

On the Privacy Afforded by Opaque Identifiers in Traffic Monitoring

Marcus Gelderie^a

Aalen University of Applied Sciences, Beethovenstr. 1, 73430 Aalen, Germany
{firstname.lastname}@hs-aalen.de

Keywords: Privacy, License Plate Recognition, Differential Privacy, Smart City, Traffic Monitoring

Abstract: We consider traffic monitoring via license plate recognition. Anonymizing license plates by substituting randomized identifiers is a common privacy enhancing strategy in this situation. However, the systematic effect of this anonymization strategy has not been fully explored. We study the information gain of an adversary upon observing such anonymized output. We find the effectiveness of randomized IDs to deteriorate with decreasing popularity of a given route. Moreover, we study the effect differential privacy has on the situation, given that an adversary must be assumed to have prior knowledge about the likelihood of various traffic patterns. We find that travel participants with a very strong preference for a given route are put most at risk.

1 INTRODUCTION


In the recent past, smart *traffic monitoring systems* (TMS) have received considerable attention from both the academic community as well as from commercial vendors (e.g. (Jain et al., 2019; Biswas et al., 2016; Bisio et al., 2022; Rana et al., 2021; Djahel et al., 2015; Gade, 2019; Bhardwaj et al., 2022)). The aim of smart traffic monitoring systems is to automatically collect data that are then used to facilitate other use-cases, such as traffic management or city planning. There are many different ways in which a TMS can be built. The kind of data that it records depends, in part, on that architecture (Rana et al., 2021). Examples include systems based on *digital image processing* (DIP) (Krishnamoorthy and Manickam, 2018), vehicle-to-X networks (Du et al., 2015), or probe vehicles (Feng et al., 2014). DIP-based approaches do not require any cooperation from the vehicles and are thus compatible with existing vehicle technology. Combined with *license-plate recognition* (LPR) (Jain et al., 2019; Du et al., 2013), they offer the possibility to track vehicles through a city and build statistics from these data.

Privacy is a particular concern when LPR is used. First, the license plate clearly is a piece of personally identifiable information (PII) — unlike, for example, a birds-view snapshot of a busy intersection. Second, indiscriminate scanning of license plates is difficult to base on explicit consent (which is relatively straightforward to do when using vehicle-to-X networks, for

instance). As a result, solutions relying on LPR usually employ some form of pseudo-anonymization (see e.g. (Gao et al., 2019)). An obvious idea is to replace the license plate by an opaque, pseudo-random ID for each vehicle, and use it to correlate individual locations at a central (e.g. cloud) service. Using pseudo-random IDs in this way is a relatively straightforward technique that is simple to implement. Augmenting with *differential privacy* (Dwork et al., 2006) is possible (Gelderie et al., 2024). Yet its security properties are not fully understood (more details below).

We study the question: “What information can be inferred from the anonymized data collected in this way, if the attacker has prior knowledge on the traffic patterns?” More specifically, we study this question in two settings: In its most basic nature, the data is simply anonymized and transmitted to a server. We call this the *anonymization scenario*. This scenario corresponds to what we have observed in commercial products and literature (see below for details). By contrast, the *differential privacy scenario* considers the case where the obfuscated data that reaches the server satisfies the requirements of differential privacy. Our main contribution is a precise information theoretic treatment of these scenarios.

In the anonymization scenario, we ask: Can the server de-anonymize the data given prior knowledge about a victim’s route preferences? This is the well-known *de-anonymization attack*. Given an opaque ID and the route which it travels (both are part of the data-set such a server consumes as part of its operation), we show that an adversarial server learns infor-

^a  <https://orcid.org/0009-0003-0291-3911>

mation inversely proportional to the routes popularity (the expected number of cars on it).

In the differential privacy scenario, the server sees data that is anonymized *and* subjected to randomized noise. In accordance with the typical definition of differential privacy, the server receives a given data set with roughly equal probability regardless of whether some specific individual drove along some route or not. We ask: What can the server ascertain about the probability of a specific individual driving along a given route under these circumstances, given that the server has prior knowledge about the overall traffic patterns (i.e. the underlying probability distribution). We show that an attacker can infer information about a target vehicle, if said vehicle shows a strong preference for a specific route. Since this is a fairly common situation, the privacy of individuals is at increased risk, even if differential privacy is guaranteed.

