




Spatial K-anonymity: A Privacy-preserving Method for COVID-19 Related Geo-spatial Technologies

Rohan Iyer¹^a, Regina Rex¹^b, Kevin P. McPherson¹^c, Darshan Gandhi¹, Aryan Mahindra¹,
Abhishek Singh^{1,2} and Ramesh Raskar^{1,2}

¹PathCheck Foundation, Cambridge, U.S.A.

²MIT Media Lab, MIT, Cambridge, U.S.A.


Keywords: Spatial K-anonymity, Contact Tracing, Spatial Privacy.


Abstract: There is a growing need for spatial privacy considerations in the many geo-spatial technologies that have been created as solutions for COVID-19-related issues. Although effective geo-spatial technologies have already been rolled out, most have significantly sacrificed privacy for utility. In this paper, we explore spatial k-anonymity, a privacy-preserving method that can address this unnecessary tradeoff by providing the best of both privacy and utility. After evaluating its past implications in geo-spatial use cases, we propose applications of spatial k-anonymity in the data sharing and managing of COVID-19 contact tracing technologies as well as heat maps showing a user's travel history. We then justify our propositions by comparing spatial k-anonymity with several other spatial privacy methods, including differential privacy, geo-indistinguishability, and manual consent based redaction. Our hope is to raise awareness of the ever-growing risks associated with spatial privacy and how they can be solved with Spatial K-anonymity.


1 INTRODUCTION

1.1 Motivation

The COVID-19 pandemic has presented a multitude of challenges, from public health to economic instability. Solutions for COVID-19 contact tracing and other purposes involve an exponential rise in the use of geo-location technologies and location-based services at every level of society which has presented its own set of ethical challenges and adoption issues (Cole, 2020; Hassandoust et al., 2020). Each individual's physical location and location history can be exploited by cybercriminals interested in mining personal information such as home, work, and school addresses as well as a daily itinerary, among others. The threat of leakage of information gets even more amplified in these scenarios since it also reflects the person's infection status. Spatial privacy risks may extend beyond the individual, compromising families, coworkers, communities, and even enterprises. Appropriate measures are necessary to ensure spatial privacy is maintained in all solutions for the COVID-19 pandemic.

^a <https://orcid.org/0000-0003-3055-9377>

^b <https://orcid.org/0000-0003-0327-8916>

^c <https://orcid.org/0000-0003-4977-8537>

1.2 Contribution

This paper presents a developmental approach to addressing contact tracing spatial privacy concerns using Spatial K-anonymity (Ghinita et al., 2010). Spatial K-anonymity has been implemented in many pre-COVID-19 use cases (Martin et al., 2015; Sweeney, 2002; Allshouse et al., 2010; Gkoulalas-Divanis et al., 2010); here, we discuss the unique applications that have emerged for spatial K-anonymity for the COVID-19 pandemic. We also compare the spatial K-anonymity approach with other existing solutions like Geo-indistinguishability and Consent Redaction to demonstrate the advantage of utilizing K-anonymity in COVID-19 contact tracing.

2 RELATED WORK

Before the COVID-19 outbreak, concerns about spatial data collection have been echoed worldwide regarding individual security and privacy. Due to these concerns, several policies, like the European Union's General Data Protection Regulation (GDPR) (Council of European Union, 2014), outline identifying location data as personal, requiring other parties to anonymize the collected data or gain explicit consent

from the user for spatial data collection.

In light of current issues, spatial data plays an important role in surveillance, data sharing, and digital contact tracing as used by various government (Park et al., 2020; Jacob and Lawarée, 2020; Jian et al., 2020) and corporate entities (Inc.,) in their efforts to combat COVID-19.

Contact Tracing (Lo and Sim, 2020) involves data collection to help public health officials better control and understand an infectious outbreak, and it also informs individuals about their probable exposures to other infected people. Manual Contact Tracing (MCT) involves caseworkers working with infected persons to trace who they have been in close contact with during their infection period. MCT is laborious, and it relies on the infected person's ability to accurately recall people there were in close contact with (Barrat et al., 2020). It is also resource-intensive and not scalable; if many people are infected, each infected person would need to be assigned a caseworker. Digital Contact Tracing (DCT) is used to address the shortcomings of MCTs. It involves using technological devices such as smartphones and smartwatches to estimate close-range proximity. It is more reliable than MCT, as it is more accurate in sensing people in close proximity (Smieszek et al., 2016). Despite the effectiveness and efficiency of DCT (hbr, 2020), there has been a lot of privacy concerns regarding spatial data collection (Rowe, 2020; Boulos et al., 2009).

