
A THREAT MODEL FOR SOFT PRIVACY ON SMART CARS

A PREPRINT

 **Mario Raciti**

IMT School for Advanced Studies Lucca
Lucca, Italy
Dipartimento di Matematica e Informatica
Università di Catania
Catania, Italy
mario.raciti@imtlucca.it

 **Giampaolo Bella**

Dipartimento di Matematica e Informatica
Università di Catania
Catania, Italy
giamp@dmi.unict.it

June 8, 2023

ABSTRACT

Modern cars are getting so computerised that ENISA’s phrase “smart cars” is a perfect fit. The amount of personal data that they process is very large and, yet, increasing. Hence, the need to address citizens’ privacy while they drive and, correspondingly, the importance of privacy threat modelling (in support of a respective risk assessment, such as through a Data Protection Impact Assessment). This paper addresses privacy threats by advancing a general modelling methodology and by demonstrating it specifically on *soft privacy*, which ensures citizens’ full control on their personal data. By considering all relevant threat agents, the paper applies the methodology to the specific automotive domain while keeping threats at the same level of detail as ENISA’s. The main result beside the modelling methodology consists of both domain-independent and automotive domain-dependent soft privacy threats. While cybersecurity has been vastly threat-modelled so far, this paper extends the literature with a threat model for soft privacy on smart cars, producing 17 domain-independent threats that, associated with 41 domain-specific assets, shape a novel set of domain-dependent threats in automotive.

Keywords privacy, risk assessment, threat modelling, automotive, LINDDUN.

1 Introduction

Smart Cars [ENISA(2019)], Smart Roads [Pompigna and Mauro(2022)] and Smart Cities [Toh and Martinez(2020)] are revolutionising the automotive domain while promising increased safety, efficiency, and sustainability. Such ecosystem generates vast amounts of data that require secure storage, transmission, and processing. Additionally, the integration of various sensors, cameras, and communication systems in modern vehicles creates new opportunities for privacy breaches, raising concerns about data protection measures and corresponding risks. Therefore, people’s privacy may be put at stake when they become car drivers. The 2023 Global Automotive Consumer Study by Deloitte [Deloitte(2023)] finds that consumers may be potentially open to sharing their personal information in exchange for access to vehicle health and maintenance updates, traffic congestion information, and updates for road safety. However, the study also reveals that trust issues, especially in Europe, are hindering consumers from sharing their personal data.

Also, privacy deserves additional scrutiny. While *hard privacy* concerns the various techniques to protect a subject’s personal data from everyone else, such as anonymisation and minimisation, *soft privacy* pertains to the range of practices to be followed for the subject to share their personal data with someone else while keeping full control, such as consent mechanisms and impact assessments. The General Data Protection Regulation (GDPR) is the European answer to the privacy needs of its citizens, and is proving inspirational for similar initiatives worldwide, along the lines of the so called “Brussels effect”. Our research rests on the observation that privacy issues in the automotive domain are not fully

understood at present, although they are certain to demand GDPR compliance. Compliance may be addressed in terms of privacy risk assessment, which in turn demands privacy threat modelling, hence the general motivation for this paper.

1.1 Context and Motivation

Threat modelling is challenging as the analyst faces various problems, such as completeness and threat explosion. On the one hand, completeness may be impactful because failing to account for specific threats would cause pitfalls to the subsequent risk assessment. On the other hand, the research of completeness might also result in threat explosion, which consists of a high number of threats that may not be relevant, feasible, or threats that are redundant. Completeness and redundancy are considered by Raciti and Bella [Raciti. and Bella.(2023)] through their systematic approach that leverages the LINDDUN state-of-the-art privacy threat modelling framework [Van Landuyt and Joosen(2020)] to elicit privacy threats for the automotive domain. The final list of threats online [Raciti and Bella(2023)] is produced by taking a domain-dependent approach and by embracing the threats from various sources, including in particular ENISA’s “Good practices for security of smart cars” [ENISA(2019)].

In this context, the motivations for the present paper are manifold:

- ENISA’s “Good practices for security of smart cars” is among the most relevant sources about car cybersecurity in Europe. However, its treatment of privacy is very limited, as we shall detail below, hence a deeper close-up is necessary.
- The privacy assessment by Raciti and Bella [Raciti. and Bella.(2023)] is very relevant. However, as the authors state, it is affected by subjectivity, hence the need for a more algorithmic approach to privacy threat elicitation, such as one of combinatoric nature.
- Hard and soft privacy, as noted above, distil out the two main spheres of privacy and demand dedicated scrutiny. While LINDDUN [Van Landuyt and Joosen(2020)] considers them to some extent, it does not specifically tailor them to the automotive domain.
- Cybersecurity and privacy are certain to intertwine but are often confused with one another. As we shall see below, existing efforts to distil them out, specifically in the automotive domain, are very few and worth expanding.

1.2 Research Question and Contributions

Following the context and motivations given above, this paper focuses on soft privacy in the automotive domain from the threat modelling perspective. It does so, by stating the following research question:

RQ What are the soft privacy threats for the automotive domain and how to elicit them?

It may be observed that the analyst may want to adopt a domain-dependent or domain-independent approach. The research question sets the scene for the former, meaning that it aims at threats that are specific for the identified target domain, i.e., automotive in this case.

This paper answers the research question by proposing an innovative methodology for privacy threat modelling. It applies a combinatoric approach that has two benefits. The first is the elicitation of domain-independent threats from analysing one or more sources from the state of the art. The second is the elicitation of domain-dependent threats, derived from the combination of a generic threat knowledge base and domain-specific assets. It also leverages LINDDUN [Deng et al.(2011)Deng, Wuyts, Scandariato, Preneel, and Joosen], which has become the de-facto standard privacy threat modelling framework and is inherently domain-independent. It adopts the mentioned ENISA report as a source of specific and comprehensive knowledge on the automotive domain.

Beside the methodology just outlined, this paper contributes a list of 17 soft privacy threats that are domain-independent but are also combined with 41 specific assets of the automotive domain, so as to produce a set of domain-dependent soft privacy threats for the automotive domain. As a first additional by-product, the paper also proposes the new list of domain-independent threats as a possible extension of the generic threat knowledge base provided by LINDDUN. A second positive side effect is the suggestion of domain-dependent threats as possible candidates to extend the existent threat taxonomy proposed by ENISA. In fact, a merge of the soft privacy threats with such a taxonomy provides a novel, comprehensive security-and-privacy-embracing threat taxonomy. In particular, this compensates for the lack of privacy-specific threats in the ENISA report, hence paving the way to a better understanding of the potential privacy risks associated with the automotive domain.

