

SPIDER: Interplay Assessment Method for Privacy and Other Values

Zoltán Ádám Mann
University of Amsterdam
Amsterdam, The Netherlands

Jonathan Petit
Qualcomm
Boxborough, MA, USA
petit@qti.qualcomm.com

Sarah M. Thornton
Nuro
Mountain View, California, USA
smthorn@alumni.stanford.edu

Michael Buchholz
Universität Ulm, Institut für Mess-, Regel- und Mikrotechnik
Ulm, Germany
michael.buchholz@uni-ulm.de

Jason Millar
University of Ottawa, Faculty of Engineering
Ottawa, Canada
jmillar@uottawa.ca

Abstract—In the design of many sociotechnical systems, ensuring people’s privacy is crucial. Available strategies, patterns, and technologies for ensuring privacy are often associated with drawbacks with respect to other values, such as security, fairness, or safety. Thus, system design entails navigating such value interactions, aiming to find solutions that reconcile privacy and other values. However, no systematic methodology is available for assessing the interplay between privacy and other values in a design.

To solve this problem, we propose SPIDER, a methodology for the systematic assessment of the interplay between privacy and other values. With SPIDER, system designers can investigate, quantify, and visualize the type (positive/neutral/negative) and strength of the interplay between privacy and other values, from different stakeholders’ point of view. This helps identify areas where further improvement of the design is needed to resolve tensions between privacy and other values. We demonstrate the application of SPIDER in the domain of Cooperative Connected Automated Mobility (CCAM) on a use case of an automated delivery vehicle.

1. Introduction

Complex sociotechnical systems, such as Cooperative Connected Automated Mobility (CCAM), rely on data collection and sharing for their safe operation. In such systems, the system’s awareness and understanding of humans is often crucial. For example, self-driving cars can only become a reality if they are able to account for the needs of involved humans (e.g., passengers, other street users) [3].

However, the system’s awareness and understanding of humans, although necessary, may raise privacy concerns. For example, CCAM systems often collect large amounts of personal data. The collected data may be valuable also beyond the primary purpose of the given system, creating incentives for service providers to collect more data than strictly necessary or using the collected data for other purposes [33]. Such violations of data privacy can lead to various forms of harm [7]. In CCAM, for example, video

feeds from the cameras of the vehicles may be needed for safety reasons, but could also be used to spy on people.

The privacy research community has devised a large array of techniques that can be used to ensure that data processing preserves people’s privacy as much as possible [2]. However, many of the existing privacy-enhancing technologies have some drawbacks, such as additional computation and communication overhead [22]. The result is a wide-spread *belief* that there is an *intrinsic tension between privacy and other values* (e.g., efficiency, usability) in the design process [28]. However, this is not always true. In some cases, there is indeed a strong tension between privacy and other values, while in other cases, the tension may be weaker or non-existent, or privacy and other values may even reinforce each other. For example, storing, transmitting, and processing less data improves both privacy and efficiency.

An important task of privacy engineering is to find ways to decrease tensions between privacy and other values [5]. We need system designs that foster privacy *and* other values at the same time. Thus, we need design methods that support the analysis of the interplay between privacy and other values, to direct designers’ attention to areas where conflicts between privacy and other values still exist. This is important, so that these conflicts can be relieved by improvements of the system design, or at least explicit and sound decisions can be made about handling the conflicts.

To this end, this paper introduces SPIDER: Systematic Privacy Interplay Deriving Ethical Requirements. SPIDER is a methodology that can be used during the system design process to systematically assess a preliminary design in terms of the interplay between privacy and other relevant values. Thereby, SPIDER helps identify areas where there are tensions between privacy and other values. System designers can then investigate how these areas could be improved, and revised designs can be re-assessed with SPIDER to evaluate the improvement. This paper introduces SPIDER in general, and validates it by applying it in the CCAM domain, to the design of a fictional automated delivery vehicle.

The remainder of this paper is structured as follows. Section 2 presents related work. Section 3 describes the SPIDER methodology, which we apply to a CCAM use

case in Section 4. We discuss our findings and lessons learned in Section 5 and conclude the paper in Section 6.

2. Related Work

We are not aware of an existing methodology for systematically assessing the interplay between privacy and other values in a system’s design. However, some related approaches exist in the fields of value-focused design and privacy impact assessments. Also, CCAM privacy has been subject of relevant research, showing examples for the interplay between privacy and other values.

Value-focused design methodologies. Different design frameworks have been proposed to help bring human values, including privacy, into technology design. Value Sensitive Design (VSD), as originally presented by Friedman et al. [12], [13], details a generic design framework that centers human values with ethical import in the design process. However, VSD, as originally described, remains quite open-ended, which can present challenges when attempting to embed it in engineering processes [24]. Thus, VSD variations have been proposed in order to formalize the framework for engineering applications. Thornton [32] explicitly traces human values to terms in a partially observable Markov decision process, such as the discretization of the state space, as well as terms in the reward function. Millar et al. [24] outline a step-by-step engineering design process in which VSD is embedded, and describe concrete steps that can be taken to help resolve value tensions that are uncovered during that process. However, their process does not provide means for quantifying those value tensions in a way that allows for a side-by-side analysis of them.

