A Provably Privacy Preserving Video Surveillance Architecture for an Assisted Living Community

Matthias Huber[2], Jörn Müller-Quade[1], Tobias Nilges[1], and Carolin Thal[1]

[1]Karlsruher Institut für Technologie,
{mueller-quade,nilges}@kit.edu,carolin.thal@student.kit.edu
[2]FZI Forschungszentrum Informatik, huber@fzi.de

Abstract: Video surveillance offers many advantages but also introduces privacy issues. We propose a video surveillance architecture that preserves privacy by providing different levels of anonymization. To this end, we model an assisted living scenario and show that our architecture can provide provable privacy under explicit assumptions while maintaining the utility of the architecture. This shows that video surveillance, tailored to specific scenarios, can be applied in a privacy preserving way.

1 Introduction

Video surveillance is increasingly common in every day life, for example in public places and buildings. While video surveillance is mostly used to improve the safety, these systems can also be used to compromise the privacy of individuals. To prevent misuse, laws have been established to provide transparency and set strong bounds on further use of the collected data. But laws cannot proactively prevent misuse, so privacy enhancing mechanisms can be applied to surveillance systems to enforce compliance with the law.

In literature, privacy enhancing mechanisms for video surveillance have been extensively studied. The proposed solutions, however, are mostly methods like obscuring faces [Bou05], relocating pixels [CC11], or different levels of anonymization [Sen+03; Sen+08]. These methods do not provide any privacy guarantees, i.e. it is not clear if and to which extend privacy is achieved. Therefore assumptions are necessary, e.g. that after application of method A the video data satisfies B, where B is a precise statement of the guarantee. It remains an open problem to provide reasonable privacy guarantees from methods.

In the case of database anonymization, where the model is far better defined, anonymization methods like k-anonymity [SS98] and l-diversity [KG06] stand in contrast to notions like differential privacy [Dwo06], that provide a clear guarantee. In [ABN08], the concept of k-anonymity is extended to obscure location information in moving objects databases. An individual only signalizes a location update if it traverses from one predefiend area to another. This is comparable to the map view that we describe in Section 4. For video data, it is unclear how to define a privacy metric which describes the information that is provided by a video, so apart from specific features that have to be hidden (e.g. height, color, gender), the quality of an anonymization seems to be an empirical measure.

Our contribution. We propose a system architecture for video surveillance of an assisted living community based on the NEST [MRV10] architecture. Here, an operator controls the surveillance system. The proposed system aids the operator to detect accidents and emergencies, while still providing reasonable privacy guarantees for the residents and employees. The system does not show any information about residents or employees if no one trespasses or needs help. Furthermore, we show that even when the system shows information about residents or employees, certain privacy guarantees are still fulfilled. These guarantees can improve the rate of consent for a surveillance system.

We make assumptions about the guarantees of the individual methods used in the system. In order to achieve a precise privacy guarantee for the overall system, we prove the privacy of the combination of these methods. Based on these guarantees, different views are provided for the operator, similar to [FS08]. In detail, the architecture consists of three levels. During normal operation, only guests are tracked on a *map*. In case of an (undetected) accident, a help gesture can be performed in front of a camera. The operator only sees an *indealized stick figure*. If an accident happens, the system will present the operator with a *silhouette view* of the video sequence of the accident.

The privacy has to hold against an operator who has detailed knowledge of the people that are surveilled. It is known that privacy methods do not compose in general, i.e. several anonymized datasets combined can still leak information about a single individual [NS08; SS98]. Thus, if we combine several methods in a complex architecture, the security crucially depends on the fact that the combination still preserves the privacy. We show that this is actually the case for our architecture.

2 Preliminaries

2.1 k-Anonymity

k-Anonymity is an anonymization method defined for tables. A table is a set of tuples $t = \{t_1, .., t_m\}$ with an associated set of attributes $A = \{A_1, .., A_n\}$, where every tuple represents an individual. k-anonymity classifies the attributes of a table into *identifier*, *quasi-identifier*, and *sensitive information*. In the following, w.l.o.g. we assume that a table contains only one sensitive information attribute.