1.3 Article Summary

The organisation of the manuscript follows a simple waterfall style. Section 2 outlines the related work and Section 3 gives an overview of LINDDUN. Section 4 describes our privacy threat modelling methodology. Section 5 demonstrates the methodology by applying it to the automotive domain along with a case study, and Section 6 concludes.

2 Related Work

The challenges implicated by threat modelling led Wuyts et al. [Wuyts et al.(2019)Wuyts, Sion, Van Landuyt, and Joosen] to highlight the problems of current knowledge bases, such as limited semantics and lack of instantiating logic. Also, the authors discussed the requirements for a privacy threat knowledge base that streamlines threat elicitation efforts.

Furthermore, it is also noteworthy to recall that the process of threat modelling inherently implies assumptions and arbitrary decisions. Landuyt et al. [Van Landuyt and Joosen(2020)] highlighted the influence of assumptions to the outcomes of the analysis during the risk assessment process, more precisely in the threat modelling phase in the context of a LINDDUN privacy threat elicitation.

In addition, several attempts were made for the purposes of threat modelling in the automotive domain. Vasenev et al. [Vasenev et al.(2019)Vasenev, Stahl, Hamazaryan, Ma, Shan, Kemmerich, and Loiseaux.] were among the first to apply an extended version of STRIDE [Microsoft(2009)] and LINDDUN [Deng et al.(2011)Deng, Wuyts, Scandariato, Premeel, and Joosen] to conduct a threat analysis on security and privacy threats in the automotive domain. In particular, the case study is specific to long term support scenarios for over-the-air updates. Moreover, this work suggests that the privacy topic in the automotive domain has not reached the same level of maturity as cybersecurity.

In general, threat modelling is part of a wider process, that is risk assessment. Wang et al. [Wang et al.(2021)Wang, Wang, Qin, Ji, Zhang, and Wang] proposed a threat-oriented risk assessment framework tailored for the automotive domain, with the aim, among the others, of overcoming assumptions and subjectivity. This framework can be considered a precursor to ISO/IEC:21434 [ISO(2021)], which was defined a year later. Also, the authors applied STRIDE and the attack tree method for the threat modelling.

Moreover, Chah et al. [Chah et al.(2022)Chah, Lombard, Bkakria, Yaich, Abbas-Turki, and Galland] applied the LINDDUN methodology to elicit and analyse privacy requirements of CAV system, while respecting the privacy properties set by the GDPR. Such attempt represents a solid baseline for the broader process of privacy risk assessment tailored for the automotive domain. Finally, Bella et al. [Bella et al.(2023)Bella, Biondi, and Tudisco] advanced a dedicated risk assessment framework for privacy risks in modern cars. They proposed a double assessment, combining an asset-oriented ISO approach with a threat-oriented STRIDE approach.

The above works addressed crucial topics such as threat elicitation, threat knowledge base, privacy threat analysis and privacy risk assessment, both in general and specifically tailored to the automotive domain. However, on the best of our knowledge, there are no other works proposing a privacy threat modelling methodology that features the application of a combinatoric approach to elicit both domain-independent and domain-dependent soft privacy threats.

3 A Primer on LINDDUN

It is convenient to provide a brief introduction to LINDDUN before proceeding with the description of the methodology. LINDDUN is a privacy threat modelling methodology, inspired by STRIDE, that supports analysts in the systematical elicitation and mitigation of privacy threats in software architectures. LINDDUN privacy knowledge support represents one of the main strength of this methodology, and it is structured according to the 7 privacy threat categories encapsulated within LINDDUN's acronym [Deng et al.(2011)Deng, Wuyts, Scandariato, Premeel, and Joosen]:

- *Linkability* An adversary is able to link two items of interest without knowing the identity of the data subject(s) involved.
- *Identifiability* An adversary is able to identify a data subject from a set of data subjects through an item of interest.
- *Non-repudiation* The data subject is unable to deny a claim (e.g., having performed an action, or sent a request).
- *Detectability* An adversary is able to distinguish whether an item of interest about a data subject exists or not, regardless of being able to read the contents itself.
- *Disclosure of information* An adversary is able to learn the content of an item of interest about a data subject.

- *Unawareness* The data subject is unaware of the collection, processing, storage, or sharing activities (and corresponding purposes) of the data subject’s personal data.
- *Non-compliance* The processing, storage, or handling of personal data is not compliant with legislation, regulation, and/or policy.

The framework considers the state-of-the-art privacy properties, according to the terminology introduced by Pfitzmann [Pfitzmann and Hansen(2010)]. These are categorised as hard privacy and soft privacy properties. In particular, unlinkability, anonymity and pseudonymity, plausible deniability, undetectability and unobservability, and confidentiality (hiding data content, including access control) are under the umbrella of hard privacy; user content awareness (including feedback for user privacy awareness, data update and expire) together with policy and consent compliance are, on the other hand, soft privacy properties.

The LINDDUN framework provides a set of threats specific to privacy, named as “threat catalogue”, in the form of threat trees. These privacy threat trees are inspired by the Secure Development Lifecycle (SDL) [Howard and Lipner(2006)] and reflect common attack patterns [Wuyts and Joosen(2015)] on the basis of state-of-the-art privacy developments, structured according to LINDDUN or STRIDE threat category and Data Flow Diagram (DFD) element type. Such trees provide a formal way to describe the security of systems based on a variety of attacks. Basically, the root node represents the ultimate goal, e.g., the threatening to a property, the children nodes embody different ways of achieving that goal, namely refinements, hence leaves represent basic-level attacks that can not be further refined. In addition, non-leaf nodes can be conjunctive (logic AND) or disjunctive (logic OR) [Schneier(1999)].