The importance of prior knowledge in studying privacy is well known. Indeed, the definition of differential privacy is intended to quantify privacy by excluding prior knowledge from the equation (see again (Dwork et al., 2014)). While a sensible approach, in practice prior knowledge needs to be taken into account when the *de-facto* privacy of a system is to be quantified. We think that our results show that simply anonymizing data and adding differential privacy does not sufficiently protect users under all circumstances. Care must be taken to restrict monitoring to sufficiently popular routes. We hope this paper can serve as a first step in studying the actual privacy guarantees that can be offered to a travel participant by modern license plate monitoring systems in practice.

In related work, the area of *traffic monitoring* is a wide field and only a subset of the existing research is pertinent to this paper. Generally, traffic monitoring systems fall into three categories (Jain et al., 2019): i) in situ (e.g. sensors embedded into the road surface) ii) vehicular (e.g. probe vehicles or vehicular networks) iii) digital image processing (DIP). Of particular relevance to this paper is the third category. Again there are many works on DIP based systems, e.g. (Du et al., 2013; Baran et al., 2014; Krishnamoorthy and Manickam, 2018; Bisio et al., 2022). Altogether, these prior works highlight the relevance of analyzing and addressing gaps in privacy guarantees offered by proposed anonymization techniques.

Differential Privacy (DP) was introduced by Dwork, Nissim, McSherry, and Smith in (Dwork et al., 2006; Dwork, 2006). DP has since seen sustained and intensive research activity, resulting in a slew of research studying various application domains (more below). Of general interest are observations on the limits of DP, particularly if the adversary is as-

sumed to have *prior knowledge* (Dwork et al., 2014).

Some works on DP investigate continual release of statistics that are built from *streams* of events (Dwork et al., 2010; Shi et al., 2011; Kellaris et al., 2014; Chan et al., 2011; Jain et al., 2023). Those works lay algorithmic foundations and provide lower bounds on the noise that is required to achieve DP. In particular, some works study the aggregation of statistics from multiply locations or agents, e.g. (Cheu et al., 2019; Corrigan-Gibbs and Boneh, 2017; Chan et al., 2012). One important aspect of these works is that they classify DP algorithms as either working in a *central model* or a *local model*. In the central model, DP is implemented by a trusted party (the *curator*) that sees all data-points in the clear. This curator then outputs a perturbed version of this statistic that meets the definition of DP. In the local model, on the other hand, one ensures that even subsets of agents involved in the collection of data (so-called *coalitions*) see only DP data. In particular, they usually ensure that the curator sees only a perturbed version of the raw data.

While the decentralized nature and focus on time-series data considered in these works is relevant to traffic monitoring, their focus is on algorithmic foundations and security guarantees in the context of DP. Prior knowledge of an adversary is not taken into account and, in some cases (e.g. smart metering) would be different from the use-case of traffic monitoring.

Many papers have tackled the problem of *privacy in traffic monitoring* from a solutions perspective (e.g. (Li et al., 2018; Qu et al., 2019; Sun et al., 2021; Gelderie. et al., 2024)). Those papers typically propose a specific privacy measure and analyze its security in the underlying model. For example, works with a focus on differential privacy (Gelderie. et al., 2024; Sun et al., 2021) prove that the definition of DP is satisfied. But by itself, DP says very little about the situation when an attacker has prior knowledge. Unfortunately, in traffic monitoring it is particularly easy for the attacker to acquire at least partial knowledge about the probabilities of certain traffic patterns (e.g. via products like Google Maps or because many cities publish such data). We must assume that the adversary has intimate knowledge of the overall traffic movement patterns and their probability distribution.

In this vein, there are efforts quantifying privacy loss (Gao et al., 2019). Gao et. al. study the privacy of LPR in a somewhat similar setting to ours. Their study is of an empirical nature and is based on a large LPR dataset. We seek to augment these results by performing a systematic mathematical analysis in the framework of information theory.

A related line of research is concerned with the privacy of vehicle data, such as trajectory information




Figure 1: Architecture Overview

(Ma et al., 2019; Zhou et al., 2018). In these works, the vehicles themselves actively participate in the data gathering. This means that those architectures have a wealth of privacy enhancing options at their disposal that cannot be easily ported to LPR settings.