Goal modeling can also be used to model privacy and other values, as well as relationships between values and actors or between values and system processes [10], [37]. Such approaches support conceptual modeling in requirements analysis and system design, contributing to a better understanding of design options and their impact on privacy and other values. However, such models do not provide a way to assess and concisely represent the interplay between privacy and other values in a design.

Privacy impact assessments. To identify privacy risks in a planned system as well as mitigation techniques, it is common best practice to perform a Privacy Impact Assessment (PIA). The standard ISO/IEC 29134 [1] gives guidelines for a process on PIA, and a structure and content of a PIA report. Panda et al. [26] applied a PIA to CCAM with the objective to ensure that the NIST Privacy Engineering Objectives (i.e., Predictability, Manageability, Disassociability) [4] are met to reduce the identified privacy risks. However, a PIA focuses on dataflows and does not consider other design considerations nor highlights interplay between values.

CCAM privacy. Data collection and sharing introduce privacy risks to CCAM systems. But data collection and sharing among CCAM parties is essential for many reasons, including efficient traffic management or for the accurate assignment of liability in the event of a collision [8]. Glancy [14] investigated the privacy of automated vehicles from a legal standpoint. She concluded that “the future success of autonomous vehicles will depend in part on how well privacy interests and autonomous vehicles can

work together”, underscoring the need to design CCAM with privacy in mind. Privacy often comes into tension with other values. Lim et al. [21] highlight the risk of the increasing power disparity between organisations that control data and individuals about whom data is collected. Kohler and Colbert-Taylor [18] discuss how privacy risks may be avoided, but at the expense of the service provider’s competitive advantage [34]. Similarly, Kountche et al. [19] highlight that the privacy of CCAM users can be undermined by the data that could be harvested from them, hence recommending the application of privacy by design principles—data minimization and strict access control. However, they recognize that the desire for privacy is generally incompatible with the desire for accountability, and that a trade-off should be found between privacy and accountability. In CCAM, location privacy is paramount and one technique to prevent tracking is to digitally sign vehicular communications with short-term identifiers to provide pseudonymity while ensuring non-repudiation [27]. However, pseudonyms have impact on CCAM performance (e.g., safety) [9], [20], [36].

These examples demonstrate how privacy inevitably interacts with other values in the CCAM design space. However, we are not aware of any methodology that would systematically address such interactions. Hence, there is a need for analytic methods that focus engineers’ attention on the interactions between privacy and other values in system design. Our methodology prompts designers to holistically investigate the interplay between privacy and other values to support privacy engineering.

3. The SPIDER Methodology

The goal of SPIDER is a systematic assessment of the interplay between the *core value*, “privacy”, and other values in the design space, such as safety and fairness, during the design of complex sociotechnical systems.

3.1. Basic notions

Our methodology (SPIDER) employs a set of concepts and relationships, as summarized in Figure 1. SPIDER assesses the interactions between privacy and a set of other values. Each interaction is evaluated in terms of interplay type and strength. The *interplay type* can be either positive (e.g., improving the other value, say security, improves privacy), negative (e.g., improving the other value, for example safety, leads to degradation in privacy), or neutral (e.g., the other value and privacy are decoupled from one another) from the perspective of the relevant stakeholder under consideration. The *interplay strength* describes the magnitude of positive or negative impact on privacy that a design consideration would have from the perspective of the relevant stakeholder. We suggest that a small set of discrete possibilities for interplay strength suffice for the analysis, such as {“weak”, “strong”}. Interplay type and interplay strength can also be combined into a single number, for example from the set $\{-2, -1, 0, 1, 2\}$, where “-2” represents a strong negative impact on privacy, “1” represents a weak positive impact on privacy, and so on.

A *design consideration* can support the rationale behind a particular design decision. For each interplay between privacy and another value, the relevant design con-

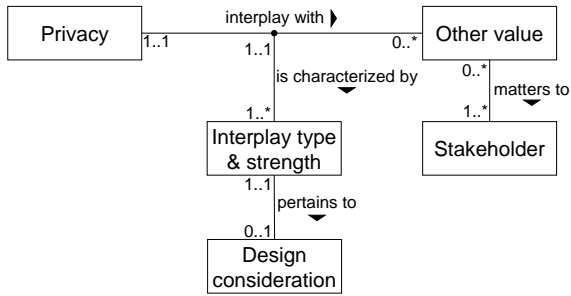


Figure 1: Overview of the used notions, with UML notation. Rectangles are types, edges are relationships between types. Relationships have a label to describe their meaning and a triangle showing the direction for interpreting the label. At the ends of an edge, cardinalities in the form min..max are shown. min is the minimum, max the maximum cardinality, and max = * means that the cardinality is unbounded. For instance, the edge between “Other value” and “Stakeholder” specifies that an “Other value” matters to at least one “Stakeholder”, while to each “Stakeholder” any number of “Other value”s may matter.

siderations are listed and assessed in detail. For a given stakeholder, when interplay types and strengths have been determined for all relevant design considerations, they can be aggregated into a single interplay type and strength between privacy and the other value. Thus, our methodology considers multiple design considerations for each (privacy, other value) pair and outputs a single representation of the interplay of that pair.