Identifiers are attributes that directly allow for identifying individuals. A table that adheres k-anonymity may not contain such an attribute. A set of nonsensitive attributes $Q \subseteq A$ that linked with external data allows for the identification of at least one individual represented in the table is called quasi-identifier.

Definition 1 ([SS98]) *A table adheres k-anonymity if for each row, there are at least k-1 other rows with the same values of the quasi-identifier attributes. The set of all rows with the same attribute values of their quasi-identifiers are called an equivalence class.*

The idea of k-anonymity is that every sensitive statement in a k-anonymous table can relate

to at least k individuals, therefore, the privacy of a single individual is protected. Consider for example the table depicted in Figure 1.

		Non-Sensitive				Sensitive
		Zip Code	Age	Nationality		Condition
1		130**	< 30	*		Heart Disease
2		130**	< 30	*		Heart Disease
3		130**	< 30	*		Viral Infection
4		130**	< 30	*		Viral Infection
5		130**	3*	*		Cancer
6		130**	3*	*		Cancer
7		130**	3*	*		Cancer
8		130**	3*	*		Cancer

Figure 1: A 4-anonymous table (taken from [KG06]), with the quasi-identifier {*Zip Code, Age, Nationality*} and the sensitive information *Condition*. This table has two 4-buckets.

The notion k-anonymity has several drawbacks [Swe02; KG06; LLV07]. It does for example not guarantee diversity of the sensitive attribute value in an equivalence class as in the last equivalence class of Figure 1. Additionally, a composition of two k-anonymous tables may compromise the privacy of an individual [GKS08].

2.2 The Network Enabled Surveillance and Tracking System

This work is based on the Network Enabled Surveillance and Tracking (NEST) system [MRV10; VKB12]. The NEST camera system is a semi-automatic surveillance system. It can track individuals over multiple rooms and cameras, as well as support operators at evaluating situations. The NEST system can process incoming video data and extract features like an individuals height, hair color, sex, color information, and movement vector. These features are stored in a feature database and are fed to a model that is presented to the operator. After the feature extraction, the NEST system deletes the video data save for a very short buffer in order to protect the privacy of the surveyed individuals. It can alarm an operator about the occurrence of previously defined situations. Therefore, it reduces on the one hand the amount of data stored. On the other hand, it supports the operator by reducing the amount of data needed to be screened.

2.3 Legal Aspects of Video Surveillance in Germany

For the system proposed in this paper, there are requirements from German privacy regulations. They state that unless there is a legal basis or the individual involved consents, data related to it may not be collected, processed, stored or used otherwise. Video surveillance is treated as collecting personal data (§4 BDSG [Bun78]).

Video surveillance of public space is allowed, however, in order to enforce domestic au-

thority or other eligible interests with prior defined purposes [Zil07]. Involved individuals have to be informed about the surveillance and collected data has to be deleted after use.

For non public space there is no such regulation. In our scenario, we assume that the individuals have to consent to the video surveillance. In this case, the system proposed in this paper can improve the acceptance, since personal data only leaves the system under prior defined circumstances.

Another legal requirement is that the system may not make autonomous decisions (§6a BDSG [Bun78]). Such an autonomous decision could be alerting paramedics or even releasing private information. Therefore, the system needs an operator.

3 System Model

In this section we give an abstract representation of an assisted living community, from which we derive requirements for the database schema and for anonymization methods.

3.1 Scenario

We consider several real world settings which are easily extensible and can occur in an assisted living community.

- **Visits.** Visitors (both private and official) to the assisted living community have to be recognized.
- **Trespassing.** We assume areas in the assisted living community that are restricted to either staff or residents, e.g. recreation rooms for the staff.
- **Help gestures.** In case of an accident other individuals have the possibility to perform help gestures towards the camera in order to alarm the staff.
- **Accidents.** If no other individual is nearby when an accident (of a resident) happens, the staff shall be alarmed.
- **Emergency.** In case of an emergency, such as a fire, the surveillance system must assist the staff to evacuate every resident.