In the larger context of *smart cities*, even more related works on privacy exist (e.g. (Husnoo et al., 2021; Hassan et al., 2019; Yao et al., 2023; Qu et al., 2019; Gracias et al., 2023)). The field is very diverse, but we are not aware of any works that investigate the effect of prior knowledge on the privacy afforded by license plate obfuscation or similar use-cases.

2 BACKGROUND

In the following, it is assumed that each vehicle is equipped with a unique and readable *license plate* l from some finite set \mathcal{L} that is recorded at certain points in the city. We call such locations *tracking points* (TPs). At each TP, a camera, or group of cameras, will record the license plate of every vehicle that passes by. This is depicted in fig. 1. TPs (depicted as circles) reside alongside roads, usually junctions (depicted as gray boxes). They report data to a *statistics server* (depicted in blue), which builds a central statistic about the number of vehicles per route.

The set V of all tracking points forms a directed graph $\mathcal{G} = (V, E)$, defined by the rule that $(v, v') \in E$ whenever v' can be reached from v without visiting a tracking point in between. We call \mathcal{G} the *city graph*. Let $\mathcal{G} = (V, E)$ be the city graph as described above. We define the set of *routes* as $\mathcal{R} = \{v_1 \dots v_l \mid (v_i, v_{i+1}) \in E, 1 \leq i < l, l \in \mathbb{N}\}$. As shown in fig. 1, two TPs can be neighbors in \mathcal{G} , even if they are not directly connected by some road. This suggests an accuracy trade-off, which we leave for future work.

The statistic server computes the target statistic $S: \mathcal{R} \rightarrow \mathbb{N}_0$, which simply counts cars on a given route. The statistic S is implicitly a function of time: $S(r)$ can vary over time for each $r \in \mathcal{R}$. In this paper we study snapshots of S at some unspecified, yet fixed

point in time t . Notably, all probability distributions that are studied depend on t . The fact that we consider snapshots is no limitation: The attacker’s advantage is then simply the maximum over her per-time-step advantages.

The information about the traffic statistic at previous times is relevant to an attacker: The victim might be known to normally drive route r_1 during rush hours, but divert to r_2 if a certain intersection is congested. In that situation, knowing whether or not said intersection was congested 20 minutes ago has an impact on the adversaries knowledge on the situation right now. However, the analysis conducted in this paper extends to situations where an attacker has knowledge about the last T time-steps for some fixed T . It merely affects the probability of vehicles colliding on IDs (see below).

The reports to the central server can expose license plate data to a central location. As a result, one usually anonymizes this data in some way. A typical strategy is to replace the license plate with some random identifier. Here, TPs transmit an opaque ID to the server instead of the license plate. As long as the same vehicle is always assigned the same opaque ID, the server can compute the same statistic. We do not consider *how* to implement such an assignment. Instead, we assume that there exists a binding of IDs, drawn from a set \mathcal{U} to license plates.

We assume this mapping of license plates to IDs is temporary: Each time an ID is assigned, this binding has a time to live (TTL) T . The TTL determines for how many hops a license plate is tracked. In particular, the routes $r \in \mathcal{R}$ that will be recorded have length at most T . Because of this, we will (abusing notation) assume that \mathcal{R} consists only of sequences of length at most T and is, in particular, finite.

The anonymization provided by the use of opaque IDs hinges on the assumption that the binding $B: \mathcal{L} \rightarrow \mathcal{U}$ is not known to an adversary. We say B is *secure*, if $B(l)$ is a uniform random variable for each $l \in \mathcal{L}$ and $B(l)$ is independent of $B(l')$ for all $l \neq l'$.

Note that our security notion does not take the implementation of such a binding into account. We focus purely on the stochastic properties of the binding itself and leave an investigation of the various technical options and cryptographic security notions of secure bindings to future work.

3 OBFUSCATION ONLY

We now state formal privacy guarantees that are afforded by TMS using randomized bindings as outlined in section 2. Our main statement in theorem 2 is

reminiscent of the well-known notions from cryptography. Of note is the role the expected value of cars per roads plays in the privacy that the system affords.

In this section, we assume opaque identifiers are chosen from the set $\mathcal{U} = \mathbb{B}^\lambda$, where $\lambda \in \mathbb{N}$ is a security parameter. We assume *secure* bindings, as defined above. However, we first study an idealized setting, where we assume that collisions on IDs do not occur: $U_l \neq U_{l'}$ for all $l \neq l'$. Such a binding is neither practical, not secure in our sense (the variables U_l are clearly not independent). We subsequently lift the result to the general case in theorem 2.