3.2. Considered values

In conducting a SPIDER analysis, it is important to consider a range of sociotechnical values (e.g., safety, efficiency, trust, and fairness). The set of values under consideration in any given analysis, and their definition, depend on the specificities of the technology under design, its intended application, and the social context within which the technology is deployed. Determining which values to consider, and how to interpret their meaning, generally requires some sociotechnical expertise, and should ideally entail broad stakeholder engagement, including both direct and indirect stakeholders drawn from a wide range of professional and personal backgrounds [24], [25]. To avoid bias, it is especially critical to include stakeholders from communities and populations in which the technology is targeted for deployment. Techniques to capture values include: (i) canonical moral values, e.g., from Haidt’s Moral Foundation Theory [15], [16]; (ii) empathy mapping, such as in [31], where the designers ask themselves how a group or individuals would feel, think, and interact with the technology through an empathetic lens; (iii) soliciting input from stakeholders, as exemplified in [24]; and using predefined lists of principles or values, such as those in the Principled Artificial Intelligence project [11]. SPIDER can be combined with any of these methodologies for determining the values to consider.

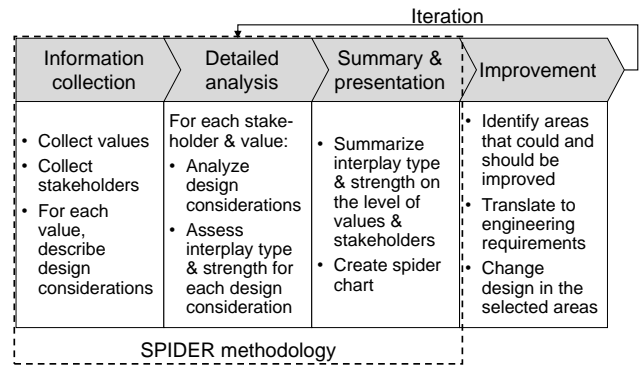


Figure 2: Overview of the SPIDER methodology within the continuous design improvement process

3.3. Process

The SPIDER methodology consists of three steps that can be embedded into a continuous design improvement (i.e., iterative design) process, as illustrated in Figure 2. We describe the steps as follows:

- 1) **Information collection.** In addition to privacy, other values relevant to the technology and application are identified and defined using one or more methods described in Section 3.2. The relevant direct and indirect stakeholders are identified and, for each value, it is determined to which stakeholders the value matters. Additionally, for every identified value, design considerations relevant to the interplay between privacy and that value are described.
- 2) **Detailed analysis.** For each such design consideration, the interplay between privacy and the other value is analyzed, documented, and assessed in terms of interplay type and strength. While doing so, the impact to privacy is regarded from the given stakeholder’s perspective. Each point of analysis is determined using the following general form: “Improving [other value] by [design consideration] increases/decreases privacy from the perspective of [stakeholder].”
- 3) **Summary & presentation.** Interplay type and strength are aggregated across design considerations for each stakeholder and value to create an overall interplay type and strength for that stakeholder and value. All value interplays for the stakeholder are visualized in a spider chart.

Results of the analysis can be used to identify opportunities and risks in the design space. This is especially apparent where strong interplays between privacy and other values are discovered. The analysis can guide design teams in their focus and in their design decision-making, since strong positive interactions indicate areas of opportunity for improving privacy and the other value, while strong negative interactions indicate potential risks to privacy. These insights can be used to derive more detailed and focused engineering requirements for improving the design while understanding the potential impact on multiple values including privacy. Embedded in an iterative design process, SPIDER can be repeated to evaluate the effect of the design changes (e.g., whether they resolved the previously identified tensions).




Figure 3: Automated Delivery Vehicle (ADV) use case

4. Case Study

In this section we illustrate the SPIDER methodology by applying it to a use case in the CCAM domain. This domain is particularly interesting because of its importance for the future of mobility, and also because it is a complex sociotechnical domain where privacy and many other values play an important role. The use case is described in Section 4.1, followed by the three steps of SPIDER in Sections 4.2–4.4.

4.1. Description of the use case

As an advanced CCAM example, we consider the use case of an automated delivery vehicle (ADV), illustrated in Figure 3. ADVs allow the cost-efficient last-mile delivery of goods to users. The ADV uses the public road infrastructure and has the ability to drive fully autonomously. Also, goods addressed to different customers/users may be loaded on the same ADV at the same time.

A typical example scenario is the following. A customer–Dani–orders medications from a local drugstore and selects robotic delivery. The ADV does not belong to the drugstore and may need to travel to the drugstore upon request. A route is calculated by the delivery service provider (DSP) for the ADV, minimizing the time to pick up and deliver all the packages. While the ADV is en route, the drugstore prepares the package, and hands over the package to the ADV on its arrival. The ADV could contain orders from multiple customers. Then, the ADV drives along its route to Dani.