Due to the complexity of these settings we assume a significant amount of false positives for every setting. Therefore it is necessary to provide an operator with presumed emergencies to decide if action has to be taken. Additionally, as mentioned in Section 2.3, no surveillance system is allowed to act on its own and human interaction (in form of an operator) is necessary.

One way to provide the operator with the necessary information is to just replay the video of the situation which caused the system to evoke an alert. This, however, provides the operator with a lot of unnecessary additional information and even allows her to identify

the individual(s) on the video stream. It is also contrary to the principle of data minimization (§3a BDSG [Bun78]), because we can provide the operator with sufficient information without revealing the complete video stream.

3.2 Agents

We assume several types of individuals that take part in the assisted living scenario and thus are filmed by the video surveillance.

- **Residents.** Individuals in this group live in the assisted living community. Their privacy is paramount, since they are filmed (nearly) at all times.
- **Staff.** This group includes the nurses that take care of the residents as well as other individuals that are working permanently in the assisted living community.
- **Guests.** Here we consider all individuals that are either visiting residents or working only for a short period of time in the assisted living community, such as doctors.
- **Operator.** The operator uses the surveillance system and evaluates every alert of the system. We assume that the operator does not appear in the surveilled area.

Definition 2 *Let $\mathcal{U}$ be the set of all individuals that move inside the assisted living community. Further let* R *be the set of residents,* S *be the set of staff members and* G *the set of guests. We write $\mathcal{U}|_{\mathsf{X}}$ to denote a certain group $\mathsf{X} \in \{\mathsf{R}, \mathsf{S}, \mathsf{G}\}$.*

Considering that the surveillance system has to keep a record for each individual and needs to extract certain features to recognize the individual, we propose (w.r.t. the surveillance system described in Section 2.2) a database scheme $d_U = [\text{PID}, \text{Age}, \text{Size}, \text{Constitution}, \text{Colors}, \text{Role}]$ for this information.

3.3 Rooms

In our scenario, we distinguish between three different types of rooms. By the term room we do not necessarily mean a single room but an area that is covered by a camera and satisfies a certain criterion.

- **Rooms of residents.** This type summarizes all private rooms of the residents as well as recreation rooms and the kitchen.
- **Restricted areas.** Rooms such as recreation rooms for the staff and the pharmacy can be restricted. but also areas outside of the assisted living area can be declared, e.g. if residents with Alzheimer's disease may not leave the grounds.
- **Unobserved areas.** Due to the privacy of the residents, there are rooms that are not monitored, for example the sanitary facilities.

Definition 3 *Let $\mathcal{R}$ be the set of all rooms in the assisted living community.*

Since we have different types of rooms we also need to assign each room access rights. It is also important to know the alignment of the camera in each room (cf. Section 5.2.2), which leads to the relation $d_R = [\text{RID}, \text{Alignment}, \text{Access Rights}]$.

3.4 System View

The surveillance system internally needs to keep track of every individual in the assisted living community. Thus the view of the system $d_{AL} = [\text{PID}, \text{Age}, \text{Size}, \text{Constitution}, \text{Colors}, \text{Role}, \text{RID}]$ is a combination of d_R and d_U. Additionally, the buffered video sequences are kept to enable recognition of accidents or help gestures.

4 Methods for Anonymization of Video Data

It is not possible to completely characterize the extractable data of a video stream, so we will take a different approach. We state features that we want to be hidden from the operator in any case, and from the architecture we derive features that are necessarily visible to let the operator do her job.

In the following, we state the assumptions we need for the anonymized video and give examples for anonymization methods that are assumed to fulfill the assumptions. To make the assumptions explicit, we first need to introduce some notation. Let $X_{u,r}$ describe a video sequence that contains a user $u \in \mathcal{U}$ in a room $r \in \mathcal{R}$. We assume the anonymization mechanism to be deterministic functions f.