For notational convenience, we write $[\cdot]$ for events defined by some function on the output of one or more random variables. For example, if X_1, \dots, X_n are random variables with identical range, we may write $[\exists i \neq j: X_i \neq X_j]$ for the event that not all X_i are equal.

We first study an idealized setting and assume that IDs are assigned via an injective function $f: \mathcal{L} \rightarrow \mathcal{U}$ that is chosen uniformly at random. Write $\mathcal{F} = \{f: \mathcal{U} \rightarrow \mathcal{L} \mid f \text{ injective}\}$. We choose $f^* \leftarrow \mathcal{F}$ uniformly and assign to each $l \in \mathcal{L}$ the ID $f^*(l)$. We use a random variable $U_l = f^*(l)$ for $l \in \mathcal{L}$.

We use some additional random variables for notational convenience. Let $l \in \mathcal{L}$ and write R_l for the random variable that assigns a route to l . If l does not currently drive at all, then R_l and U_l take the special value $\perp \notin \mathcal{U} \cup \mathcal{L}$. Finally, we write $D = \{l \in \mathcal{L} \mid U_l \neq \perp\}$ for the set of cars that drive.

Note that $U_l = f^*(l)$ is defined only if l drives, although $f^*(l)$ is always defined. This is because in our model (see section 2), IDs are only assigned to cars that drive. In that situation, U_l and R_l *cannot be independent*. They are only conditionally independent on the events “ l drives” = $[l \in D]$ or “ l does not drive” = $[l \notin D]$.

For $u \in \mathcal{U}$, write $R_u = \{w \in V^* \mid w \text{ consistent with } u\}$ for the random variable that denotes the set of all sequences that can be associated with u . If u is not currently assigned to any license plate, then $R_u = \emptyset$. If there is a collision on u , then R_u may contain more than one element (including sequences of vertices that are not valid walks in \mathcal{G}). Since we assume, for the moment, that IDs are assigned using an injective function f^* , R_u is either empty, or a singleton containing one valid route from $\mathcal{R} \subsetneq V^*$. We therefore write $R_u = r$ whenever it is ensured that R_u is a singleton.

For $r \in \mathcal{R}$, write $N_r = |\{l \in \mathcal{L} \mid R_l = r\}|$ for the random variable denoting the number of vehicles on route r . Note that we make *no* assumption about the distribution of R_l and U_l . However, $\Pr[U_l = u \mid l \in D] = \Pr[f^*(l) = u] = |\mathcal{U}|^{-1} = 2^{-\lambda}$.

Lemma 1. *For every $r \in \mathcal{R}$, $u \in \mathcal{U}$ and ev-*

ery $l \in \mathcal{L}$ with $\Pr[R_u = r] \neq 0$, it holds that $\Pr[U_l = u \mid R_u = r] = \frac{\Pr[R_l = r]}{\mathbb{E}[N_r]}$, where the probabilities are taken over the uniform choice of $f^ \leftarrow \mathcal{F}$ and over the randomness of R_l .*

This and all other proofs are omitted due to space constraints, but can be found in the full version of this paper (Gelderie, 2024).

The assumption that f^* is chosen uniformly from the set \mathcal{F} of all injective functions from \mathcal{L} to \mathcal{U} is, of course, impractical. It is more natural to chose on $u \in \mathcal{U}$ per $l \in \mathcal{L}$ uniformly at random (as is done whenever we work with Variant 2, Version 4 UUIDs, for example).

If we draw IDs from the set $\mathcal{U} = \mathbb{B}^\lambda$ uniformly at random, we may deal with collisions. This case can be dealt with in the usual way using well-known Birthday Paradox probability bounds. This is the content of the theorem below.

In the following, we recall that $R_u \subseteq V^*$ is defined to be the set of all sequences of vertices that are consistent with $u \in \mathcal{U}$. This set can now contain more than one element.

Theorem 2. *Let $\lambda \in \mathbb{N}$. For every $r \in \mathcal{R}$, $u \in \mathcal{U}$ and every $l \in \mathcal{L}$ with $\Pr[r \in R_u] \neq 0$ it holds that*

$$\left| \Pr[U_l = u \mid r \in R_u] - \frac{\Pr[R_l = r]}{\mathbb{E}[N_r]} \right| \leq \frac{|\mathcal{L}|^2}{2^\lambda}$$

where the probabilities are taken over U_l and R_l .