In the privacy domain, there have been efforts to ensure user spatial privacy rights are respected. State-of-the-art methods like k-anonymity (Gkoulalas-Divanis et al., 2010), l-diversity (Machanavajjhala et al., 2006), and t-closeness (Li et al., 2007) are some of the few that are applied to various technologies that utilize spatial data. K-anonymity assures that any query on data will give limited access and will maintain the anonymity of at least k individuals. This is usually achieved by using the concepts of purging or generalization. l-diversity covers for the cons of K-anonymity by ensuring that there is significant attribute diversity between equivalence groups. However, it is difficult to accomplish with data having numerous Quasi Identifiers. t-closeness solves this problem by ensuring that the distribution of each sensitive attribute in an equivalence group is close to the distribution in the complete dataset. Thus, these methods help in maintaining user privacy immensely.

Many spatial data protection measures have also been used to safeguard users' spatial privacy rights in various COVID-19 contact tracing applications. Some of the privacy measures applied as outlined in (Wirth et al., 2020) involve pseudonymization or anonymization of spatial data, aggregation of geo-

spatial and temporal data, data minimization, transparency, explicit consent, and ability to revoke consent by the users.

BlueTrace (Bay et al., 2020) protects the users' privacy by using local storage for geo-location history and through the users' explicit consent to data location. They also provide the ability for a user to recount their given consent. Another contact tracing application, Epione (Trieu et al., 2020), guarantees more robust privacy protection using a cryptographic protocol that securely delivers users' status to health authorities and ensures that users' contact information is not revealed to other non-desired parties. These are just some of the methods used to guarantee users' spatial data protection in digital contact tracing. A more unusual approach to this problem is TrackCovid (Yasaka et al., 2020), which avoids handling personal location data for contact tracing and uses anonymization via QR codes at specified checkpoints. (Chan et al., 2020) highlights possible privacy issues that can arise from the various methods of COVID-19 related contact tracing and suggests a Bluetooth pseudonymization approach to combat these privacy issues.

3 COVID-19-BASED APPLICATIONS OF SPATIAL K-ANONYMITY

3.1 Contact Tracing

Location-based privacy is growing in popularity due to the data collection mechanisms of smartphones and other location-agnostic devices. During the COVID-19 pandemic, spatial privacy has proven to be of paramount concern to government officials, public health experts, and other medical professionals. In some countries, it is predicted that because of the consequences of sharing contact-tracing data, over 70% of patients who have contracted COVID-19 were at risk of being identified by location, gender and/or age (Jung et al., 2020).

GPS provides an advantageous solution to contact tracing than other solutions, given other technologies, such as outbreak response tools or symptom tracking tools. Particularly, GPS based solutions can transmit data in background mode via smartphones and wearable devices, provide public space and infection proximity data, and have cross-platform functionality inherent to the data. When digital contact tracing employs a GPS-based solution for tracing individuals, it often queries highly specific geo-location data of a user. This data is quite sensitive as it tracks

the real time location of a user and stores it onto a database that could be potentially used for tracking and tracing purposes. Spatial k-anonymity employs certain methodologies that convert highly accurate geo-location data into a more granular form while preserving the effectiveness of location-based contact tracing.

Imagine person X queries location-based data from address Y. An attacker may use this information to associate person X's real-time location with high accuracy. Spatial k-anonymity solves this problem by employing methods that ensure user X's location-based query has a probability larger than $1/k$, where k is a defined anonymity requirement.

A widely used method to calculate k involves masking a user's location by deploying specific geomasking techniques and relocating the user to a new buffer location. Geomasking is a general category of methods that is used to mask or change the geographical location of a user or patient in an unpredictable way that maintains that user or patient's relationship to a particular geo-location and their status as it concerns a disease infection or outbreak (Stinchcomb, 2004). The k-anonymity value is the number of all potential residential locations inside this buffer area, including the original user location. The value of $1/k$ is thus a measure of the person's disclosure risk in question (the probability of being re-identified) after geomasking. The smaller the $1/k$, the more difficult it is to re-identify the person's real location.