During the drive, the ADV is using its array of sensors (cameras, lidar, radar, etc.) to monitor the road traffic, detect various objects and people, and make predictions about their behavior so it can navigate safely and efficiently. All the while, the ADV, and all of its data, are being monitored by a human operator and stored remotely in case an unusual situation needs to be mitigated (e.g., the ADV fails in some way, cannot navigate an unusual traffic situation like a roadblock, etc.). While en route, a pedestrian suddenly runs across the street (not at a crosswalk or intersection) causing the ADV to rapidly stop in the lane of travel. The event is recorded for later analysis by the DSP. The ADV arrives at Dani’s location and Dani is notified. Dani walks to the curb to meet the ADV, is authenticated, then takes their package from the ADV. The ADV requests that Dani pose for a picture with both the package and their face in view of the ADV’s camera, as evidence of successful delivery.

4.2. Step 1: Information Collection

The values identified for this case study and their definitions are shown in Table 1. As mentioned in Section 3.2,

TABLE 1: Values interacting with privacy in the use case

Value	Definition
Safety	The system or component can operate within the defined safety constraints (e.g., no harms to the CCAM user or any other road user).
Fairness	The CCAM system or component can be exposed to or equally used by anybody, independent of age, gender, etc.
Trustworthiness	Users of the CCAM system or component can rely on the system’s fulfilling its intended task(s), and that they and their data provided to the system or components are safe and secure.
Usability	Ease of interaction for the CCAM user.
Security	Prevention, detection, and/or mitigation of possible physical or logical attacks on the CCAM system or components.
Functionality	Ability of the system or component to execute one or more intended tasks for CCAM.
Efficiency	The realization of the component or system consumes less resources (e.g., compute power, energy etc.) than other realizations.
Collaboration	Ability of stakeholders to collaborate (e.g., share data) with each other.

ideally one would directly engage stakeholders to identify a list of relevant values. For this case study, we used an empathy mapping approach with the stakeholders taken from Kargl et al. [17].

From the numerous stakeholders of a CCAM system, we selected two for illustration purposes: end-user and bystander. The end-user is the customer (Dani) who purchased the goods and requested robotic delivery. The bystander is an indirect stakeholder who may or may not approve of CCAM technology. The DSP is not considered a stakeholder for this analysis given they are the entity conducting the SPIDER analysis.

4.3. Step 2: Detailed Analysis

To illustrate our methodology, we selected a subset of values presented in Table 1 – namely security, fairness, and safety – and examined several design considerations for these values. We quantified each consideration separately and averaged them (rounding away from zero, i.e., with ceiling in positive and flooring in negative direction) to get an aggregated rating per stakeholder and value.

4.3.1. Security. For the interplay between privacy and security, we assessed three design considerations: physical security of the ADV; authentication; and data security.

Physical security of the ADV requires some monitoring of the external environment in order to prevent theft or damage, which may negatively affect bystanders’ or other road users’ privacy. On the other hand, improved physical security of the ADV prevents others from seeing the goods in the ADV, thus improving privacy of the end-user. Therefore, for the end-user, improving physical security would increase privacy—a “+2” rating (strong positive)—whereas it would decrease privacy—a “−1” rating (weak negative)—for the bystander.

Both business partner and user authentication are required when the ADV arrives at the shop and customer location, respectively. Authentication may require biometric data collection (e.g., fingerprints or facial recognition),

TABLE 2: Ratings of the interplay between **security** and privacy for different stakeholders

Design consideration	Stakeholder	
	End-user	Bystander
Physical security	+2	-1
Authentication	-2	0
Data security	+2	+2
Aggregate	+1	+1

or simply a picture of the person retrieving the package. Improving security with authentication thus decreases privacy—a “-2” rating (strong negative)—for the end-user. For the bystander, authentication has no effect, resulting in a “0” rating (neutral).

In-vehicle data (e.g., package, customer, sensor, and location data), and data sent to the DSP, say for teleoperation or fleet management, need to be secured against tampering or eavesdropping. This also protects any personal data relating to the end-user or bystander. Thus, improving data security increases privacy—a “+2” rating (strong positive)—for the end-user and bystander.

Averaging the ratings, security would rate “+1” (weak positive) ($\lceil(2 - 2 + 2)/3\rceil$) for the end-user and also “+1” (weak positive) ($\lceil(-1 + 0 + 2)/3\rceil$) for the bystander. Table 2 shows the ratings. Overall, improving security is positive for the privacy of both the end-user and bystander.

4.3.2. Fairness. The interplay between fairness and privacy entails several nuanced ethical considerations. We examine a few design considerations to explore these interactions more closely. One design consideration relates to how the ADV routes through a city because that may determine what types of individuals are exposed to the risk of the ADV operating near them. In particular, given that many cities within the US have been structured based on historically racist policies [30], the ADV may unintentionally further exacerbate these opaque systemic biases when deploying at scale across a city or even across a country. Hence, DSPs may have to take additional actions to compensate for historically enacted systemic bias. To improve fairness across demographics, the ADV and DSP may need to collect aggregated demographic data (demo data for short). We note that this may be a better area for regulators to step in such that DSPs and other ADV and related Original Equipment Manufacturers (OEMs) do not need to measure routing fairness. Scoring the privacy impact of collecting aggregated demo data to improve fairness, as an example, could take the form of “-1” (weak negative) for both the end-user and bystander. This is because collecting aggregated demo data tends to carry the risk of re-identification of users in the dataset.