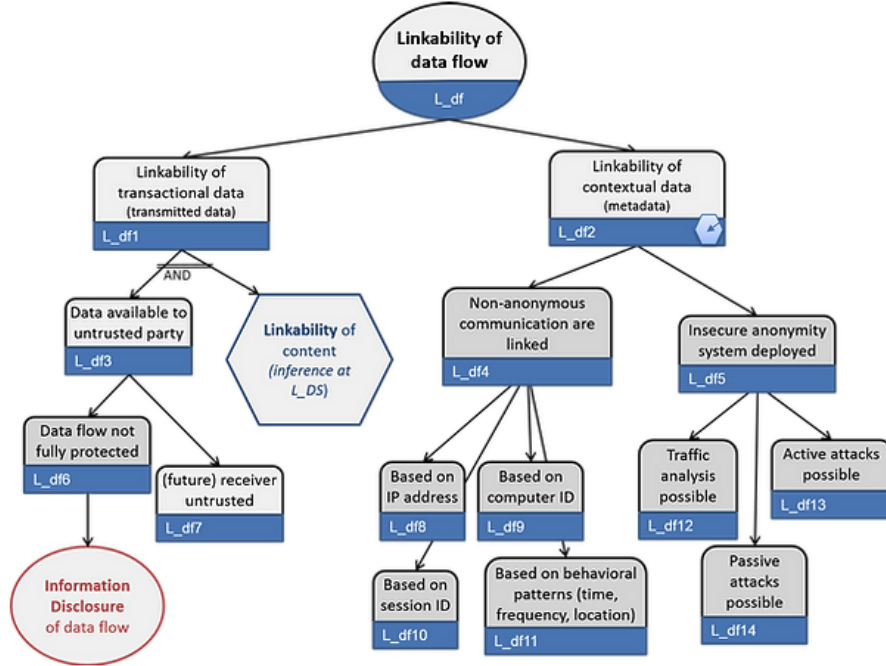


Figure 1: Example of a LINDDUN threat tree: Linkability of Data Flow.

An example tree is presented in Figure 1 for the Linkability of Data Flow, which can be achieved through L_{df1} “Linkability of the contextual data (transmitted data)”, e.g., IP address necessary for communication, and L_{df2} “Linkability of the transactional data (metadata)”, that is the actual data transmitted. Both of these provide various attack paths which are not necessarily limited to the LINDDUN property analysed, namely Linkability of data flow could lead to the Information Disclosure of data flow if we consider the following path, assuming the Linkability of content, namely inference, at data store applies:

$$L_{df1} \rightarrow L_{df3} \wedge L_{DS} \rightarrow L_{df6} \rightarrow ID_{df}$$

Despite LINDDUN threat trees may lack some semantics and have minimal selection criteria to express potential threats [Wuyts et al.(2019)Wuyts, Sion, Van Landuyt, and Joosen], they still aim at providing a valuable overview of potential threat types that seeks to be general, hence suitable for a privacy threat analysis of any application domain.

Moreover, the application of LINDDUN may lead to a high number of threats that may not be relevant, feasible, or important, thereby being labor-intensive and time-consuming, especially for complex or large systems. Hence, the advantage of having a catalogue of privacy threats, which are broad and applicable to various domains, may result in the problem of threat explosion.

4 A Privacy Threat Modelling Methodology

Privacy is a complex and multifaceted concept that may be interpreted in different ways in different contexts, yet in the first place it is a fundamental human right. In fact, NIST defined it as *the right of a party to maintain control over and confidentiality of information about itself* [NIST(1992)]. In a GDPR fashion, we may summarise privacy as the right of an individual, that is, the data subject, to control or influence what information related to them may be collected, processed and stored, and by whom and to whom that information may be disclosed.

Privacy is frequently intertwined with security, as privacy concerns frequently arise in relation to security matters. While threat modelling has traditionally been approached from a security perspective, it should be emphasised that privacy and security are two distinct concepts, hence they cannot be used interchangeably. In fact, the analyst should not rely on risk assessment models that are only security-centric, as these may not be sufficient to capture all aspects of privacy risks. A challenge for all privacy threat modelling approaches is how to consider the impact on data subjects involved in the privacy threat. This aspect is stressed in law and regulations compliance, i.e., in the Data Protection Impact Assessment (DPIA), required under the GDPR, to help identify, assess, and mitigate privacy risks associated with data processing activities. From a data subject's perspective, a DPIA is an important aspect of privacy threat modelling, as it ensures that organizations consider the potential impact of their data processing on individuals' privacy rights and take appropriate measures to address any risks. Arguably, a DPIA would benefit from a privacy threat model.

We advance a privacy threat modelling methodology that incorporates both domain-independent and domain-specific knowledge and considers the potential consequences on the privacy of individuals as its cornerstone. The pivotal notion that our methodology relies upon is the use of a combinatoric approach to elicit both domain-independent threats and domain-dependent threats. In particular, the former embody a generic threat knowledge base that consists of what is already known at present, whilst the domain-specific threats are derived from the first. Furthermore, our methodology identifies at least four variables that contribute to model privacy threats, i.e., the specific privacy property, the threat agents, the application domain and the level of detail. The combinatoric approach as well as each of the identified variables are discussed below.

4.1 The Combinatoric Approach

The approach consists of three steps:

1. Domain-Independent Threat Elicitation
2. Domain-Dependent Asset Collection
3. Domain-Dependent Threat Elicitation

The first step involves the collection of domain-independent threats from relevant sources. The second step consists of the collection of a list of assets for the target domain from relevant sources. The third and last step aims at producing a list of domain-specific threats. In particular, for each domain-independent threat elicited in Step 1, this step associates the assets enumerated in Step 2, hence obtaining a list of domain-dependent threats. The sheer association expresses the object of the threat that was domain-independent in the first place, thereby making it domain-dependent.

For the sake of convenience, the steps are noted through the creation of a table that should contain at least the following columns: Source, Threat, Assets. In particular, the source reflects the origin document or catalogue where the given threat was derived from, whereas the threat and assets columns are self-explanatory.

It is noteworthy that the subjectivity through this approach reduces to the collections to be done through the first two steps. Prose is not meant to be re-interpreted and modified by the analyst following step 3. By contrast, the sibling approach [Raciti. and Bella.(2023)] declares its subjectivity in the specific wording of each threat.

4.2 The Specific Privacy Property

Privacy relates to the control that individuals have over their personal information, including how it is collected, used, and shared. According to the state of the art [Danezis(2008), Deng et al.(2011)Deng, Wuyts, Scandariato, Preneel, and Joosen], we can distinguish between two degrees of privacy, namely hard privacy and soft privacy. We identify in

such properties the first element of scrutiny to build a privacy threat model, and contend that each of them deserves a specialised treatment.