3.2 Heat Map

The incorporation of heatmaps is primarily for better visualization and understanding of the travel history of the user over a period of time. With the help of the heatmap feature the user can realise the travel history and make optimum decisions while planning their next travel journey.

Case 1: Restricted to an Individual User. In this case the user can get detailed information about their travel history and they can use this information to further make their travel plans.

The user can map a circle or any arbitrary shape to distinguish the regions as a public or private one. Based on the segregation, the user can be shown only their travel history in the public region. This would help them to understand their travelling trends. Also, there can be potentially marked several points of danger and caution on the map and accordingly notify the user if they have crossed paths during their journey. A track of the last 7-14 days should be kept ideally since that would help to keep track of the exposure period.

The primary objective of this is to keep the user

aware of their surroundings and help them make concrete decisions while making their next travel plan. The first option should be the most suited and feasible as well.

Case 2: Combination of All the Users Data. In this case, the methodology is to club the details of all the users around a specific user and with the help of that a combined heat map can be plotted. This would potentially help the user get a good understanding of their surroundings. Again, classification of zones can be implemented here with the RED region as a danger zone and alert the user if they are travelling by that particular area.

Many of the new metrics are based on uncertainty. In this paper here the authors suggested a test how easily and quickly the location of the user can be traced and identified from the anonymous collection of users. It was observed that the major focus was on the path a user might have taken. Thus, one of the key parameters metric to track would be to identify the size of the cluster, larger the size, larger the ambiguity and more difficult to track the individual user.

A few suggestions would be to showcase the exact coordinates as street or even the town/ city coordinates, this would add a high amount of ambiguity in the process of identifying the exact location of the user thus helping them stay protected and preserve their privacy.

4 COMPARING SPATIAL K-ANONYMITY WITH OTHER PRIVACY-PRESERVING METHODS

In this section we will give a detailed description of the various spatial privacy approaches: K-anonymity, geo-indistinguishability, and consent redaction. We will also show how these approaches differ from each other and spatial k-anonymity, highlighting their strengths and weaknesses in digital contact tracing and other COVID-19 related applications.

K-anonymity (Gkoulalas-Divanis et al., 2010) is a well known anonymization technique for spatial privacy. This technique refers to the process of transforming a data set that contains personal or identifying information to a dataset that strips away any information that could identify an individual, making it hard for other parties to determine the identity of an individual from the transformed data set. In this paper, we will be focusing on the K-Anonymity algorithm for protecting spatial privacy. K-anonymity is characterized as the number of individuals, which