Another design consideration relates to real-time ADV sensing. Sensors on an ADV collect large amounts of data via various sensing modalities, which may have different limitations. It is well-documented that camera-only based object detection algorithms have difficulty discerning skin tone [35]. The DSP may need to collect detailed individual demo data in order to measure how discriminatory their object classification algorithms are against various demographics. Further, the ADV may need redundant sensing modalities to reduce the effect of the shortcomings of

TABLE 3: Ratings of the interplay between **fairness** and privacy for different stakeholders

Design consideration	Stakeholder	
	End-user	Bystander
Aggregated demo data collection	-1	-1
Individual demo data collection	-2	-2
ADV prediction bias	0	-2
Aggregate	-1	-2

individual sensing modalities. Thus collecting individual demo data to improve fairness has a “-2” (strong negative) effect on end-user and bystander privacy.

Additionally, the ADV may make behavior predictions about bystanders, and hence, behave differently as a function of demographics, such as age and gender. For example, age may help predict the behavior of bystanders such as to discern between the erratic motion of a child versus the more predictable motion of an adult. The end result is that the ADV is likely to behave more cautiously around bystanders it identifies as children compared to those identified as adults. This may be a socially acceptable bias in the ADV behavior, but it is important to call out that it is a biased behavior towards a particular demographic. If there is a need for the ADV to discern gender, then that could once again lead to certain demographics experiencing different ADV behavior than another. So improving fairness by introducing prediction biases decreases privacy for bystanders—a “-2” rating (strong negative)—while having a neutral “0” impact to the end-user’s privacy with the assumption they are at home and not behaving as a bystander in the scenario.

As shown in Table 3, fairness and privacy have a weak negative interplay for the end-user ($\lfloor(-1 - 2 + 0)/3\rfloor = -1$) and strong negative for the bystander ($\lfloor(-1 - 2 - 2)/3\rfloor = -2$). Therefore, these design considerations underscore a potential tension between privacy and fairness.

4.3.3. Safety. To assess the interplay between privacy and safety, we analyze four design considerations. First, sharing data between ADV OEMs for greater testing coverage improves safety across all CCAM (because they can all share scenarios (see SafetyPool¹)), but may come with privacy concerns due to sharing sensor data that may include personal data of end-users and/or bystanders. Thus, scoring the design consideration of sharing testing data to improve safety has a “-2” (strong negative) privacy impact for the end-user and bystander.

A second consideration relates to collision risk models, which are often based on age [23] or gender as proxies for physical stature. These collision risk models can be used by the ADV designer when validating the safety of the system. However, this requires personal identifiable information about individuals (e.g., age) involved in prior collision risk measurements. Therefore, this rates “-2” for the bystander. We score the end-user as “0” (neutral) because it is assumed the ADV would be stationary at the curbside during user interaction, making collision risk modeling irrelevant.

1. <https://www.safetypool.ai>

TABLE 4: Ratings of the interplay between **safety** and privacy for different stakeholders

Design consideration	Stakeholder	
	End-user	Bystander
Sharing testing data	-2	-2
Collision risk models	0	-2
Remote monitoring	-2	-2
Information about goods	-1	0
Aggregate	-2	-2

Remote monitoring of the ADV is another design consideration. For safety purposes, remote operators may get access to video feeds, and could potentially learn detailed information about individuals in the ADV’s operating environment, especially if they regularly monitor the same area. For the end-user, a score of “-2” is given because tele-operators could see them during the entire delivery interaction (e.g., to assist end-users). For the bystander, we also rate it as “-2” because of the potential surveillance.

Transported goods could be hazardous, in which case the DSP would have to display some package information on the ADV for bystander and first-responder safety. Displaying information of this sort could alert bystanders to the nature of an end-user’s purchases, resulting in a reduction of privacy of “-1” (weak negative). Bystanders’ privacy is not affected, hence is rated “0” (neutral).

As shown in Table 4, safety has a strong negative interplay with privacy for both stakeholders (end-user: $\frac{[-2 + 0 - 2 - 1]}{4} = -2$, bystander: $\frac{[(-2 - 2 - 2 + 0)]}{4} = -2$). Therefore, safety features must be carefully designed with privacy in mind, and designers should justify their design decisions.

4.4. Step 3: Summary & presentation

In this step, we combine the rating for each value in a spider chart, allowing one to quickly identify the interplay between privacy and other values. Figure 4 shows the interplay between privacy and other values for the end-user and the bystander. Note that we also rated the other values not discussed in Section 4.3 (usability, functionality, collaboration, efficiency, and trustworthiness), but those analyses are omitted for conciseness.

We observe that the end-user exhibits 25% positive, 25% neutral, and 50% negative interplays. Especially, the four negative interplays are: between privacy and functionality (e.g., camera used at package delivery), between privacy and efficiency (e.g., sharing user data to store’s suppliers for production forecasting), between privacy and usability (e.g., user interface must be minimal and hence privacy settings are not directly available), and between privacy and safety (see discussion in Section 4.3).

The bystander has 37.5% positive, 25% neutral, and 37.5% negative interplays. We note that the bystander’s privacy is less negatively impacted than the end-user’s, which makes sense as it is an indirect stakeholder.