Hard privacy refers to data minimisation, based on the assumption that personal data is not disclosed to third parties. The threat model includes service provider, data holder, and adversarial environment, where strategic adversaries with certain resources are motivated to breach privacy, similar to security systems [Deng et al.(2011)Deng, Wuyts, Scandariato, Preneel, and Joosen]. Examples of hard privacy include anonymisation of data, data minimisation, and data retention policies, including the use of algorithms such as k-anonymisation, t-closeness, differential privacy. For instance, a company that collects user data may anonymise the data before sharing it with third parties, ensuring that the users' identities remain protected. Similarly, a company may limit the amount of data it collects, processes, or stores to a minimum, and may have policies in place for deletion once processing that data is no longer necessary.

Soft privacy, on the contrary, is based on the assumption that the data subject lost control of their personal data and has to trust the honesty and competence of data controllers [Deng et al.(2011)Deng, Wuyts, Scandariato, Preneel, and Joosen]. Examples of soft privacy include transparency and consent mechanisms, data subject access rights, and privacy impact assessments. For example, a company may provide clear and concise privacy notices to inform users about how their data will be used and shared. They may also obtain users' consent before using their data for purposes beyond the original scope. Additionally, companies may, and should to be GDPR-compliant, provide users with the right to access, modify, or delete their personal data.

In summary, while hard privacy focuses on minimising the risks associated with the collection and retention of personal data, soft privacy focuses on the appropriate use and sharing of personal data while respecting individuals' rights to control their data. It is clear that, in addition to hard privacy and soft privacy, *cybersecurity* plays a major, complementary role in terms of protection against the unauthorised access of data.

Our methodology pays specific attention to both incarnations of privacy, but is described only on soft privacy in this paper due to space constraints.

4.3 The Threat Agents

Threat agents are individuals, groups, or systems that have the capability to exploit vulnerabilities and cause harm to a system or organisation. In threat modelling, understanding the capabilities, motivations, and objectives of threat agents is crucial for identifying and prioritising threats. There are different types of threat agents, including insiders, outsiders, script kiddies, hackers, and nation-state actors.

In the context of privacy threat modelling, we refer to a threat agent as any entity, individual or group, who poses a threat to an individual's privacy. Unlike the security literature, which traditionally refers to such entities as "adversaries" or "attackers", here the term threat agent also includes other sources of risks to privacy, as a threat agent is less security-annotated and not limited to malicious actors only. In fact, we also consider three additional actors directly from GDPR, i.e., data controller, data processor, and third party as threat agents. Therefore, a threat agent can be one or more of the following entities:

- *Attacker* Anyone, including an insider, or anything, including malware, acting with malicious intent to compromise a system to breach users' privacy.
- *Data controller* The natural or legal person, public authority, agency or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data;
- *Data processor* A natural or legal person, public authority, agency or other body which processes personal data on behalf of the controller;
- *Third party* A natural or legal person, public authority, agency or body other than the data subject, controller, processor and persons who, under the direct authority of the controller or processor, are authorised to process personal data.

This taxonomy allows us to better model the data subject's perspective during the threat modelling exercise. For instance, it might be relevant to understand whether to classify a threat against data as a data breach or data leak [ENISA(2022)]. A data breach is defined as *any breach of security leading to the accidental or unlawful destruction, loss, alteration or unauthorised disclosure of or access to personal data transmitted, stored or otherwise processed* (article 4.12 GDPR), while a data leak is an event that can cause the unintentional release of sensitive, confidential or protected data due to, for example, misconfigurations, vulnerabilities or human errors. Many use these notions interchangeably, yet they entail fundamentally different concepts that mostly lie in how they happen and what is the agent causing them. Data breach may be understood as an intentional attack brought by an attacker with the goal of gaining unauthorised access and the release of sensitive, confidential or protected data. In other words, a data breach is a deliberate and forceful attack

against a system or organisation with the intention of stealing data. On the other hand, a data leak may be understood as not including intentional attacks, and it may be attributed to, e.g., data controller and data processor.

Our methodology assumes that one or a combination of multiple threat agents may be involved in the threat scenario.

4.4 The Application Domain

The application domain in threat modelling identifies two prevailing approaches, i.e., domain-dependent and domain-independent. Domain-dependent threat modelling is specific to a particular application domain, such as healthcare, finance, or automotive, and it takes into account the unique characteristics of the domain itself, thus it may be more accurate and effective. On the other hand, domain-independent threat modelling is not specific to any application domain and can be applied to a wide range of systems. This approach uses generic threat categories, such as spoofing, tampering, and repudiation in STRIDE, to identify and prioritise threats. A generic threat knowledge base comes particularly useful in situations where there is no prior knowledge of the system or domain. LINDDUN, for example, is a domain-independent methodology.

A combination of the two approaches may offer a more effective and efficient analysis, picking the advantages of both. In addition, we may obtain the same results from a domain-dependent approach by starting from a generic threat knowledge base and associating it to domain-specific characteristics, i.e., assets. This is what, for instance, ENISA did in the study reported in “Good practices for security of smart cars” [ENISA(2019)]. They addressed the domain-dependent vs domain-independent dilemma by taking the combinatoric approach we mentioned above. Our methodology inherits the combinatoric approach, i.e., the combination of domain-independent threats with domain-dependent assets to elicit domain-dependent threats.

4.5 The Level of Detail

The last element of scrutiny derives from the level of detail — of the statement describing a threat. For example, “Unchanged default password” is certain to be more detailed than (the more abstract) threat “Human error”. Normally, the analyst strives to choose a consistent level of detail till the end of the exercise. The level of detail becomes relevant in the context of threat modelling and, subsequently, in risk assessment exercises with respect to the likelihood estimation of a threat.

The concepts of hyponym and hypernym are semantic relations between terms with respect to their hierarchy. Both play an important role in understanding the level of detail in a statement. In particular, hypernyms are broader, more general terms that encompass a category or set of items, whilst hyponyms are more specific, narrower terms that belong to a category or set. The relationship is asymmetric, meaning that while a hypernym can include many hyponyms, that is, a hypernym is an umbrella term, a hyponym can only have one hypernym.

A *higher level of detail* (hyponym) implies the analyst is able to estimate the likelihood of such a threat with more precision. However, an excessive level of detail leads to the degeneration of the threat to a “measurable event”, hence to an exact assignment of the likelihood, that is either the bottom or the top in the given range. If the analyst’s aim is to obtain a checklist of measurable events, a higher level of details represents the best option. However, the most appropriate level of detail, that is, the choice of employing hypernyms or hyponyms, should be considered within the main picture, and the analyst will choose it with some inevitable bias.