This result formalizes the intuitive notion that driving anonymously along a very popular route does not leak much information about the vehicle l . Conversely, if the route is unpopular, then the system leaks information consistent with the probability of vehicle l driving that route. In particular, we capture intuitively obvious observations, such as: If vehicle l is parked at a remote location every night, then the expected number of cars on all routes leading to that location is close to 1. In this event the driver can be de-anonymized: $\Pr[R_{l'} = r] \approx 0$ for all $l' \neq l$.

4 OBFUSCATION & DP

It is well-known (see e.g. (Dwork et al., 2014)) that the security guarantees afforded by DP make no statement about the knowledge an adversary might draw from *prior knowledge*. If a traffic monitoring system provides (ϵ, δ) -DP, this just says that any two possible adjacent inputs produce the same output with almost equal probability. But of course, those two adjacent inputs could reasonably have very different prior probabilities, meaning that an attacker can infer significantly more information about the nature of the

input from observing the output than one might reasonably expect from the notion of DP.

We recall the definition of differential privacy (Dwork et al., 2006; Dwork, 2006). Note that *adjacency* has not yet been defined; we give a use-case specific definition below.

Definition 1. Let $\epsilon, \delta > 0$. Let \mathcal{M} be a randomized algorithm. Then \mathcal{M} is said to have (ϵ, δ) -differential privacy, if for all adjacent inputs $x, x' \in \text{dom}(\mathcal{M})$ and all subsets $S \subseteq \text{range}(\mathcal{M})$ it holds that $\Pr[\mathcal{M}(x) \in S] \leq \exp(\epsilon) \cdot \Pr[\mathcal{M}(x') \in S] + \delta$. The algorithms \mathcal{A} is referred to as the *curator*.

In this section, we investigate the effect that an (ϵ, δ) -DP curator \mathcal{M} has on the conclusions an adversary can rationally draw from the observed outputs. Recall that we assign routes to vehicles via the random variable $R_l \in \mathcal{R} \uplus \{\perp\}$ for each $l \in \mathcal{L}$, where $R_l = \perp$ means that the vehicle with license plate l is not driving at all.

First, we introduce some notation. In what follows, it is convenient to interpret \perp as just another route. We define $\hat{\mathcal{R}} = \mathcal{R} \uplus \{\perp\}$. We denote by $X = (R_l)_{l \in \mathcal{L}}$ the random vector that assigns a route (or \perp) to every vehicle. Clearly X is a complete representation of the input to a curator in the central DP model. We recognize values of x of X as functions $x: \mathcal{L} \rightarrow \hat{\mathcal{R}}$. Write $\mathfrak{X} = \hat{\mathcal{R}}^{\mathcal{L}}$ for the set of all possible values X might take on. Note that since \mathcal{L} and \mathcal{R} are finite (the latter because of the length-bound enforced by the TTL), the set \mathfrak{X} is also finite. Let $p: \mathfrak{X} \rightarrow [0, 1]$ the distribution of X .

To reason about DP, we need to clarify adjacency in our context: Let $l \in \mathcal{L}$ and let $x, x' \in \mathfrak{X}$. We call x and x' *a-adjacent* (or simply *adjacent*), if $x(s) = x'(s)$ for all $s \neq l$, and $x'(l) \neq x(l)$. We write $x \bowtie_l x'$.

Recall that the *information content* of x is defined as $I(x) \stackrel{\text{def}}{=} -\log(p(x))$. We can now define:

Definition 2 (Preference Gap). Let $l \in \mathcal{L}$ and $r \in \hat{\mathcal{R}}$. The quantity $\sigma(l, r) = \sup_{x \in A_r} \sup_{x' \bowtie_l x} I(x') - I(x)$ is called the *preference gap* of l and r . The quantity $\sigma(l) = \sup_{r \in \hat{\mathcal{R}}} \sigma(l, r)$ is called the *preference gap* of l .

With the convention that $\exp(\infty) = \infty$, we have:

Theorem 3. Let \mathcal{M} be an (ϵ, δ) -DP curator on \mathfrak{X} .