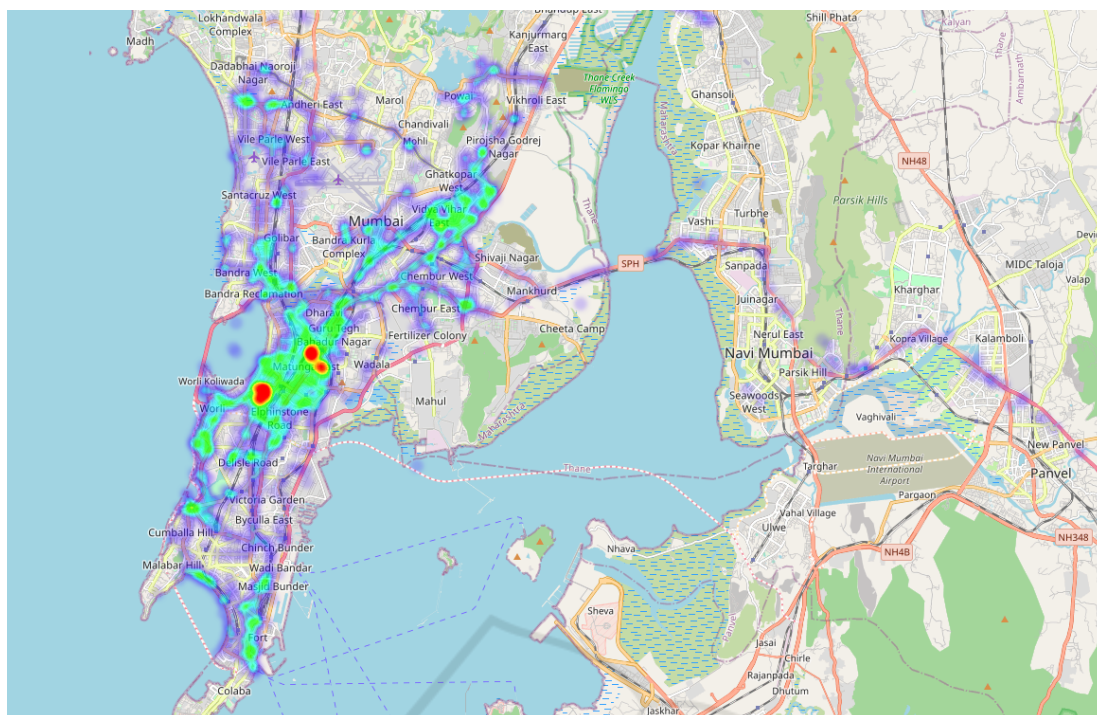


Figure 1: An example of a user’s location history represented as a heat map (Red-orange → heavily visited, Green → regularly visited, Blue → least visited).

is a distinct value ‘k’, associated with alike attributes like gender, city, etc such that one particular individual cannot be identified from the data set. Since the dataset consists of geographical identifiers as well, the traditional k-anonymity algorithm can be expanded to spatial k-anonymity. Spatial k-anonymity is a widely used method to evaluate the degree of geoprivacy obtained following the implementation of a particular geo masking technique.

Spatial k-anonymity works the best for conventional geo-spatial datasets wherein only one geographical area is identified for each individual, however, it has important shortcomings when dealing with individual GPS datasets. Spatial K-anonymity algorithm is highly regarded since rather than having a single geographic location that can disclose a person’s identity, it provided multiple several everyday activity locations, for example, home, offices, and shopping venues that can be easily identified from a GPS dataset, which makes it difficult to zero down a single person. Having data of a single location associated with a person can make an individual vulnerable by providing a wealth of geographic data to hackers and can be utilized for similar malpractices. However, since most people spend only a small amount of time at home and spend a significant amount of time away from home, having multiple locations increases the scope of user privacy. It is important to consider

the disclosure risk associated with all daily activity locations in a detailed manner.

Geo-indistinguishability (Andrés et al., 2013) is a metric used for location-based systems that protect the user’s exact location while allowing approximate information, typically needed to obtain a specific desired service, to be released. The underlying motivation for geo-indistinguishability comes from differential privacy but more handcrafted for the spatial data. It is focused on protecting the user’s location within a radius, with a level of privacy that depends on that corresponds to a generalized version of the well-known concept of differential privacy. The value must be optimized in consideration of both privacy and utility, and it must be large enough to fit the notion of privacy. That is, an attacker with potential auxiliary information should not be able to pinpoint the user’s identity. For example, must also be small enough to allow digital contact tracing to accurately identify users who have been in contact with a user who has tested positive. Like spatial K-anonymity, geo-indistinguishability guarantees spatial privacy needs for location-based systems such as digital contact tracing apps. (Vepakomma et al., 2020) uses differential privacy to release a sketching data structure of location data. Aside from this, there have not been any application of geo-indistinguishability in digital contact tracing.

Manual Consent Redaction is a privacy-preserving methodology that enables users to revoke given consent to data collection for specific locations or all locations tracked by geo-spatial technologies. This technique offers users greater control over their data, as they can choose to opt-in or out of the service or select what data points to exclude from external servers. In manual consent redaction, the user can choose to revoke their consent for geo-spatial technologies to access all their locations or just for specific locations. If the user desires the latter, they must manually select which locations to exclude from the data collection process, and these would also be removed from the external servers. Choosing the former means the user can not use the geo-spatial technology, and all their data should be removed from the servers. This approach gives user's control over their privacy and is a highly interpretable notion of privacy from a user standpoint. Its disadvantages includes the risk of human error and it has less privacy-utility tradeoff.