Our methodology addresses this variable by keeping the level of detail of the source knowledge base unaltered. This choice does not introduce subjectivity.

5 Demonstration in the Automotive Domain

We apply the privacy threat modelling methodology described above to address the research question. In particular, we set the variables discussed through Section 4 as follows:

P Soft Privacy

T Attacker, Data Controller/Processor, Third Party

D Domain-Dependent

L Abstract

5.1 Domain-Independent Threat Elicitation

Soft privacy is the target property, therefore we must consider the LINDDUN threats that refer to such property, i.e., U(nawareness) and N(on-compliance). For each node of the U and N property trees, we annotate the pertaining threat in a table. For the sake of demonstration, Figure 2 illustrates the threat tree for the U(nawareness) property.




Figure 2: LINDDUN threat tree: Unawareness.

As one of the aims in eliciting the list of soft privacy threats is completeness, we may also want to extend the list of domain-independent threats by adding two other sources. In particular, we include the 8 threats that were found by Raciti and Bella [Raciti. and Bella.(2023)] to be outstanding with respect to LINDDUN. In detail, they account for the 2 threats from the ENISA report that fall under the “Legal” category, i.e., “Failure to meet contractual requirements” and “Violation of rules and regulations/Breach of legislation/ Abuse of personal data”, and 6 threats from the “Calculation of the complete Privacy Risks list v2.0” [OWASP(2021)] document, i.e., “Consent-related issues”, “Inability of user to access and modify data”, “Insufficient data breach response”, “Misleading content”, “Secondary use”, “Sharing, transfer or processing through 3rd party”. These threats relate to soft privacy as per the definition of soft privacy we covered previously in Section 4.2. It is noteworthy to mention that this addition is still possible without consequences on the domain variable, as such threats are general privacy threats, i.e., they ignore domain specific entities. Hence, such threats can be analysed in relation with (virtually) any application domain. Table 1 illustrates our results, which represent a first by-product of this paper as generic, domain-independent soft privacy threats, following this step.

It is convenient to provide a brief and generic explanation of the threats, referring to the descriptions provided by their sources. In particular, U(nawareness) means being unaware of the consequences of sharing information. Users are often not aware of the impact of sharing data on social networks or with other services. Lack of feedback and awareness tools can cause users to be unaware of the information they are sharing. It is important that privacy support is user-friendly and that users have the ability to modify privacy settings. Data subjects are often unaware of what data a system has collected and stored about them, so they should always have the ability to review their own data.

Furthermore, N(on-compliance) means not being compliant with legislation, regulations, and corporate policies. Non-compliance can lead to fines, loss of image and credibility. Users should have control over their data and be able to decide who has access to it through consent. Privacy policies and consents need to be stored correctly and consistently to ensure compliance. Lack of user-friendly policy management and insufficient implementation can lead to non-compliance with user consent requirements. Clear and transparent notice of applied policies is necessary to inform users about data collection, storage, and processing.

Table 1: Domain-independent soft privacy threats elicited in Step 1.

Source	Threat
U	Providing too much personal data
	Unaware of stored data
	No/insufficient feedback and awareness tools
	No user-friendly privacy support
	Unable to review personal information (data accuracy)
N	Attacker tampering with privacy policies and makes consents inconsistent
	Incorrect or insufficient privacy policies
	Inconsistent/insufficient policy management
	Insufficient notice
ENISA	Failure to meet contractual requirements
	Violation of rules and regulations/Breach of legislation/Abuse of personal
OWASP	Consent-related issues
	Inability of user to access and modify data
	Insufficient data breach response
	Misleading content
	Secondary use
	Sharing, transfer or processing through 3rd party

In addition, the “Legal” threat category by ENISA highlights that not following contractual requirements can result in negative consequences such as financial loss, safety issues, privacy breaches, and operational disruptions. Moreover, non-compliance with international or European regulations and laws, e.g., GDPR and UNECE regulations, may have a significant impact on users’ privacy, such as the unauthorised exposure of their personal information.

Finally, the threats originating from OWASP raise concerns about legal, technological and organisational aspects that focus on real-life risks. These include: collection of data without the user’s consent or aggregating consent; the lack of ability for users to change or delete incorrect data relating to them; not informing the affected persons (data subjects) about a possible breach or data leak; user misled through confusing, ambiguous, or misleading content or descriptions of their privacy practices; operators’ use personal data in ways unrelated to the original purpose of the collection, without even informing affected people of the change; providing user data to any third-party, without obtaining the user’s consent, and/or lack of clarity about policies or protocols employed by the third party.

5.2 Domain-Dependent Asset Collection

For Step 2, we leverage two sources from the state of the art, i.e., the assets identified in the work proposed by Bella et al. [Bella et al.(2023)Bella, Biondi, and Tudisco] and the ENISA taxonomy of the key assets in the automotive domain. The former presents the following list of assets: “Personally Identifiable Information”, “Special categories of personal data”, “Driver’s behaviour”, “User preferences”, “Purchase information”, “Smartphone data”, “GPS data”, “Vehicle information”, “Vehicle maintenance data”, “Vehicle sensor data”. The report by ENISA focuses on Automated Driving System-Dedicated Vehicle (ADS-DS) [SAE(2021)], i.e., semi-autonomous and autonomous cars, and V2X communications. Figure 3 depicts such SAE vehicles automation levels. The focus of the study is on smart cars that, as connected systems, have the necessary capabilities to autonomously perform all driving functions under certain (or all) conditions, and are able to communicate with their surroundings including other vehicles, pedestrians and Road-Side Units (RSU). Moreover, the key concepts analysed by ENISA do not only concern passenger cars but also commercial vehicles (e.g. buses, coaches, etc.), including self-driving, ride-sharing vehicles that can be shared with other users.

The assets proposed by ENISA are categorised in: “Car sensors and actuators”, “Decision Making Algorithms”, “Vehicle Functions”, “Software management”, “Inside vehicle Communication Components”, “Communication Networks and Protocols”, “Nearby External Components”, “Network and Domain Isolation Features”, “Servers”, “Systems and Cloud Computing”, “Information”, “Humans”, “Mobile Devices”. For the sake of brevity, we only quote the descriptions of the assets under the “Information” category, which are inherently most likely to be affected by soft privacy threats:

- *Sensors data* This asset refers to data that is gathered by the different smart car sensors and which will be transmitted to the appropriate ECU for processing.