Definition 4 *Let f be an anonymization method for video sequences X. We call f a* k-anonymizer *iff for all $r, r' \in \mathcal{R}$ and for all $\mathcal{U}$ with $|\mathcal{U}| > n(f)$ it holds that for all $u \in \mathcal{U}$, there exists $v_i \in \mathcal{U}, 1 \leq i \leq k-1$, such that $f(X_{u,r})$ is indistinguishable from $f(X_{v_i,r'})$.*

By $n(f)$ we denote that depending on the anonymization method it might become necessary to have a larger set of individuals to compare to. The indistinguishability can only be based on empirical data, because no metric for the privacy of video data is known. Note also that the definition of an anonymizer implies that the rooms are pairwise indistinguishable as well.

Corollary 1 *Every k-anonymizer f implies k-anonymity on any feature set that can be extracted from video sequences anonymized with f.*

We propose three different anonymizers.

Map. The anonymizer f_M shows all *guests* of the assisted living community as dots on a map. In principle every individual is indistinguishable, i.e. $k = |\mathcal{U}_\mathsf{G}|$. All residents and the staff are not shown at all, and the guests cannot be identified by the dots (cf. Figure 3).

There are two attacks on this anonymization: the movement pattern of the depicted individuals, and associations of guests and residents. Movements patterns for guests are a small concern, and it might be possible to hide guests that regularly visit certain residents.

Avatar. The anonymizer f_A is based on an *idealized* stick figure on white background (cf. Figure 2(b)), which is extracted via gait recognition from the video stream. Apart from movements this method is assumed to remove most features of the originial individual. We assume that there are always atleast 2 individuals that are indistinguishable based on their gestures, i.e. $k = 2$.

One might assume that gait recognition can be used by an adversary to uniquely identify each individual, similar to [HMH14], who use waving gestures for authentication. [HMH14] need very good cameras and use a non-idealized image, so their approach will probably fail. Additionally, several samples are necessary to train the distinguisher, and an identity has to be associated manually.

Silhouette. We propose to use a combination of silhouette and pixelization/soft-focus as anonymizer f_S to anonymize accidents. The silhouette of the individual is determined, and the environment cut out. The silhouette is filled and the contours are pixelated/put in soft-focus. This removes many of the body features, while it is still possible to observe an accident (cf. Figure 2(a)). To obtain the required anonymity the degree of pixelization can be adjusted. Still, Definition 4 requires that at least two individuals are indistinguishable given the anonymized video sequence.

Pose-estimation or gait recognition will not work on the anonymized video sequence, because people have to be positioned properly in front of a high resolution camera. Certain specific properties of elderly persons like walking aids have to be either anonymized, or another person must also have such a walking aid. This is the assumption we have to make, which in general implies that the amount of residents has to be fairly large.

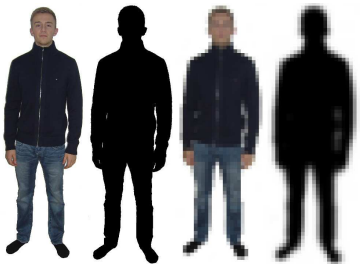

(a) Plain picture with silhouette, soft-focus effect, and a combination of both anonymization methods applied to it.

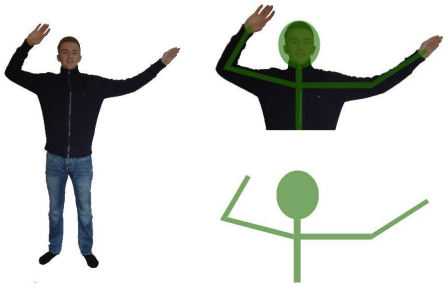

(b) A stick figure replacing an individual that forms a help gesture with its arms.

Figure 2: Two anonymization methods for video data.

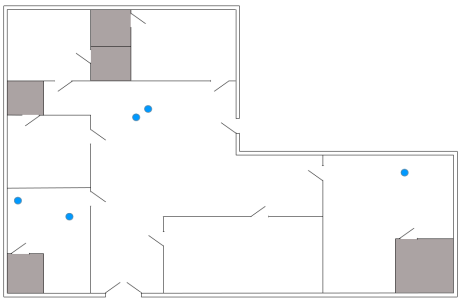

Figure 3: Example map with five individuals on it. Grey highlighted areas are unsupervised by camera. In our system, the maximal resolution for the position of agents are individual rooms.