Most contact tracing apps only give the users the option of fully consenting to data collection or revoking their consent. Using this approach, the user can only use the contact tracing app if they agree to data collection. Otherwise, they can not specify if they want only specific locations redacted. Unlike the other privacy-preserving methods above, manual consent redaction alone does not fully guarantee spatial privacy. It is more restrictive, as revoking consent bars a user from using the technology. For example, (Bay et al., 2020) uses consent revocation in combination with other privacy-preserving methods to ensure users' privacy. The more user-friendly approach, where users can exclude specific data points if they choose to, does not guarantee the remaining data collected are non-identifying. i.e., attackers can still identify users based on their DAL even if they redact private locations from data collection.

5 DEVELOPMENTAL SOLUTION

The idea of a spatial k-anonymity development solution is that it takes a list of user's locations (L), and creates a bound portion of anonymized geographical locations (anonymized-L) that is shared with at least K-1 other users.

Pulling heavily on the previous work of (Cassa et al., 2006), which addresses the issues of using an anonymizing algorithm while preserving the fidelity of containing an outbreak through epidemiology data, the features of our algorithm are as follows:

1. **Population Density-based Gaussian Spatial Skew.** We hope to "blur" any location the user deems as private. As such, we need to take the user's addresses that are private (represented in longitude and latitude), and skew using a random offset based on Gaussian distribution which has a standard deviation that is inversely proportional to the area's population density. This allows for smaller transposition of people in high density areas while maintaining k-anonymity.

- **Individual Point Gaussian Blur.** Points will have to be skewed a minimum distance. To do this, we develop two randomly selected variables, σ_x and σ_y which represent standard deviations from a normal distribution to give the distance and direction of a user's displacement. From here, two displacement values $((d_x, d_x))$. Gaussian blur, a concept used in image processing, takes a float or integer representation and applies noise in line with a Gaussian distribution.

2. **Combining with the Anonymization Algorithm of Cassa et al.** Assuming we have some context data like US Census or other state-sponsored data collection, we want anonymization skewed inversely with population density. That is, if someone lives in a highly-dense area, they should move a shorter distance in the anonymized mode, whereas rural participants move larger distances in the anonymized mode.

3. **Anonymization Multipliers.** In order to control for population and age-based density issues, multipliers from the Cassa et al. paper will be employed. They are as follows:
 APDM = Age-based population density multiplier
 TPDM = Total population density multiplier
 AGPD = Average age group population density
 UBGAD = User's block group age density
 ATPD = Average total population density
 UBGPD = User's block group population density
 AM = Age Multiplier
 APD = Age-based population density
 CM = Combined multiplier

$$APDM = AGPD/UBGAD \quad (1)$$

$$TPDM = ATPD/UBGPD \quad (2)$$

Equations (1) and (2) combine to form a combined multiplier that looks like this:

$$CM = AM * APD + (1 - AM) * TPDM \quad (3)$$

The overall Anonymization Multiplier is as follows, with c being a scaling factor that can alter the overall skew of a coordinate pair:

$$CM = AM * APDM + (1 - AM) * TPDM \quad (4)$$

These solutions should prove to hold up to traditional spatial K-anonymity privacy standards that any nefarious effort to identify the user's location can be pinpointed with no greater than $1/K$ probability (Ghinita et al., 2010).

6 TECHNICAL SOLUTION AND DISCUSSION

In order to exemplify our developmental solution, a technical solution is implemented in Python using an SQLite database, as one might find in a mobile application to hold a small dataset of 50 locations in the Government Center neighborhood of Boston, Massachusetts. The application of the code can be found on Github in a Jupyter Notebook (McPherson, 2020)

In summation, the approach follows a similar production-level approach one might find in a mobile application: starting with an SQLite database, the data is transformed into a dataframe, while each longitude and latitude are Gaussian blurred and transformed from the estimates above.

To test the hypothesis that the latitude and longitude coordinates were skewed in a Gaussian manner, a Shapiro-Wilk test is provided at the end of the notebook, as provided by the `scipy` library. The Shapiro-Wilk test estimated a p-value greater than 0.05 (0.118) for latitudinal values, meaning that the null hypothesis cannot be rejected. We assume then that the latitudinal values are transformed in a Gaussian manner. However, the Shapiro-Wilk test concluded that the null hypothesis for longitudinal values should be rejected ($p=0.00$). However, it can be noted that there appears to be an outlier in the distribution, which probably affected the transformation (see In[30] in notebook).