Figure 3: SAE vehicles automation levels as defined in SAE J3016. Source: ENISA

- *Keys and certificates* This asset refers to the different keys and certificates used for security purposes (such as authentication, securing the exchanges, secure boot, etc.). Keys are stored in devices embedded in the vehicle (e.g. ECU) and/or in servers depending on their use.
- *Map data* This asset refers to the information about the car environment. Map data allows us to increase the passenger safety by correlating its information with the sensor data. Contrary to GNSS, which gives only information about the geolocalisation, map data gives information about the surrounding environment.
- *V2X information* This asset refers to the various information exchanged via V2X communications (e.g. emergency vehicle approaching, roadworks/collision warning and traffic information).
- *Device information* This asset refers to the various information related to a device embedded in a smart car (e.g. ECU, TCU) or connected devices (e.g. smartphones, tablet). This includes information such as type, configuration, firmware version, status, etc., of different smart car sensors and which will be transmitted to the appropriate ECU for processing.
- *User information* This asset refers to smart cars user (e.g. driver, passenger, etc.) information such as name, role, privileges and permissions.

Moreover, soft privacy is part of privacy, which is related to security, thereby all the assets proposed by ENISA may be virtually involved in the execution of the combinatoric approach.

During the execution of this step, we identified some redundancy between the two sources, that is, some assets proposed by Bella et al. are *embraceable* with the ENISA taxonomy. Thereby, we explicitly picked the following assets from the paper contribution: “Special categories of personal data”, “Driver’s behaviour”, “User preferences”, “Purchase information”. The remaining assets, according to our scrutiny, are already contemplated in ENISA taxonomy.

At this point, one may argue that also threats from Bella et al. could be considered, with respect to the previous step. However, those threats were mapped to hard privacy properties in Raciti and Bella’s contribution, hence they are not relevant for the target of this paper, that is, soft privacy. Yet, they would be significant for an exercise tailored to that particular property.

5.3 Domain-Dependent Threat Elicitation

In the last step, we associate the results of the previous steps. For each domain-independent threat elicited in Step 1, we assign the assets from Step 2 that we deem to be potentially affected by that particular threat. In general, a threat may apply to multiple assets, therefore for some threat-asset pairs we annotate multiple ENISA assets or, in case all assets are affected, we add the label “All assets” for the sake of brevity. In particular, most assets that we deem to be potentially affected by the threats fall under the ENISA category “Information”.

While Table 2 summarises the results, which represent the second by-product of this paper, following the last step, we may present an exemplification of each domain-dependent threat, with the additional aim of providing the reason behind the various threat-asset(s) associations:

- *Providing too much personal data* Users may be required to provide an excessive amount of personal data to access certain features on the infotainment system of the vehicle or to use a connected service, potentially putting their privacy at risk.
- *Unaware of stored data* Users may not be aware of what data is being collected and stored by the onboard systems of their car, leading to potential privacy violations.

- *No/insufficient feedback and awareness tools* OEM and/or car-related services may not provide adequate feedback or awareness tools to inform users about what data is being collected and how it is being used, leading to potential privacy violations.
- *No user-friendly privacy support* OEM and/or car-related services may lack user-friendly privacy support, such as clear and accessible privacy policies or easy-to-use privacy controls, leaving users unsure about how to protect their privacy.
- *Unable to review personal information (data accuracy)* Users may be unable to review or correct their personal information that is being collected and stored by the onboard systems of the car, potentially leading to inaccurate data being used for decision-making.
- *Attacker tampering with privacy policies and makes consents inconsistent* Attackers may attempt to tamper with privacy policies of the car or manipulate a user’s consent, leading to potential privacy violations.
- *Incorrect or insufficient privacy policies* OEM and/or car-related services may have privacy policies that are incorrect or insufficient, leaving users unaware of what data is being collected and how it is being used.
- *Inconsistent/insufficient policy management* OEM and/or car-related services may not have effective policy management processes in place, leading to potential privacy violations.
- *Insufficient notice* OEM and/or car-related services may not provide sufficient notice to users about what data is being collected and how it is being used, potentially leading to privacy violations.
- *Failure to meet contractual requirements* OEM and/or car-related services may not meet the contractual privacy requirements agreed upon with users, leading to potential privacy violations.
- *Violation of rules and regulations/Breach of legislation/Abuse of personal data* OEM and/or car-related services may violate rules and regulations or abuse personal data, potentially leading to privacy violations.
- *Consent-related issues* Users may not fully understand what they are consenting to when using connected services or providing personal data to their car’s onboard systems, potentially leading to privacy violations.
- *Inability of user to access and modify data* Users may be unable to access or modify their personal information that is being collected and stored by their car’s onboard systems, potentially leading to privacy violations.
- *Insufficient data breach response* OEM and/or car-related services may not have effective processes in place for responding to data breaches, potentially leading to privacy violations.
- *Misleading content* OEM and/or car-related services may provide misleading content to users about what data is being collected and how it is being used, potentially leading to privacy violations.
- *Secondary use* Personal data collected by OEM and/or car-related services may be used for secondary purposes without the user’s consent, potentially leading to privacy violations.
- *Sharing, transfer or processing through 3rd party* Personal data collected by OEM and/or car-related services may be shared, transferred or processed by third parties without the user’s consent or knowledge, potentially leading to privacy violations.

5.4 Case Study

This section presents a case study that relies on the latest breaking news and articles about privacy incidents in the automotive domain. In particular, we employ classical web searches as a source of relevant information by building queries as “privacy automotive”, “automotive breach”, “smart car privacy”, et similia, in the *News* search filter offered by Google. If we matched some news with a soft privacy threat from the previous exercise, then we would be able to give some statistics about the occurrences of such threat, hence inferring an estimation of its likelihood. For the sake of brevity, we only present some illustrative examples of news that matched with one or more of the proposed soft privacy threats.

A data breach at Toyota Motor’s Indian business [Reuters(2023)] might have exposed some customers’ personal information. “Toyota Kirloskar Motor (TKM) has been notified by one of its service providers of an incident that might have exposed personal information of some of TKM’s customers on the internet”. This piece of news perfectly embodies two threats that we find in Table 2, i.e., “Insufficient data breach response” and “Violation of rules and regulations/Breach of legislation/Abuse of personal data”, because of GDPR compliance.