For every $l \in \mathcal{L}$, $r \in \hat{\mathcal{R}}$ and $y \in \text{range}(\mathcal{M})$ with $\Pr[Y = y] > 0$ it holds that: $\Pr[R_l = r \mid Y = y] \leq (\exp(\epsilon) \Pr[R_l \neq r \mid Y = y] + \frac{\delta}{\Pr[Y=y]}) \cdot \exp(\sigma(l))$

If $\Pr[R_l = r] \neq 0$, then: $\Pr[R_l \neq r \mid Y = y] \leq |\mathcal{R}| \cdot (\exp(\epsilon) \Pr[R_l = r \mid Y = y] + \frac{\delta}{\Pr[Y=y]}) \cdot \exp(\sigma(l))$.

5 CONCLUSION

We have considered traffic monitoring using anonymized license plates. In this context, we have studied the question, how prior knowledge about the overall probability distributions of the general driving behavior or individual participants affects the privacy guarantees of such traffic monitoring systems. We extended this study to systems that provide differential privacy.

When no DP is involved, the knowledge an adversary has about the probabilities of individual driving behavior can greatly increase the confidence in unmasking attacks, where an adversary tries to identify the individual behind a certain opaque ID. Specifically, if the route in question is very unpopular, the risk of unmasking is high.

We then studied how the guarantees of DP, stated in terms of neighboring data-sets, generalize to guarantees about the likelihood of a particular individual driving on a certain route. We found that the unevenness of the underlying probability distribution of traffic patterns can degrade the assurances made by the DP mechanism significantly.

An interesting open question for future work is to what extent the picture changes when the adversary has only *partial* knowledge of the probability distributions involved. For instance, an adversary may have information about the probability of a given number of vehicles per route at a given time, but not the probabilities of individual vehicles being on that route.

REFERENCES

- Baran, R., Ruść, T., and Rychlik, M. (2014). A smart camera for traffic surveillance. In *Multimedia Communications, Services and Security: 7th International Conference*, pages 1–15. Springer.
- Bhardwaj, V., Rasamsetti, Y., and Valsan, V. (2022). Traffic control system for smart city using image processing. *AI and IoT for Smart City applications*, pages 83–99.
- Bisio, I., Garibotto, C., Haleem, H., Lavagetto, F., and Sciarrone, A. (2022). A systematic review of drone based road traffic monitoring system. *IEEE Access*, 10.
- Biswas, S. P., Roy, P., Patra, N., Mukherjee, A., and Dey, N. (2016). Intelligent traffic monitoring system. In *Proc. o. 2nd Intl. Conf. on Computer and Communication Technologies*, pages 535–545. Springer.
- Chan, T.-H. H., Shi, E., and Song, D. (2011). Private and continual release of statistics. *ACM Transactions on Information and System Security*, 14(3):1–24.
- Chan, T.-H. H., Shi, E., Song, D., and Song, D. (2012). Optimal lower bound for differentially private multi-party aggregation. *Embedded Systems and Applications*.