Finally, it is illustrated that, when following the probability distribution of a Gaussian distribution, the estimated k-anonymity (see In[32] in notebook) increases non-monotonically by the equation of:

Estimated k-anonymity = location population density * area of neighborhood * probability in Gaussian standard distribution.

This allows for k-anonymity to increase nonlinearly even as the standard distribution increases linearly. This, paired with the factors that are context and geographically specific, allow k-anonymity

to vary in magnitude across space and time. Although knowledge of various characteristic values like population density likely would compromise the algorithm in some way, the algorithm randomizes the magnitude of skew by inverse variation to other personal and hard-to-obtain privacy data like patient age and exact location.

One main limitation of our study is the number of data points our team was able to obtain from various census blocks, as well as those census block's characteristic information regarding age and population density. A follow-up study should be performed on the power of the data in regards to its normality test, as well as obtaining more data points to be able to apply the algorithm to actual confirmed COVID-19 patients.

7 CONCLUSIONS

Currently, we have built an application to demonstrate the concept of spatial privacy protection of the user. For the same we managed to build a working front end and back end product. In future, we plan to up-scale this and make it production ready to be deployed for real life use cases. For the same we plan to shift from using SQLite as our database to others such as Postgres for handling heavy load. Also, we plan to add several new features such as heatmap generation for the user to help them understand their travel patterns, mark points/places of heavy concentration of users and declare it to be a zone of high risk, also identify and mark the points which are public places. All these features would help the user to plan their travel journey better and also help understand the surroundings they reside in a better way.

ACKNOWLEDGEMENTS

Special thanks to PathCheck Foundation for supporting our paper, including Ramesh Raskar, Aryan Mahindra, and Haris Nazir.

REFERENCES

- (2020). how digital contact tracing slowed covid-19 in east asia-2020.
- Allshouse, W. B., Fitch, M. K., Hampton, K. H., Gesink, D. C., Doherty, I. A., Leone, P. A., Serre, M. L., and Miller, W. C. (2010). Geomasking sensitive health data and privacy protection: an evaluation using an

