s.n.
- Azaria, A. a. E. A. a. V. T. a. L. A., 2016. *Medrec: Using blockchain for medical data access and permission management*. s.l., IEEE, pp. pp. 25--30.

- Belguith, S. et al, 2018. Phoabe: Securely outsourcing multi-authority attribute based encryption with policy hidden for cloud assisted iot. *Computer Networks*, Volume 133, pp. pp. 141-156.
- Chase, M., 2007. Multi-authority attribute based encryption. In: *Theory of cryptography conference*. s.l.:Springer, pp. pp. 515--534.
- Dewhurst, D. R. et al, 2020. Divergent modes of online collective attention to the COVID-19 pandemic are associated with future caseload variance. *arXiv preprint arXiv:2004.03516*.
- Dodds, P. S. et al, 2020. Long-term word frequency dynamics derived from Twitter are corrupted: A bespoke approach to detecting and removing pathologies in ensembles of time series. *arXiv preprint arXiv:2008.11305*.
- Gencoglu, O. and Gruber, M., 2020. Causal Modeling of Twitter Activity During COVID-19. *arXiv preprint arXiv:2005.07952*.
- Guo, H., Meamari, E. and Shen, C.-C., 2019. *Multi-authority attribute-based access control with smart contract*. s.l., s.n.
- Kumar, R. et al, 2020. Blockchain-federated-learning and deep learning models for covid-19 detection using ct imaging. *arXiv preprint arXiv:2007.06537*.
- Lamos, V. et al, 2020. Tracking COVID-19 using online search. *arXiv preprint arXiv:2003.08086*.
- Liu, B. et al, 2020. Experiments of federated learning for covid-19 chest x-ray images. *arXiv preprint arXiv:2007.05592*.
- Meamari, E., Guo, H., Shen, C.-C. and Hur, J., 2020. Collusion Attacks on Decentralized Attributed-Based Encryption: Analyses and a Solution. *arXiv preprint arXiv:2002.07811*.
- Microsoft white paper (2012), Differential Privacy for Everyone, http://download.microsoft.com/download/D/1/F/D1F0DFF5-8BA9-4BDF-8924-7816932F6825/Differential_Privacy_for_Everyone.pdf
- Qian, H., Li, J., Zhang, Y. and Han, J., 2015. Privacy-preserving personal health record using multi-authority attribute-based encryption with revocation. *International Journal of Information Security*, pp. pp. 487--497.
- Sandler, M. et al, 2018. *Mobilenetv2: Inverted residuals and linear bottlenecks*. s.l., s.n., pp. pp. 4510--4520.
- Sharma, A. and Muttou, S. K., 2018. Spatial image steganalysis based on resnext. In: *2018 IEEE 18th International Conference on Communication Technology (ICCT)*. s.l.:IEEE, pp. pp.1213--1216.
- Tian, H. et al, 2019. Privacy-preserving public auditing for secure data storage in fog-to-cloud computing. *Journal of Network and Computer Applications*, Volume 127, pp. pp. 59--69.

Wang, L., Lin, Z. Q. and Wong, A., 2020. Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images. *Scientific Reports*, Volume 10, pp. pp. 1--12.

Yom-Tov, E., Lamos, V., Cox, I. J. and Edelstein, M., 2020. Providing early indication of regional anomalies in COVID19 case counts in England using search engine queries. *arXiv preprint arXiv:2007.11821*.

Yu, G. et al, 2020. Enabling Attribute Revocation for Fine-Grained Access Control in Blockchain-IoT Systems. *IEEE Transactions on Engineering Management*.

Zhang, J., Xue, N. and Huang, X., 2016. A secure system for pervasive social network-based healthcare. *IEEE Access*, Volume 4, pp. pp. 9239-9250.