5 Privacy Compliant Video Surveillance

Based on the model presented in Section 3 and the anonymization methods for video data presented in Section 4, we can now describe the behavior of the system. In Section 5.1, we describe the views the system presents to the operator under different circumstances. Based on this description, we provide a security analysis of the system in Section 5.2.

5.1 Behavior of the System

Depending on the situation, the operator only gets the information needed to assess it. The default view the system presents to the operator is the map view without individuals. In order to show accidents after they happened, the video stream is buffered in the video buffer (cf. Section 2.2) for a fixed time. Depending on the situation, the system presents additional information to the operator as follows:

- **Default.** The default view is the map view without any individuals.
- **Visitors.** If visitors are present, the system shows them as points on the map. Information about the position of visitors is necessary in order to prevent trespassing.
- **Trespassing of restricted areas.** If an individual trespasses to a restricted area, the operator needs to know the identity of the individual as well as her position. Therefore the trespassing individual is highlighted in the map view. Additionally the system presents the video stream of the trespassing individual to the operator.
- **Help gesture.** If the system recognizes a help gesture of an individual, it presents the operator the avatar view of the individual. If the operator verifies the help gesture, in order to assess the situation fully, the system presents the buffered video stream of the individual to the operator. If the operator recognizes a falsa alarm, the system switches back to default.

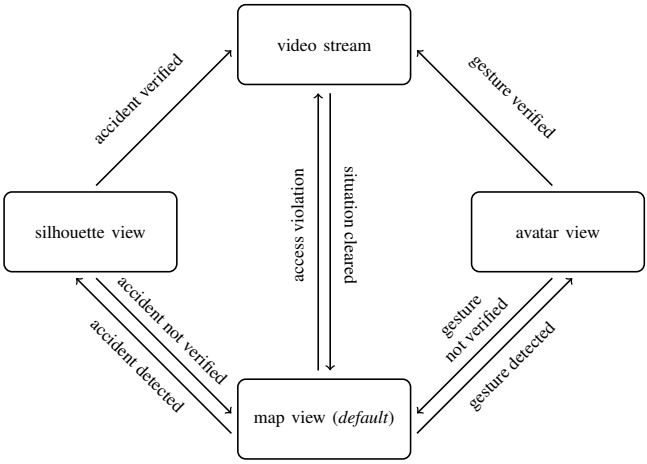

Figure 4: The behavior of the system. Depending on the situation, the system presents different views to the operator.

- **Accident.** If the system recognizes an accident of an individual, it presents the silhouette view of the accident to the operator. If the operator verifies the accident, the system presents the video stream, in order to support the operator assessing the situation fully and to react accordingly. If the operator does not verify the accident, the system switches back to default.
- **Emergencies.** Emergencies such as a fire are recognized by separate detectors and are issued manually. In case of an emergency, the system is informed about it. It then shows all camera streams to the operator.

In order to prevent misuse of false positives, every action of the operator that activates a plain video view is logged, adding a control instance for the operator.

5.2 Security Analysis

We will start our security analysis by first defining the adversary. Then we prove that the composition of views does not yield any more information than the individual views.

5.2.1 Adversary

Figure 5 depicts all possible attacks on the surveillance architecture. The attack on the connection between cameras and the server can be handled with standard techniques like encrypted connections and tamper-resilient cameras. We can also assume that the database server is placed in a secure location with proper access control.

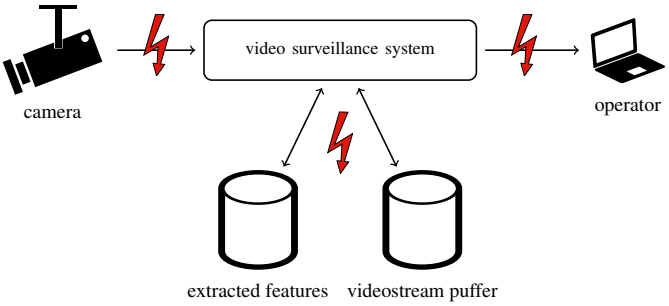

Figure 5: Possible weak points of the surveillance system that an adversary can attack.