We found that Safety and Functionality are the most challenging values w.r.t. privacy (both have negative interplays for the two stakeholders), and hence, where special attention should be paid. It is worth noting that no opposite interplays were observed (i.e., positive interplay for one

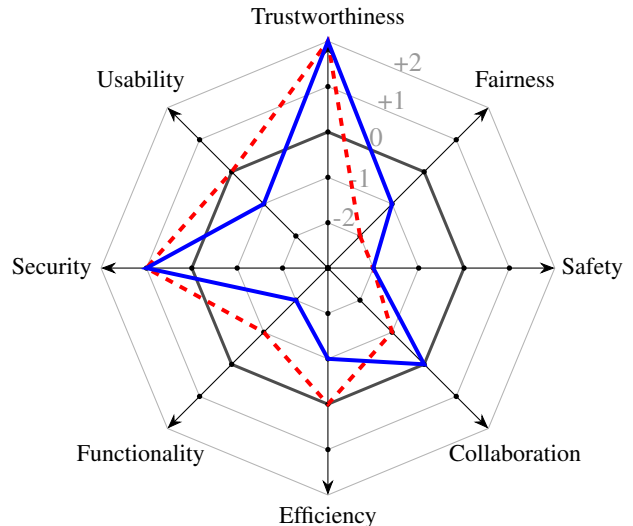


Figure 4: Interplay between privacy and each value for two stakeholders (blue solid: **end-user**; red dashed: **bystander**). Solid black line represents the neutral (0) rating.

stakeholder, and negative interplay for another). This is particularly promising as dealing with conflicting stakeholder needs would have been more challenging (because the designer would have to favor one stakeholder). Thanks to the spider chart, an ADV designer can identify that only 37.5% of the values have positive interplay with privacy, encouraging them to deploy privacy enhancing techniques specifically optimized for improving the other interplays.

5. Discussion

In this section we discuss the lessons learned from applying SPIDER, as well as possibilities for extending and further improving SPIDER in the future.

5.1. Lessons learned

Subjectivity of ratings. It should be noted that interplay ratings, to some extent, mirror the subjective assessment of the experts performing the analysis. This subjectivity cannot be completely removed in the context of values. Handling subjectivity during the process is especially important for questions that are not yet settled by law, standards, or best practices. However, SPIDER fosters objectivity by ensuring that the rationales of rating decisions are transparent (i.e., documented) and justified. If multiple experts perform the analysis, especially those with sociotechnical expertise, the differences in their perceptions become explicit; thus, SPIDER fosters constructive dialogue about values during the design process.

Interplay amongst stakeholders. Although SPIDER aims at mapping the interplay among values, as a side-effect it also gives insights into the relation between different stakeholders’ view on values. In particular, SPIDER makes it apparent if the interplay between privacy and another value is positive for one stakeholder and negative for another stakeholder. Making such conflicts between different stakeholders’ perception of values explicit could foster the mitigation of such conflicts.

Keeping the focus on rating the interplay. When performing the detailed analysis step of SPIDER, we noticed that it is easy to fall into the trap of rating the “other value” itself instead of its interplay with privacy. For example, when looking at fairness, one could be tempted to rate the impact of different design considerations on fairness for different stakeholders. During the analysis, it is important to be consistent in considering how an improvement to “other values” will impact privacy. Therefore, we envision that SPIDER analyses could be supported by providing a set of questions, specifically asking about the interplay between privacy and the given value, similar to the card deck of the privacy engineering approach LINDDUN GO².

Output of SPIDER. As mentioned earlier, the spider chart and the underlying considerations can be used to derive engineering requirements for making the system design more holistic in terms of designing for both privacy and the other values simultaneously. It is important to note that, in addition to engineering requirements, also non-technical requirements—e.g., for business processes related to running the service—may result from the analysis.

5.2. Possible extensions

Other rating aggregation functions. When aggregating the results from multiple design considerations, different aggregation methods can be used. Averaging, as done in this paper, is one possibility. A more conservative approach would be to take the minimum of the values: for example, if there is at least one design consideration leading to a rating of “-2”, then the result would be “-2”. This would lead to a much stricter overall assessment. Using a weighted average, where the assessor can determine individual weights for each design consideration, would yield more flexibility, but could also be overwhelming. More experience is needed to find out what aggregation method works best.

Deployment of Privacy Enhancing Technologies. Deploying PETs may influence the ratings. For example, using secure multi-party computation could allow private data sharing [6]. This could improve the interplay rating between privacy and collaboration. Or, blurring bystanders’ faces by using a video anonymizer [29] would improve the interplay rating between privacy and safety for remote monitoring (see Table 4). One could generate spider charts with and without considering PETs to observe their respective effects. But a systematic way to assess the effect of PETs on the ratings would help designers to determine the appropriate PETs to deploy in a given situation.

Multi-value interplay. In its current form, SPIDER is focused on the pairwise interplay between values (privacy and one other value at a time). With this, SPIDER can capture most of the interesting interactions among values. However, in some cases, the interplay of more than two values could be of interest. For example, the implementation of an improvement in functionality may also improve fairness, while increasing privacy, thus leading to a positive interplay among the triple of privacy, functionality, and fairness. Future work could examine how capturing

such multi-value interplay could be integrated into SPIDER without significantly increasing its complexity.