Furthermore, we also find another news that represents the threat “Violation of rules and regulations/Breach of legislation/Abuse of personal data”. The Dutch Data Protection Authority (DPA) investigated Tesla’s camera-based “Sentry Mode” security system [Europe(2023)], which is designed to protect the vehicle against theft or vandalism while it is parked. It does this by taking footage with four cameras on the outside of the vehicle. This specific threat

has now received a mitigation measure from the manufacturer, as the company altered security cameras to be more privacy-friendly and avoid GDPR violations. Originally, when Sentry Mode was enabled, this system was on by default. The cameras continuously filmed everything around a parked Tesla and stored one hour of footage each time.

In addition, we also found a review [underscored(2023)] that perfectly matched with the implications related to the threat “No user-friendly privacy support”. The article discusses a suggestion for a new feature to be added to the Ring Car Cam. The author proposes an Alexa-based voice command that would temporarily turn off the interior camera and microphone. This suggestion is based on the author’s wife’s volunteer work, which involves discussing private and privileged information about children’s legal cases on the phone. The author’s wife currently uses the physical privacy shutter to prevent the camera from recording video and audio inside the car. However, she sometimes forgets to flip the shutter up or down. Therefore, the author proposes a hands-free privacy trigger that would allow the user to enable or disable privacy mode with a voice command. This feature would eliminate the need for the user to physically manipulate the shutter, making it easier to maintain privacy while driving.

5.5 Findings

In this section, we discuss the findings from the previous experiment. The application of the combinatoric method to the automotive domain yielded notable results, which are available in Table 2, as stated above. In particular, we produced a novel, refined list of soft privacy threats that are domain-dependent. In fact, if we associate the generic threat knowledge base pertaining to soft privacy, collected at the end of Step 1, with the automotive-specific assets collected at the end of Step 2, we obtain domain-specific soft privacy threats for modern cars with a homogeneous level of detail and dependent on the automotive domain. This answers the research question. A confirmation of the concreteness and relevance of such threats for the automotive domain was proven, although informally, with the employment of web searches. Furthermore, the ENISA report is more tailored to security, in terms of the *property* variable, thereby the new list of threats may be taken as possible candidates to expand the ENISA report, hence enriching the broader threat knowledge base in the automotive domain over soft privacy. Moreover, we cannot claim that no more valid candidates exist, yet the list of threats results complete with respect to the state-of-the-art privacy threat knowledge base, i.e., LINDDUN threat catalogue and automotive asset taxonomy, i.e., ENISA. Nonetheless, these threats remain valid candidates for the international community’s evaluation.

6 Conclusions

This paper faced the challenge of privacy threat modelling focusing specifically on soft privacy and on the automotive domain, as its research question states. The research question found an answer in the advancement of a threat modelling methodology that is general for privacy, as it revolves around (at least) four variables. Moreover, while the methodology ultimately aims at eliciting domain-dependent threats, the combinatoric approach it takes has two additional benefits. One is that, by collecting threats from relevant albeit scattered sources, it does not change their level of detail hence it does not introduce (additional) subjectivity. The other one is that, by taking a combinatoric approach through threat elicitation, it produces a set of domain-independent threats but, at the same time, a set of automotive domain-dependent threats.

The methodology is then demonstrated on soft privacy, yielding 17 domain-independent threats, which can be tailored to 41 assets to become automotive domain-dependent. We argue that soft privacy has not received such an in-depth treatment thus far, considering that our set of threats is appreciably larger than both LINDDUN’s and ENISA’s, and it clearly introduces more facets of soft privacy. A natural follow-up of our work is to analyse the application of the combinatoric approach proposed by our methodology to different tuples of variables, for example addressing hard privacy at a very high level of detail.

In particular, to achieve a fuller account on the level of detail during the elicitation process, we expect to be able to leverage modern, intelligent techniques from the area of Natural Language Processing. For example, we are currently looking at the concepts of hypernym and hyponym towards the definition of a metrics to assist the analyst through the choice of a consistent level of detail for threat descriptions. An interpretation of the findings of this paper is that (soft) privacy is finally threat-modelled as extensively as it deserves and, in particular, as cybersecurity traditionally has been. In consequence, the risks for “*natural persons with regard to the processing of personal data and on the free movement of such data*” [GDPR(2016)], particularly when those natural persons drive smart cars, can be now assessed more precisely than before.

Table 2: Results of the application of our methodology to the automotive domain: list of soft privacy threats.

Source	Threat	Assets
U	Providing too much personal data	User information, Special categories of personal data
	Unaware of stored data	Map data, V2X information, Device information, User information, Special categories of personal data, User preferences, Purchase information
	No/insufficient feedback and awareness tools	Map data, Device information, User information, Special categories of personal data, Driver's behaviour, User preferences, Purchase information
	No user-friendly privacy support	Sensors data, Map data, V2X information, Device information, User information, Special categories of personal data, Driver's behaviour, User preferences, Purchase information
N	Unable to review personal information (data accuracy)	User information, Special categories of personal data
	Attacker tampering with privacy policies and makes consents inconsistent	Sensors data, Key and certificates, Map data, V2X information, Device information, User information, Special categories of personal data, Driver's behaviour, User preferences, Purchase information
	Incorrect or insufficient privacy policies	All assets
	Inconsistent/insufficient policy management	All assets
ENISA	Insufficient notice	Sensors data, Key and certificates, Map data, V2X information, Device information, User information, Special categories of personal data
	Failure to meet contractual requirements	All assets
	Violation of rules and regulations/Breach of legislation/Abuse of personal data	All assets
	Consent-related issues	All assets
OWASP	Inability of user to access and modify data	Map data, V2X information, Device information, User information, Special categories of personal data, User preferences, Purchase information
	Insufficient data breach response	Sensors data, Key and certificates, Map data, V2X information, Device information, User information, Special categories of personal data, User preferences, Purchase information
	Misleading content	Map data, V2X information, Device information, User information, Special categories of personal data, User preferences
	Secondary use	All assets
	Sharing, transfer or processing through 3rd party	Sensors data, Key and certificates, Map data, V2X information, Device information, User information, Special categories of personal data, Driver's behaviour, User preferences, Purchase information