The main point of attack that cannot be covered with standard techniques is the system view that the operator gets while supervising the system. She can try to correlate all the views to learn the identity of an individual in the assisted living community. We have to assume that the adversary has complete knowledge about all individuals, including clothes, behaviour and handicaps. Also, we assume that the operator stores everything that is depicted for later use.

5.2.2 Security Analysis

In Section 5.1 we defined four different views that an operator might see depending on the situation in the assisted living community. Obviously, the plain view of the video data directly identifies any individual. If plain video data is shown to the operator, however, either the individual wants to be (help gesture, accident, emergency) or needs to be (trespassing) identified. Thus we do not consider the unanonymized view in our security proof.

The remaining views of the operator are map view, avatar view and silhouette view. This leads to the following composition scenarios: (1) Map view $\leftrightarrow$ avatar view, (2) Avatar view $\leftrightarrow$ silhouette view and (3) Map view $\leftrightarrow$ silhouette view.

We will now prove that each of the above composition scenarios guarantees the privacy of any individual in the assisted living community. Please note, that due to the architecture of the video surveillance system, the operator $\mathcal{A}$ is either provided with the avatar view of a video sequence or a silhouette view but never with both views of the same video sequence.

Theorem 1 *Assuming that f_M, f_A and f_S are $|\mathcal{U}|$-, 2- and $|\mathcal{U}|_{\mathsf{G}}| + 2$-anonymizers respectively, as described in Section 4, an adversarial operator $\mathcal{A}$ cannot uniquely identify any individual in $\mathcal{U}$.*

The above theorem implies that any feature set extracted from the combined anonymized views will be 2-anonymous. This means that at least two individuals in the assisted living community are plausibly depicted in the anonymized view, independent of the information the operator might have gathered.

Proof 1 *Let $\mathcal{U}$ be the set of individuals present with $|\mathcal{U}| = n$, $\mathcal{U}|_{\mathsf{G}}$ be the set of guests present with $|\mathcal{U}|_{\mathsf{G}}| = g$, X_u be an video sequence of individual $u \in \mathcal{U}$, $d_M = f_M(d_{AL})$ be the map view (the RIDs of the guests), $d_A = f_A(X_u)$ be an avatar view of X_u, and $d_S = f_S(X'_{u'})$ be an silhouette view of $X'_{u'}$. Furthermore, let f_M and f_A provide n-anonymity and let f_S provide k-anonymity with $n > k \geq g + 2$. We show that the composition of (1) d_M and d_A, (2) d_A and d_S, and (3) d_M and d_S is 2-anonymous, respectively.*

The first two cases are straightforward: the anonymizer f_A provides 2-anonymity. Since the set of video sequences that are anonymized by f_A is disjunct from all video sequences that are anonymized by other anonymizers (due to the architecture, help gestures are only anonymized by f_A), there is no correlation between d_M and d_A respectively d_A and d_S.

Let $G_{f_S(X_u)}$ be the set of guests that could be represented by the silhouette view $f_S(X_u)$ and let $g' = |G_{f_S(X_u)}|$. Then, the composition of the map view and the silhouette view does not yield any advantage over $(k - g')$-anonymity:

The operator $\mathcal{A}$ might be able to extract the camera alignment from a silhouette view $f_S(X_u)$. In the worst case, all guests are in rooms with a different camera alignment. Since the map view provides the operator $\mathcal{A}$ with this information, he knows that the silhouette view can not be a representation of any guest. In this case, the composition of d_M and d_S provides $(k - g')$-anonymity.

Therefore, and since from $n > k \geq g' + 2$ it follows $k - g' \geq 2$, the surveillance system provides at least 2-anonymity. This concludes the proof.

It turns out that the camera positions are the only factor that contributes to the correlation between map view and silhouette view, so given that all cameras are in the same position for each room, we can give an even stronger guarantee.