Broadening the definition of privacy. In our use case, we focused the SPIDER analysis on one direct (i.e., end-user) and one indirect (i.e., bystander) *human* stakeholder. Indeed, when thinking about privacy, it is natural to focus on people. However, complex sociotechnical systems involve a broader set of stakeholders, including legal entities such as DSP or city council. Investigating companies as stakeholders in SPIDER raises the debate to what extent the notion of privacy can be applied to legal persons. In the United States at least, companies have the right to claim a reasonable expectation of privacy in their records³, whereas the General Data Protection Regulation of the European Union is limited to data about natural persons. Since the techniques for protecting confidential business data are often similar to those for protecting personal data, it would be interesting to broaden the assumed definition of privacy in SPIDER to consider non-human stakeholders in order to capture the interplays of the sociotechnical system holistically.

Directionality of interplay. In the presented application of SPIDER, we rated the interplay using the following rule: “Improving [other value] by [design consideration] increases/decreases privacy”. This rule implies that the experts performing the SPIDER assessment only look at the impact of improving [other value]. However, it could be beneficial to also investigate “improving privacy by [design consideration] increases/decreases [other value]”. Analyzing both directions could confirm the directionality of the interplays, and potentially result in a more comprehensive understanding of the interplay between privacy and other values in the design.

Generalization to other core values. This paper has focused on privacy as the core value. In future, the methodology could be generalized to put other values in the center, or to look at all pairs of values.

6. Conclusions and Future Work

Building privacy-friendly systems is challenging, partly because of the interplay between privacy and other values, all of which hold some importance when designing complex sociotechnical systems. In this paper, we introduced SPIDER, a systematic methodology for assessing, quantifying, and visualizing the interplay between privacy and other values in system design. Contrary to the widespread belief that privacy would always conflict with other values, we have shown that the interplay between privacy and other values can be of different type and strength. Applying SPIDER to the realistic use case of an automated delivery vehicle has yielded first insights into the practical applicability of the methodology. In particular, we found that applying SPIDER fostered deep thinking about values and the design considerations for each value, thus yielding a comprehensive and balanced view on how privacy interacts with different values in the system design. The outputs of SPIDER, encompassing the detailed ratings of design considerations with their documented (i.e., transparent) rationale, as well as the

2. <https://linddun.org/go-getting-started/>

3. <https://www.dataprotectionreport.com/2019/02/companies-right-to-privacy/>

spider chart as visual overview, offer a promising basis for further improvement of the design.

The presented work opens several interesting paths for future research. First, SPIDER could be extended to also include the improvement and focused re-evaluation of system designs, possibly by incorporating SPIDER into an existing system design methodology. Second, it would be important to devise artifacts, such as checklists, templates, and digital tools, to support the execution of SPIDER. Finally, the application and evaluation of SPIDER in the context of real-world engineering projects would likely yield valuable insights about possible improvements.

Acknowledgements. The work presented in this paper grew out of discussions by the authors at 2023 Dagstuhl Seminar 23242 “Privacy Protection of Automated and Self-Driving Vehicles” [17]. The contributions and opinions expressed by each author are their own, and do not necessarily reflect the views of their employers.

