ETHICAL ANALYSIS OF GEOSOCIAL DATA TO BALANCE SOCIAL AND INDIVIDUAL INTERESTS

D. Burghardt and A. Dunkel

Institute of Cartography, TU Dresden, Germany
firstname.lastname@tu-dresden.de

Keywords: ethical cartography, geosocial media, geovisual analysis, privacy protection, HyperLogLog

Introduction

Geosocial media platforms enable sharing of information on location-based events and spatial phenomena. In addition to the discourse on socially relevant topics, personal views and preferences are shared equally. Research and use of this data in the interest of society and the individuals requires a responsible, ethical approach (Zhang et al., 2022). According to Olteanu et al. (2019) an ethical approach means that individual autonomy is respected, that research should be beneficial and non-maleficence and should aim on the ideal of justice.

Individual autonomy can be ensured by a declaration of informed consent. With active, purpose-oriented participation (e.g. citizen science projects, hash tag campaigns), this is easier to achieve, but more difficult if not impossible when evaluating millions of social media posts. Even if users make data publicly available (also referred to as the "public data" argument) and agree to the terms of service that the data will be used by third parties, Williams et al. (2017) calls for researchers to consider user expectations, the impact of context collapse and the functioning of algorithms with combining potentially sensitive personal data.

Context collapse are described by Crawford and Finn (2015) on the example of tweeting information about location, food, water needs, personal well-being and health status of friends and family in a crisis situation, thereby opening the risk that this data could be used discriminatory in areas such as employment, property and health insurance. The perception of privacy cannot be clearly defined, but depends on the context and shifts with the circumstances under which personal information is made available (Nissenbaum, 2010). Furthermore, an “aggregation effect” (Solove, 2012) relates to the combination of different data sources, which can lead to privacy relevant insights without the person’s knowledge. Crawford and Schultz (2014) provide an example on “predictive privacy harms” based on a New York Times article[1]. The article describes how customer data from a retail chain was analysed using data mining techniques, allowing to predict which female customers were pregnant, even if they had not announced this publicly.

---

[1] Charles Dugigg, How Companies Learn Your Secrets (2012) https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html

This shifts the focus towards the collection, use and application of data. According to Malhotra et al. (2004), "[t]he very act of data collection […] is the starting point of various information privacy concerns" (p. 338). An ethical research approach requires a balance between social and individual interests - the boundary of privacy is not rigid, but depends on the topic, place, time and user characteristics. Cartographers and geoscientists are responsible here 1.) to develop methods that protect the privacy of users 2.) but also allow flexibility of methods in order to shift application- and context-dependent boundaries. Importantly, even data that is used in ethically sound applications, making use of the data in the users´ personal interests, may be re-purposed, for example, to compromise user privacy. Our primary intent is to minimise these risks and maximise collective benefits, where the application of data is the shared interest of user and society. This is achieved at data collection time, by a new data store, limiting what data is collected and how it can be analysed.

Method

Protecting privacy of users with HyperLogLog

HyperLogLog is an algorithm we use to break up geo-social media posts into quantitative, statistical information units, so called HLL sets. Spatial, temporal, thematic and social information units can be generated, for which the number of different users (user count) or contributions (post count) is statistically encoded (see Fig. 1). When using the HLL sets, it can be estimated with an accuracy of 2% how many users have shared information at a location or in a specific period of time, or how many posts have been published on a keyword or topic.

This data transformation is irreversible, that is it is not possible to reconstruct individual elements from HLL set. This prevents correlating information for tracking single people across contexts (Katsomallos et al. 2019), such as linking posts from a single user at different times, locations and topics. Instead, only estimates about the quantity (e.g. how many users) are available. Therefore, the use of HLL together with dissolving original content into its individual components already provides strong benefits to user privacy, when compared to the use of raw data.