- Cheu, A., Smith, A., Ullman, J., Zeber, D., and Zhilyaev, M. (2019). Distributed differential privacy via shuffling. In *38th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 375–403. Springer.
- Corrigan-Gibbs, H. and Boneh, D. (2017). Prio: Private, robust, and scalable computation of aggregate statistics. In *14th USENIX symposium on networked systems design and implementation*, pages 259–282.
- Djahel, S., Doolan, R., Muntean, G.-M., and Murphy, J. (2015). A communications-oriented perspective on traffic management systems for smart cities: Challenges and innovative approaches. *IEEE Communications Surveys & Tutorials*, 17(1):125–151.
- Du, R., Chen, C., Yang, B., Lu, N., Guan, X., and Shen, X. (2015). Effective urban traffic monitoring by vehicular sensor networks. *IEEE Transactions on Vehicular Technology*, 64(1):273–286.
- Du, S., Ibrahim, M., Shehata, M., and Badawy, W. (2013). Automatic license plate recognition (alpr): A state-of-the-art review. *IEEE Transactions on Circuits and Systems for Video Technology*, 23(2):311–325.
- Dwork, C. (2006). Differential privacy. In *International colloquium on automata, languages, and programming*, pages 1–12. Springer.
- Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, pages 265–284. Springer.
- Dwork, C., Naor, M., Pitassi, T., and Rothblum, G. N. (2010). Differential privacy under continual observation. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, page 715–724.
- Dwork, C., Roth, A., et al. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407.
- Feng, Y., Hourdos, J., and Davis, G. A. (2014). Probe vehicle based real-time traffic monitoring on urban roadways. *Transportation Research Part C: Emerging Technologies*, 40:160–178.
- Gade, D. (2019). Ict based smart traffic management system “ismart” for smart cities. *International Journal of Recent Technology and Engineering*, 8(3):1000–1006.
- Gao, J., Sun, L., and Cai, M. (2019). Quantifying privacy vulnerability of individual mobility traces: A case study of license plate recognition data. *Transportation Research Part C: Emerging Technologies*, 104:78–94.
- Gelderie, M. (2024). On the privacy afforded by opaque identifiers in traffic monitoring (extended version). https://www.hs-aalen.de/uploads/publication/file/11363/opaque_ident_ext.pdf.
- Gelderie, M., Luff, M., and Brodschelm, L. (2024). Differential privacy for distributed traffic monitoring in smart cities. In *Proceedings of the 10th International Conference on Information Systems Security and Privacy*, pages 758–765.
- Gracias, J. S., Parnell, G. S., Specking, E., Pohl, E. A., and Buchanan, R. (2023). Smart cities—a structured literature review. *Smart Cities*, 6(4):1719–1743.
- Hassan, M. U., Rehmani, M. H., and Chen, J. (2019). Differential privacy techniques for cyber physical systems: a survey. *IEEE Communications Surveys & Tutorials*, 22(1):746–789.
- Husnoo, M. A., Anwar, A., Chakraborty, R. K., Doss, R., and Ryan, M. J. (2021). Differential privacy for iot-enabled critical infrastructure: A comprehensive survey. *IEEE Access*, 9:153276–153304.
- Jain, N. K., Saini, R., and Mittal, P. (2019). A review on traffic monitoring system techniques. *Soft computing: Theories and applications: Proceedings of SoCTA 2017*, pages 569–577.
- Jain, P., Raskhodnikova, S., Sivakumar, S., and Smith, A. (2023). The price of differential privacy under continual observation. In *International Conference on Machine Learning*, pages 14654–14678. PMLR.
- Kellaris, G., Papadopoulos, S., Xiao, X., and Papadias, D. (2014). Differentially private event sequences over infinite streams. *Proceedings of the VLDB Endowment*, 7(12):1155–1166.
- Krishnamoorthy, R. and Manickam, S. (2018). Automated traffic monitoring using image vision. In *2018 Second International Conference on Inventive Communication and Computational Technologies (ICICCT)*, pages 741–745.
- Li, Y., Zhang, P., and Wang, Y. (2018). The location privacy protection of electric vehicles with differential privacy in v2g networks. *Energies*, 11(10):2625.
- Ma, Z., Zhang, T., Liu, X., Li, X., and Ren, K. (2019). Real-time privacy-preserving data release over vehicle trajectory. *IEEE transactions on vehicular technology*, 68(8):8091–8102.
- Qu, Y., Nosouhi, M. R., Cui, L., and Yu, S. (2019). Privacy preservation in smart cities. In *Smart cities cybersecurity and privacy*, pages 75–88. Elsevier.
- Rana, S., Rana, S., Garg, U., Garg, U., Gupta, N., and Gupta, N. (2021). Intelligent traffic monitoring system based on internet of things. *2021 International Conference on Computational Performance Evaluation (ComPE)*.
- Shi, E., Chan, T.-H. H., Rieffel, E., Chow, R., and Song, D. (2011). Privacy-preserving aggregation of time-series data. In *Proceedings of the 18th Annual Network & Distributed System Security Symposium*, volume 2.4.
- Sun, Y.-E., Huang, H., Yang, W., Chen, S., and Du, Y. (2021). Toward differential privacy for traffic measurement in vehicular cyber-physical systems. *IEEE Transactions on Industrial Informatics*, 18(6).
- Yao, A., Li, G., Li, X., Jiang, F., Xu, J., and Liu, X. (2023). Differential privacy in edge computing-based smart city applications: Security issues, solutions and future directions. *Array*, page 100293.
- Zhou, Z., Qiao, Y., Zhu, L., Guan, J., Liu, Y., and Xu, C. (2018). Differential privacy-guaranteed trajectory community identification over vehicle ad-hoc networks. *Internet Technology Letters*, 1(3):e9.

Acknowledgements This work was done as part of a project funded by the German Federal Office for Information Security under project funding reference number 01MO23006.