Corollary 2 *Assuming that f_M, f_A and f_S are $|\mathcal{U}|$-, 2- and 2-anonymizers respectively, as described in Section 4, and the cameras position is the same for every camera, an adversarial operator $\mathcal{A}$ cannot uniquely identify any individual in $\mathcal{U}$.*

References

[ABN08] Osman Abul, Francesco Bonchi, and Mirco Nanni. "Never Walk Alone: Uncertainty for Anonymity in Moving Objects Databases". In: *ICDE*. 2008, pp. 376–385.

[Bou05] T. E. Boult. "PICO: Privacy through Invertible Cryptographic Obscuration". In: *Proceedings of the Computer Vision for Interactive and Intelligent Environment*. CVIIE '05. 2005, pp. 27–38. ISBN: 0-7695-2524-5.

[Bun78] Deutscher Bundestag. *Bundesdatenschutzgesetz*. `http://www.gesetze-im-internet.de/bundesrecht/bdsg_1990/gesamt.pdf`. accessed 06/04/2013. 1978.

[CC11] Janusz Cichowski and Andrzej Czyzewski. “Reversible video stream anonymization for video surveillance systems based on pixels relocation and watermarking”. In: *ICCV Workshops*. 2011, pp. 1971–1977.

[Dwo06] Cynthia Dwork. “Differential Privacy”. In: *ICALP (2)*. 2006, pp. 1–12.

[FS08] Sven Fleck and Wolfgang Straqer. “Smart Camera Based Monitoring System and Its Application to Assisted Living”. In: *Proceedings of the IEEE* 96.10 (2008), pp. 1698–1714.

[GKS08] Srivatsava Ranjit Ganta, Shiva Prasad Kasiviswanathan, and Adam Smith. “Composition Attacks and Auxiliary Information in Data Privacy”. In: *KDD*. März 2008, pp. 265–273.

[HMH14] Eiji Hayashi, Manuel Maas, and Jason I. Hong. “Wave to me: user identification using body lengths and natural gestures”. In: *CHI*. 2014, pp. 3453–3462.

[KG06] Daniel Kifer and Johannes Gehrke. “l-Diversity: Privacy Beyond k-Anonymity”. In: *In ICDE*. 2006, p. 24.

[LLV07] Ninghui Li, Tiancheng Li, and Suresh Venkatasubramanian. “t-Closeness: Privacy Beyond k-Anonymity and l-Diversity”. In: *ICDE*. 2007, pp. 106–115.

[MRV10] Jürgen Moßgraber, Frank Reinert, and Hauke Vagts. “An Architecture for a Task-Oriented Surveillance System: A Service- and Event-Based Approach”. In: *Proceedings of the 2010 Fifth International Conference on Systems*. IEEE Computer Society, 2010, pp. 146–151.

[NS08] Arvind Narayanan and Vitaly Shmatikov. “Robust De-anonymization of Large Sparse Datasets”. In: *Proceedings of the 2008 IEEE Symposium on Security and Privacy*. SP ’08. 2008, pp. 111–125. ISBN: 978-0-7695-3168-7.

[Sen+03] Andrew Senior et al. “Blinkering Surveillance: Enabling Video Privacy through Computer Vision”. In: *IBM Research Report* 22886 (2003).

[Sen+08] Andrew W. Senior et al. “Enabling Video Privacy through Computer Vision.” In: *IEEE Security and Privacy* 3.3 (Aug. 18, 2008), pp. 50–57.

[SS98] Pierangela Samarati and Latanya Sweeney. “Generalizing Data to Provide Anonymity when Disclosing Information (Abstract)”. In: *PODS*. 1998, p. 188.

[Swe02] Latanya Sweeney. “k-Anonymity: A Model for Protecting Privacy”. In: *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10.5 (2002), pp. 557–570.

[VKB12] Hauke Vagts, Erik Krempel, and Jürgen Beyerer. “User-Centric Protection and Privacy in Smart Surveillance Systems”. In: *Future Security*. 2012, pp. 237–248.

[Zil07] Martin Zilkens. “Videoüberwachung - Eine rechtliche Bestandsaufnahme”. In: *Datenschutz und Datensicherheit* 31.4 (2007), pp. 279–283.