In addition to the quantitative analysis, HLL sets can be related using set operations (intersection and union) in order to implement more complex spatial, temporal and content-related analyses. HLL sets of the same or different types can be combined or compared, for example, by intersecting a spatial and thematic HLL set, it is possible to study how many users have shared information on a specific topic at a specific place. The ability to compare and intersect different sets has a natural tendency to produce greater errors for more fine-grained queries, up to producing random results for single-user queries. In this respect, HyperLogLog is suitable for the processing of personal data.

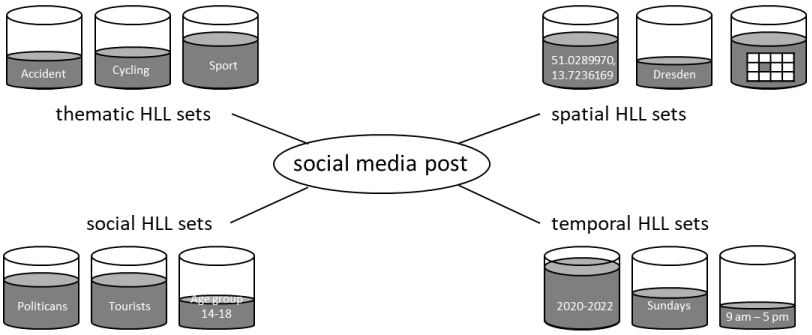

Figure 1: Social media posts can be broken down into different spatial, temporal, thematic or social HLL sets, for which the number of different users or number of posts is statistically encoded.

From a data processing point of view, the generated HLL sets are heavily compressed, so that much less storage space is required compared to the original data, e.g. 1.5 KByte for a set of 1 billion elements. Furthermore, very high-performance queries over extensive data sets are possible. The granularity of the information units and the partitions can be flexibly defined, both in terms of space and time as well as in terms of thematic breadth and depth. There are a number of parameter settings that can be adjusted during the creation of HLL sets, to allow fine-tuning what can be done and to which degree (see Dunkel et al. 2021).

Flexibility of the HLL methods in order to consider context

Notwithstanding these baseline benefits to privacy, there are edge cases that require special handling. For instance, even the existence of a single specific term, a specific time, or location (etc.), may provide hints that can be repurposed or combined with other (e.g. external) information to compromise user privacy in certain situations. Following the principle that different data must be treated differently (Almås et al. 2018), we seek to contribute to a systematic approach to fine-tuning privacy preservation and analytical flexibility.

There are two main approaches to adjust privacy–utility tradeoffs with HLL. First - stop and allow lists can be used during the generation of the HLL set to enable context-dependent data protection through filtering. Second - threshold values can be defined flexible to influence granularity of HLL sets and based on that, the degree of anonymity. Table 1 lists example contexts for each context in the framework, where accuracy (utility) may be traded in favour of a higher degree of privacy, similar to the broader data sensitivity spectrum proposed by Rumbold & Pierscionek 2018.

Type of context	**Example of sensitive context factor**	**Reference**
Spatial context	home location	Georgiadou et. al. (2019) Kim et al. (2021)
	hospitals	Ağır et al. (2016), Kim & Kwan (2021)

	related to specific events (concert grounds, party locations)	Such et al. (2017)
Temporal context	night times	Nikas et al. (2018)
	past and archived content, time collapse	Brandtzaeg & Lüders (2018)
	during specific events (e.g. new year, Thanksgiving, 4th July)	Such et al. (2017)
Thematic context	activists, protesters, dissidents	Uldam (2018)
	health issues (e. g. related to diabetes or corona)	Markovic et al. (2021)
Social context	children	Steinberg (2017); Marwick & boyd (2014)
	LGBTQ+[2]	Birnholtz (2020)
	personal, social relationships	Houghton & Joinson (2010)
	minorities (race and religion)	Mashhadi et al. (2021)

Table 1: Example of sensitive context factors for which no data analysis might be carried out.

Whether stop lists or allow lists are preferable depends on the context of application. Allow lists are more restrictive and require less effort from the analysts, by automatically excluding all terms, times, or locations (etc.) that are not explicitly considered beforehand. For the spatial context, for instance, unless worldwide data is required, allow lists are frequently used, to limit data collection to a specific area, region or place (etc.). Conversely, stop lists can be added selectively on top, to exclude places that are known to be related to vulnerable groups or sensitive contexts (e.g. hospitals, party locations). In a similar way, filter lists for specific terms, hashtags or emoji can be defined for the thematic context.