References

- [1] ISO/IEC JTC 1/SC 27. ISO/IEC 29134:2023 - Guidelines for privacy impact assessment. 2023.
- [2] Carlisle Adams. *Introduction to Privacy Enhancing Technologies: A Classification-Based Approach to Understanding PETs*. Springer Nature, 2021.
- [3] Jonn Aksen and Benjamin K Sovacool. The roles of users in electric, shared and automated mobility transitions. *Transportation Research Part D: Transport and Environment*, 71:1–21, 2019.
- [4] Sean Brooks, Michael Garcia, Naomi Lefkovitz, Suzanne Lightman, and Ellen Nadeau. An introduction to privacy engineering and risk management in federal systems. Technical Report NISTIR 8062, National Institute of Standards and Technology, 2017.
- [5] Ann Cavoukian. Privacy by design: The 7 foundational principles. Information and privacy commissioner of Ontario, Canada, 2009.
- [6] Daphnee Chabal, Dolly Sapra, and Zoltán Ádám Mann. On achieving privacy-preserving state-of-the-art edge intelligence. 4th AAAI Workshop on Privacy-Preserving Artificial Intelligence, 2023.
- [7] Danielle Keats Citron and Daniel J Solove. Privacy harms. *Boston University Law Review*, 102:793, 2022.
- [8] Vasant Dhar. Equity, safety, and privacy in the autonomous vehicle era. *Computer*, 49(11):80–83, 2016.
- [9] David Eckhoff and Christoph Sommer. Marrying safety with privacy: A holistic solution for location privacy in vanets. In *IEEE Vehicular Networking Conference*, pages 1–8, 2016.
- [10] Golnaz Elahi and Eric Yu. Modeling and analysis of security tradeoffs – a goal oriented approach. *Data & Knowledge Engineering*, 68(7):579–598, 2009.
- [11] J. Fjeld, N. Achten, H. Hilligoss, A. Nagy, and M. Srikumar. Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI. Technical report, Berkman Klein Center for Internet & Society, 2020.
- [12] B. Friedman, P.H. Kahn, and A. Borning. Value sensitive design and information systems. In *The handbook of information and computer ethics*, pages 69–101. 2008.
- [13] Batya Friedman and Peter H Kahn. Human values, ethics, and design. In *The Human-Computer Interaction Handbook*, pages 1177–1201. 2003.
- [14] Dorothy J Glancy. Privacy in autonomous vehicles. *Santa Clara L. Rev.*, 52:1171, 2012.
- [15] Jesse Graham, Jonathan Haidt, and Brian A Nosek. Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, 96(5):1029–1046, 2009.
- [16] Jonathan Haidt. *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. Vintage, 2012.
- [17] Frank Kargl, Ioannis Krontiris, Jason Millar, André Weimerskirch, and Kevin Gomez. Privacy Protection of Automated and Self-Driving Vehicles (Dagstuhl Seminar 23242). *Dagstuhl Reports*, 13(6):22–54, 2024.
- [18] William J Kohler and Alex Colbert-Taylor. Current law and potential legal issues pertaining to automated, autonomous and connected vehicles. *Santa Clara Computer & High Tech. LJ*, 31:99, 2014.
- [19] Djibrilla Amadou Kountche, Marwan El-Bekri, Pedro J Fernández Ruiz, Jose Santa, Antonio F Gómez-Skarmeta, Joao Almeida, Qiang Tang, and Nuno Cruz. Cybersecurity and data privacy aspects in 5g-mobix. <https://hal.science/hal-03749024/>, 2022.
- [20] Stéphanie Lefevre, Jonathan Petit, Ruzena Bajcsy, Christian Laugier, and Frank Kargl. Impact of v2x privacy strategies on intersection collision avoidance systems. In *IEEE Vehicular Networking Conference*, pages 71–78, 2013.
- [21] Hazel Si Min Lim and Araz Taeihagh. Autonomous vehicles for smart and sustainable cities: An in-depth exploration of privacy and cybersecurity implications. *Energies*, 11(5):1062, 2018.
- [22] Zoltán Ádám Mann, Florian Kunz, Jan Laufer, Julian Bellendorf, Andreas Metzger, and Klaus Pohl. RADAR: Data protection in cloud-based computer systems at run time. *IEEE Access*, 9:70816–70842, 2021.
- [23] Harold J Mertz, Priya Prasad, and Annette L Irwin. Injury risk curves for children and adults in frontal and rear collisions. *SAE transactions*, pages 3563–3580, 1997.
- [24] J. Millar, D. Paz, S. M. Thornton, C. Parisi, and J. C. Gerdes. A framework for addressing ethical considerations in the engineering of automated vehicles (and other technologies). *Proceedings of the Design Society: DESIGN Conference*, 1:1485–1494, 2020.
- [25] Helen Nissenbaum. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2009.
- [26] Sakshyam Panda, Emmanouil Panaousis, George Loukas, and Konstantinos Kentrotis. Privacy impact assessment of cyber attacks on connected and autonomous vehicles. In *International Conference on Availability, Reliability and Security*, pages 1–9, 2023.
- [27] Jonathan Petit, Florian Schaub, Michael Feiri, and Frank Kargl. Pseudonym schemes in vehicular networks: A survey. *IEEE communications surveys & tutorials*, 17(1):228–255, 2014.
- [28] David E Pozen. Privacy-privacy tradeoffs. *The University of Chicago Law Review*, pages 221–247, 2016.
- [29] Zhongzheng Ren, Yong Jae Lee, and Michael S. Ryoo. Learning to anonymize faces for privacy preserving action detection. In *European Conference on Computer Vision*, 2018.
- [30] Richard Rothstein. *The color of law: A forgotten history of how our government segregated America*. Liveright Publishing, 2017.
- [31] Waralak Vongdoiwang Siricharoen. Using empathy mapping in design thinking process for personas discovering. In *Context-Aware Systems and Applications, and Nature of Computation and Communication*, pages 182–191, 2021.
- [32] Sarah Thornton. *Autonomous Vehicle Motion Planning with Ethical Considerations*. PhD thesis, Stanford University, 2018.
- [33] Tjerk Timan and Zoltan Mann. Data protection in the era of artificial intelligence: Trends, existing solutions and recommendations for privacy-preserving technologies. In *The Elements of Big Data Value*, pages 153–175. Springer, 2021.
- [34] Darrell M West. Moving forward: self-driving vehicles in china, europe, japan, korea, and the united states. *Center for Technology Innovation at Brookings: Washington, DC, USA*, 2016.
- [35] B. Wilson, J. Hoffman, and J. Morgenstern. Predictive inequity in object detection. *arXiv preprint arXiv:1902.11097*, 2019.
- [36] Qianhong Wu, Josep Domingo-Ferrer, and Ursula Gonzalez-Nicolas. Balanced trustworthiness, safety, and privacy in vehicle-to-vehicle communications. *IEEE Transactions on Vehicular Technology*, 59(2):559–573, 2009.
- [37] Eric Yu and Luiz Cysneiros. Designing for privacy and other competing requirements. In *2nd Symposium on Requirements Engineering for Information Security (SREIS)*, 2002.