- e911 database. *Geocarto International*, 25(6):443–452.
- Andrés, M. E., Bordenabe, N. E., Chatzikokolakis, K., and Palamidessi, C. (2013). Geo-indistinguishability: differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security, CCS '13*, pages 901–914, New York, NY, USA. Association for Computing Machinery.
- Barrat, A., Cattuto, C., Kivelä, M., Lehmann, S., and Saramäki, J. (2020). Effect of manual and digital contact tracing on COVID-19 outbreaks: a study on empirical contact data.
- Bay, J., Kek, J., Tan, A., and Hau, C. S. (2020). Bluetrace: A privacy-preserving protocol for community-driven contact tracing across borders.
- Boulos, M., Curtis, A. J., and AbdelMalik, P. (2009). Musings on privacy issues in health research involving disaggregate geographic data about individuals. *International Journal of Health Geographics*, 8(1):46.
- Cassa, C. A., Grannis, S. J., Overhage, J. M., and Mandl, K. D. (2006). A context-sensitive approach to anonymizing spatial surveillance data: Impact on outbreak detection. *Journal of the American Medical Informatics Association*, 13(2):160–165.
- Chan, J., Foster, D., Gollakota, S., Horvitz, E., Jaeger, J., Kakade, S., Kohno, T., Langford, J., Larson, J., Sharma, P., Singanamalla, S., Sunshine, J., and Tessaro, S. (2020). Pact: Privacy sensitive protocols and mechanisms for mobile contact tracing.
- Cole, C. (2020). The safety of privacy: Increased privacy concerns may prevent effective adoption of contact tracing apps.
- Council of European Union (2014). Council regulation (EU) no 269/2014.
- Ghinita, G., Zhao, K., Papadias, D., and Kalnis, P. (2010). A reciprocal framework for spatial k-anonymity. *Information Systems*, 35(3):299–314.
- Gkoulalas-Divanis, A., Kalnis, P., and Verykios, V. S. (2010). Providing k-anonymity in location based services. *ACM SIGKDD Explorations Newsletter*, 12(1):3–10.
- Hassandoust, F., Akhlaghpour, S., and Johnston, A. C. (2020). Individuals' privacy concerns and adoption of contact tracing mobile applications in a pandemic: A situational privacy calculus perspective. *Journal of the American Medical Informatics Association*.
- Inc., A. Apple and google partner on covid-19 contact tracing technology.
- Jacob, S. and Lawarée, J. (2020). The adoption of contact tracing applications of COVID-19 by european governments. *Policy Design and Practice*, pages 1–15.
- Jian, S.-W., Cheng, H.-Y., Huang, X.-T., and Liu, D.-P. (2020). Contact tracing with digital assistance in taiwan's COVID-19 outbreak response. *International Journal of Infectious Diseases*, 101:348–352.
- Jung, G., Lee, H., Kim, A., and Lee, U. (2020). Too much information: Assessing privacy risks of contact trace data disclosure on people with COVID-19 in south korea. *Frontiers in Public Health*, 8.
- Li, N., Li, T., and Venkatasubramanian, S. (2007). t-closeness: Privacy beyond k-anonymity and l-diversity. In *2007 IEEE 23rd International Conference on Data Engineering*, pages 106–115.
- Lo, B. and Sim, I. (2020). Ethical framework for assessing manual and digital contact tracing for COVID-19. *Annals of Internal Medicine*.
- Machanavajjhala, A., Gehrke, J., Kifer, D., and Venkatasubramanian, M. (2006). L-diversity: privacy beyond k-anonymity. In *22nd International Conference on Data Engineering (ICDE'06)*, pages 24–24.
- Martin, M. K., Helm, J., and Patyk, K. A. (2015). An approach for de-identification of point locations of live-stock premises for further use in disease spread modeling. *Preventive Veterinary Medicine*, 120(2):131–140.
- McPherson, K. P. (2020). Application of gaussian skew to privacy model. <https://github.com/PrivateKit/spatialprivacy>. Accessed Feb 22, 2021.
- Park, Y. J., Choe, Y. J., Park, O., Park, S. Y., Kim, Y.-M., Kim, J., Kweon, S., Woo, Y., Gwack, J., Kim, S. S., Lee, J., Hyun, J., Ryu, B., Jang, Y. S., Kim, H., Shin, S. H., Yi, S., Lee, S., Kim, H. K., Lee, H., Jin, Y., Park, E., Choi, S. W., Kim, M., Song, J., Choi, S. W., Kim, D., Jeon, B.-H., Yoo, H., and and, E. K. J. (2020). Contact tracing during coronavirus disease outbreak, south korea, 2020. *Emerging Infectious Diseases*, 26(10):2465–2468.
- Rowe, F. (2020). Contact tracing apps and values dilemmas: A privacy paradox in a neo-liberal world. *International Journal of Information Management*, 55:102178.
- Smieszek, T., Castell, S., Barrat, A., Cattuto, C., White, P. J., and Krause, G. (2016). Contact diaries versus wearable proximity sensors in measuring contact patterns at a conference: method comparison and participants' attitudes. *BMC Infectious Diseases*, 16(1).
- Stinchcomb, D. (2004). Procedures for geomasking to protect patient confidentiality.
- Sweeney, L. (2002). Achieving k-nonymity Privacy Protection Using Generalization and Suppression. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(05):571–588.
- Trieu, N., Shehata, K., Saxena, P., Shokri, R., and Song, D. (2020). Epione: Lightweight contact tracing with strong privacy.
- Vepakomma, P., Pushpita, S. N., and Raskar, R. (2020). Dams meta-estimation of private sketch data structures for differentially private contact tracing.
- Wirth, F. N., Johns, M., Meurers, T., and Prasser, F. (2020). Citizen-centered mobile health apps collecting individual-level spatial data for infectious disease management: Scoping review. *JMIR Mhealth Uhealth*, 8(11):e22594.
- Yasaka, T. M., Lehrich, B. M., and Sahyouni, R. (2020). Peer-to-peer contact tracing: Development of a privacy-preserving smartphone app. *JMIR mHealth and uHealth*, 8(4):e18936.