Figure 2: A thematically sensitive emoji on drug use at selected locations

[2] lesbian, gay, bisexual, transgender, queer, and others)

For thematic contexts, the openness of possible references complicates defining holistic stop lists ahead of time. As an example, Figure 2 shows a map generated from terms, hashtags and emoji used on Geosocial Media (Twitter, Flickr, Instagram) at a public vantage point and park. The syringe emoji (💉) could indicate drug use, which may lead to further on-site investigation by (e.g.) authorities, with potential unexpected consequences from the users perspective. Obviously, this is an edge case for social-individual privacy because both positive (society) and negative (user) consequences are imaginable. One solution would be to assign the specific emoji to a thematic broader emoji class, e.g. the umbrella group of "medical emoji"[3]. As another solution, the 💉 emoji could be classified ahead of time, for increased sensitivity, leading to (e.g.) a greater spatial granularity reduction on data ingestion, or exclusion, preventing having to deal with this ambiguous ethical edge case in advance.

Lastly, as the second approach to enable systematic user privacy with HLL, threshold values may be defined, similar to what is known from other disciplines, such as the HIPAA Privacy Rules for health data publications (Malin et al. 2011) or census statistics (Szibalski, 2007, p. 142). Allshouse et al. (2010), for instance, use geomasking in combination with k-anonymity, to define a lower threshold of k=5 (people), which is a rule of thumb size in geoprivacy (Kamp et al. 2013). Comparable best-practice threshold values could be defined for HLL sets of different sizes (e.g. suggestions by Desfontaines et al. 2019), with smaller sets indicating lessor privacy protection due to a scarce context collapse. In the spatial context this could be implemented by using quadtrees, for example, to split and aggregated social data into sub-sections (quads), based on pre-defined thresholds, where the resolution is automatically decreased for areas of lessor data density.

Discussion and Conclusion

HLL features several characteristics that make it particularly suited as an intermediate, privacy-aware component for location aware applications, such as VGI and geosocial data, at data collection time. By allowing to split data into comparable subsets of spatial, temporal, thematic and social units, HLL supports the privacy principle of treating different data differently. However, since the HLL algorithm only allows cardinality estimation, its application to the spatial domain requires consideration of addition components, methods and risk mitigation strategies. A range of needs to gradually tune privacy–utility trade-offs at various stages of data processing have been discussed in this paper.

Stop and allow lists are proposed to accommodate different types of context flexible. In particular sensitive topics can be excluded, specific user groups can be protected and sensitive locations and periods of time can be treated with appropriate caution. Furthermore, threshold values can be defined related to the minimal number of elements, which should be contained within a HLL set. If threshold values are not reached, HLL units will need to be aggregated or generalised further. For spatial units the quadtree level, for instance, could be lowered, and for thematic units a broader thematic group could be selected. Such thematic groups can be derived from thematic classifications or semantic

[3] Unicode Consortium, unicode.org/emoji/charts-13.0/full-emoji-list.html#medical

hierarchies. As it cannot always be defined rigid ahead of time, the HLL algorithm also allows posterior increase of privacy protection for already collected data in databases (a union operation; in the context of databases known as roll-ups).

Balancing social-individual and privacy-utility trade-offs requires a discourse on what is deemed acceptable. Possible discrimination as a result of privacy protection must also be taken into account. If information about groups worthy of protection is excluded, this must be compensated, for instance by consideration of different methods of data collection, in order to avoid disadvantages or non-consideration in analysis and decisions as far as possible.

Acknowledgements:

This work was supported by the German Research Foundation as part of the priority programme "Volunteered Geographic Information: Interpretation, Visualisation and Social Computing" (VGIscience, priority programme 1894).