References

- ENISA. Good Practices for Security of Smart Cars. <https://www.enisa.europa.eu/publications/smart-cars>, 2019.
- Andrea Pompigna and Raffaele Mauro. Smart roads: A state of the art of highways innovations in the smart age. *Engineering Science and Technology, an International Journal*, 25:100986, 2022. ISSN 2215-0986. doi:<https://doi.org/10.1016/j.jestch.2021.04.005>. URL <https://www.sciencedirect.com/science/article/pii/S2215098621000872>.
- Cano Toh, Sanguesa and Martinez. Advances in smart roads for future smart cities. In *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, volume 476, 2233, 2020. doi:10.1098/rspa.2019.0439.
- Deloitte. 2023 global automotive consumer study. <https://www.deloitte.com/global/en/Industries/automotive/perspectives/global-automotive-consumer-study.html>, 2023.
- Mario Raciti and Giampaolo Bella. How to model privacy threats in the automotive domain. In *Proceedings of the 9th International Conference on Vehicle Technology and Intelligent Transport Systems - VEHITS*, pages 394–401. INSTICC, SciTePress, 2023. ISBN 978-989-758-652-1. doi:10.5220/0011998800003479.
- Dimitri Van Landuyt and Wouter Joosen. A descriptive study of assumptions made in linddun privacy threat elicitation. In *Proceedings of the 35th Annual ACM Symposium on Applied Computing, SAC '20*, page 1280–1287, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450368667. doi:10.1145/3341105.3375762. URL <https://doi.org/10.1145/3341105.3375762>.
- Mario Raciti and Giampaolo Bella. Github repository with complete outcomes. <https://github.com/tsumarios/LINDDUN-threats-completeness>, 2023.
- Mina Deng, Kim Wuyts, Riccardo Scandariato, Bart Preneel, and Wouter Joosen. A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. *Requirements Engineering*, 16(1):3–32, Mar 2011. ISSN 1432-010X. doi:10.1007/s00766-010-0115-7. URL <https://doi.org/10.1007/s00766-010-0115-7>.
- Kim Wuyts, Laurens Sion, Dimitri Van Landuyt, and Wouter Joosen. Knowledge is power: Systematic reuse of privacy knowledge for threat elicitation. In *2019 IEEE Security and Privacy Workshops (SPW)*, pages 80–83, 2019. doi:10.1109/SPW.2019.00025.
- Alexandr Vasenev, Florian Stahl, Hayk Hamazaryan, Zhendong Ma, Lijun Shan, Joerg Kemmerich, and Claire Loiseaux. Practical security and privacy threat analysis in the automotive domain: Long term support scenario for over-the-air updates. In *Proceedings of the 5th International Conference on Vehicle Technology and Intelligent Transport Systems - VEHITS*, pages 550–555. INSTICC, SciTePress, 2019. ISBN 978-989-758-374-2. doi:10.5220/0007764205500555.
- Microsoft. The STRIDE threat model. <https://learn.microsoft.com/en-us/previous-versions/commerce-server/ee823878%28v=cs.20>, 2009. URL <https://learn.microsoft.com/en-us/previous-versions/commerce-server/ee823878%28v=cs.20>.
- Yunpeng Wang, Yinghui Wang, Hongmao Qin, Haojie Ji, Yanan Zhang, and Jian Wang. A systematic risk assessment framework of automotive cybersecurity. *Automotive Innovation*, 4(3):253–261, Aug 2021. ISSN 2522-8765. doi:10.1007/s42154-021-00140-6. URL <https://doi.org/10.1007/s42154-021-00140-6>.
- ISO. ISO/IEC 21434 - road vehicles — cybersecurity engineering, 2021.
- Badreddine Chah, Alexandre Lombard, Anis Bkakra, Reda Yaich, Abdeljalil Abbas-Turki, and Stéphane Galland. Privacy threat analysis for connected and autonomous vehicles. *Procedia Computer Science*, 210:36–44, 2022. ISSN 1877-0509. doi:<https://doi.org/10.1016/j.procs.2022.10.117>. URL <https://www.sciencedirect.com/science/article/pii/S1877050922015733>. The 13th International Conference on Emerging Ubiquitous Systems and Pervasive Networks (EUSPN) / The 12th International Conference on Current and Future Trends of Information and Communication Technologies in Healthcare (ICTH-2022) / Affiliated Workshops.
- Giampaolo Bella, Pietro Biondi, and Giuseppe Tudisco. A double assessment of privacy risks aboard top-selling cars. *Automotive Innovation*, pages 1–18, 01 2023. ISSN 2522-8765. doi:10.1007/s42154-022-00203-2. URL <https://doi.org/10.1007/s42154-022-00203-2>.
- Andreas Pfitzmann and Marit Hansen. A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management. http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf, August 2010. URL http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf. v0.34.
- Michael Howard and Steve Lipner. *The Security Development Lifecycle*, volume 34. Microsoft Press, 06 2006. ISBN 0735622140. doi:10.1007/s11623-010-0021-7.

- Kim Wuyts and Wouter Joosen. Linddun privacy threat modeling: a tutorial, Jul 2015. Technical Report (CW Reports),.
- Bruce Schneier. Attack trees. *Dr. Dobb's journal*, 24(12):21–29, 1999.
- NIST. Foundations of a Security Policy for Use of the National Research and Educational Network (NISTIR 4734). <https://nvlpubs.nist.gov/nistpubs/Legacy/IR/nistir4734.pdf>, 1992.
- G. Danezis. Introduction to Privacy Technology. http://www0.cs.ucl.ac.uk/staff/G.Danezis/talks/Privacy_Technology_cosic.pdf, 2008.
- ENISA. ENISA Threat Landscape 2022. <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2022>, 2022.
- OWASP. Top 10 Privacy Risks. <https://owasp.org/www-project-top-10-privacy-risks/>, 2021.
- SAE. Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles (J3016_202104). https://www.sae.org/standards/content/j3016_201806/, 2021.
- Reuters. Toyota's indian unit warns of a possible customer data breach. <https://www.reuters.com/technology/toyotas-indian-unit-warns-possible-customer-data-breach-2023-01-01/>, 2023.
- Automotive News Europe. Tesla escapes fine from dutch watchdog after automaker alters security cameras. <https://europe.autonews.com/automakers/tesla-alters-cameras-avoid-dutch-fine-over-privacy-violations>, 2023.
- CNN underscored. The ring car cam takes ring's great security smarts on the road. <https://edition.cnn.com/cnn-underscored/reviews/ring-car-cam>, 2023.
- GDPR. Regulation (EU) 2016/679 General Data Protection Regulation. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32016R0679>, 2016.