References

Ağır, B., Huguenin, K., Hengartner, U. & Hubaux, J.-P. (2016). On the Privacy Implications of Location Semantics. Proceedings on Privacy Enhancing Technologies, vol.2016, no.4, 2016, pp.165-183. https://doi.org/10.1515/popets-2016-0034

Allshouse, W. B., Fitch, M. K., Hampton, K. H., Gesink, D. C., Doherty, I. A., Leone, P. A., Serre, M. L. & Miller, W. C. (2010). Geomasking sensitive health data and privacy protection: An evaluation using an E911 database. Geocarto International, 25(6), 443–452. https://doi.org/10.1080/10106049.2010.49649

Almås, I., Attanasio, O., Jalan, J., Oteiza, F. & Vigneri, M. (2018). Using data differently and using different data. Journal of Development Effectiveness, 10(4), 462–481. https://doi.org/10.1080/19439342.2018.1530279

Birnholtz, J., Kraus, A., Zheng, W., Moskowitz, D.A., Macapagal, K. & Gergle, D. (2020). Sensitive Sharing on Social Media: Exploring Willingness to Disclose PrEP Usage Among Adolescent Males Who Have Sex With Males. Social Media + Society. https://doi.org/10.1177/2056305120955176

Brandtzaeg, P. B. & Lüders M. (2018). Time Collapse in Social Media: Extending the Context Collapse. Social Media + Society. https://journals.sagepub.com/doi/10.1177/2056305118763349

Crawford, K. & Schultz, J. M. (2014). Big data and due process: Toward a framework to redress predictive privacy harms. Boston College Law Review, 55(1), 93–128. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2325784

Crawford, K. & Finn, M. (2015). The limits of crisis data: analytical and ethical challenges of using social and mobile data to understand disasters. GeoJournal 80, 491–502. https://link.springer.com/article/10.1007/s10708-014-9597-z

Desfontaines, D.; Lochbihler, A. & Basin, D. (2019). Cardinality Estimators do not Preserve Privacy. Proc. Priv. Enhancing Technol. 2, 26–46. https://doi.org/10.48550/arXiv.1808.05879

Dunkel, A., Löchner, M. & Burghardt, D. (2020). Privacy-aware visualization of volunteered geographic information (VGI) to analyze spatial activity: A benchmark implementation. ISPRS International Journal of Geo-Information, 9(10). https://doi.org/10.3390/ijgi9100607

Georgiadou, Y., De By, R. A. & Kounadi, O. (2019). Location privacy in the wake of the GDPR. ISPRS International Journal of Geo-Information, 8(3). https://doi.org/10.3390/ijgi8030157

Houghton, D. J. & Joinson, A. N. (2010) Privacy, Social Network Sites, and Social Relations, Journal of Technology in Human Services, 28:1-2, 74-94, https://doi.org/10.1080/15228831003770775

Kamp, M., Kopp, C., Mock, M., Boley, M. & May, M. (2013). Privacy-Preserving Mobility Monitoring Using Sketches of Stationary Sensor Readings. In: Blockeel, H., Kersting, K., Nijssen, S., Železný, F. (eds) Machine Learning and Knowledge Discovery in Databases. ECML PKDD 2013. Lecture Notes in Computer Science, vol 8190. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-40994-3_24

Katsomallos, M., Tzompanaki, K. & Kotzinos, D. (2019). Privacy, space, and time: A survey on privacy-preserving continuous data publishing. Journal of Spatial Information Science, 19(19), 57–103. https://doi.org/10.5311/JOSIS.2019.19.493

Kim, J. & Kwan, M.-P. (2021). An examination of people's privacy concerns, perceptions of social benefits, and acceptance of COVID-19 mitigation measures that harness location information: a comparative study of the US and South Korea. ISPRS International Journal of Geo-Information, 10(1) https://doi.org/10.3390/ijgi10010025

Kim, J., Kwan, M.-P., Levenstein, M. C. & Richardson, D. B. (2021). How do people perceive the disclosure risk of maps? Examining the perceived disclosure risk of maps and its implications for geoprivacy protection. Cartography and Geographic Information Science, 48(1), 2021, 2–20. https://doi.org/10.1080/15230406.2020.1794976

Malhotra, N. K.; Kim, S. S. & Agarwal, J. (2004). Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. Inf. Syst. Res., 15, 336–355. https://www.jstor.org/stable/23015787

Malin, B., Benitez, K. & Masys, D. (2011). Never too old for anonymity: A statistical standard for demographic data sharing via the HIPAA Privacy Rule. Journal of the American Medical Informatics Association, 18(1), 3–10. https://doi.org/10.1136/jamia.2010.004622

Markovic, R., Vejmelka, L. & Kljucevic., Z. (2021) Impact of COVID 19 on the Use of Social Networks Security Settings of Elementary and High School Students in the Split-Dalmatia County. 2021 44th International Convention on Information, Communication and Electronic Technology (MIPRO), pages 1476-1482. https://doi.org/10.23919/MIPRO52101.2021.9597179

Marwick, A. E. & boyd, d. (2014). Networked privacy: How teenagers negotiate context in social media. New Media & Society, 16, 1051–1067. https://doi.org/10.1177/1461444814543995

Mashhadi, A., Winder, S. G., Lia, E. H. & Wood, S. A. (2021). No Walk in the Park: The Viability and Fairness of Social Media Analysis for Parks and Recreational Policy Making. ICWSM 2021. Retrieved from https://ojs.aaai.org/index.php/ICWSM/article/view/18071

Nikas, A., Alepis, E. & Patsakis, C. (2018). I know what you streamed last night: On the security and privacy of streaming. Digital Investigation, 25, 78–89. https://doi.org/10.1016/j.diin.2018.03.004

Nissenbaum, H. (2010). Privacy in context: Technology, policy and the integrity of social life. Stanford, CA: Stanford University Press.

Olteanu, A., Castillo, C., Diaz, F. & Kıcıman, E. (2019). Social Data: Biases, Methodological Pitfalls, and Ethical Boundaries. Frontiers in Big Data, https://doi.org/10.3389/fdata.2019.00013

Rumbold, J. M. M. & Pierscionek, B. K. (2018). What Are Data? A Categorization of the Data Sensitivity Spectrum. Big Data Research, 12, 49–59. https://doi.org/10.1016/j.bdr.2017.11.001

Solove, D. J. (2012). Privacy self-management and the consent dilemma. The George Washington University Law School. Public Law and Legal Theory Paper & Legal Studies Research Paper No. 2012-141. https://ssrn.com/abstract=2171018

Steinberg, S. B. (2017). Sharenting: Children's Privacy in the Age of Social Media, 66 Emory L. J. 839 (2017). https://scholarlycommons.law.emory.edu/elj/vol66/iss4/2

Such, J. M., Porter, J., Preibusch, S. & Joinson, A. (2017). Photo privacy conflicts in social media: A large-scale empirical study. Conference on Human Factors in Computing Systems - Proceedings, 2017-May, 3821–3832. https://doi.org/10.1145/3025453.3025668

Szibalski, M. (2007). Kleinräumige Bevölkerungs- und Wirtschaftsdaten in der amtlichen Statistik Europas. Statistisches Bundesamt, Wiesbaden (Hrsg.): Auszug aus Wirtschaft und Statistik.

Uldam, J. (2018). Social media visibility: challenges to activism. Media, Culture & Society, 40(1):41-58. https://doi.org/10.1177/0163443717704997

Williams, M. L., Burnap, P. & Sloan, L. (2017). Towards an Ethical Framework for Publishing Twitter Data in Social Research: Taking into Account Users' Views, Online Context and Algorithmic Estimation. Sociology, 51(6), 1149–1168. https://doi.org/10.1177/0038038517708140

Zhang, H., McKenzie, G., Tomko, M., Egorova, E. & Kim, J. (2022). Report from the First Workshop on Cyber Ethics in Platial Research. In: FB Mocnik and R Westerholt (eds.), Proceedings of the 3rd International Symposium on Platial Information Science (PLATIAL'21), pp. 87–92. https://doi.org/10.5281/zenodo.6413002