

PriPoCoG: Empowering End-Users' Data Protection Decisions

Jens Leicht^a, Julien Lukasewycz^b and Maritta Heisel^c

Paluno - The Ruhr Institute for Software Technology, University of Duisburg-Essen, Germany

{jens.leicht, julien.lukasewycz, maritta.heisel}@uni-due.de

Keywords: General Data Protection Regulation, User Interfaces, Consent Management, Privacy Policy Customization, Policy Languages, Tool Support, Privacy Policy Visualization.

Abstract: The General Data Protection Regulation (GDPR) demands data controllers to provide transparent information about data processing to data subjects. This information is mostly provided in the form of textual privacy policies. These policies have many disadvantages, such as their inconsistent structure and terminology, their large scope, and their high complexity. For this reason, data subjects are likely to accept the agreement even if they do not fully agree with the data processing contained in it; this phenomenon is known as the privacy paradox. To overcome these disadvantages, we propose a user interface based on the results from a thorough literature review and a group interview. By not relying on a completely textual approach, we reduce the mental effort required from data subjects and increase transparency. We utilize the Prolog - Layered Privacy Language (P-LPL), which allows data subjects to customize privacy policies. Our work extends the compliance checks of P-LPL with compatibility checks for customized privacy policies. The proposed interface provides graphical representations for privacy policies, aligning with different mental models of data subjects. We provide a prototype to demonstrate the proposed theoretical concepts.


1 INTRODUCTION


Over the years, textual privacy policies became the de facto standard of informing data subjects about data processing. While the General Data Protection Regulation (GDPR) (European Parliament and Council of the European Union, 2016) demands transparency and an informed consent, textual policies often remain vague in their formulations and overwhelming for the data subject based on their large scope and complexity. This leads to data subjects having difficulties in fully comprehending privacy policies. Current privacy policies only offer two types of acceptance for the data subject: either fully accept or fully decline the policy. Further customizability of the content or a partial acceptance is not offered by data controllers. Both the length and complexity of the policy, as well as the absence of any customizability options, lead to the privacy paradox (Norberg et al., 2007). It states that data subjects generally have a high interest in privacy, while at the same time not being willing to invest the amount of work this implies. A high cost in time and effort will lead to data subjects renouncing


their rights by simply accepting the full policy. This is confirmed by a study by Ibdah et al. arguing that only 3.6% of data subjects are willing to read a monolithic textual privacy policy (Ibdah et al., 2021).

We present our privacy policy interface, leveraging the features of the PriPoCoG framework (Leicht et al., 2022). Our interface provides transparency through increased comprehensibility, while at the same time providing more details about the data handling. We also empower data subjects with customizable privacy policies, similar to the concept known from cookie banners. Instead of accepting or declining the full policy, data subjects are able to accept only those parts of a policy that they agree with. To achieve this customizability and ensure compatibility of the resulting policies with requirements of data controllers, we extend the P-LPL component of the PriPoCoG framework with compatibility checks.

We first show how our interface integrates into the PriPoCoG framework in Section 2. Next, we present our contribution in Section 3 by explaining how we planned and implemented our prototype user interface. We then compare our interface against related work in Section 4. Finally, we conclude our work and provide ideas for future research in Section 5.

^a  <https://orcid.org/0009-0003-5612-5590>

^b  <https://orcid.org/0000-0002-6850-4788>

^c  <https://orcid.org/0000-0002-3275-2819>

2 PriPoCoG FRAMEWORK

This work extends the Privacy Policy Compliance Guidance (PriPoCoG) framework (Leicht et al., 2022). Figure 1 shows an overview of the framework, with the privacy policy interface highlighted in orange. PriPoCoG covers the whole privacy policy life cycle, from policy creation, GDPR compliance checks, to policy management and enforcement; and now policy presentation to and customization by the data subject.

The framework supports *Policy Authors* regarding GDPR compliance by providing a *Privacy Policy Editor* (Leicht and Heisel, 2024). The editor creates privacy policies using the Prolog-Layered Privacy Language (*P-LPL*), which is the main component of the framework. It formalizes parts of the *GDPR* that are concerned with privacy policies. P-LPL also implements an extended version of the *Layered Privacy Language* by Gerl (Gerl, 2020). The policy editor enables policy authors to reuse work from *Threat Modelers* by importing information about data flows from the data flow diagrams (DFDs) created using the *DFD-Editor* (Leicht et al., 2023). This editor improves the privacy policy definition process and can also be used outside the framework as a stand-alone DFD-editor. Policy authors, as well as *Data Protection Authorities*, can get compliance feedback from the policy editor. The *Privacy Policy Management* component manages the large number of policies that a data controller receives from their data subjects (Leicht and Heisel, 2025). The framework uses Privacy Policy Based Access Control (*P2BAC*) to ensure that the privacy policies are actually enforced by the data controllers and data processors (Leicht and Heisel, 2023).

The current work on the *Privacy Policy Interface* connects the *Data Subject* to the framework. The interface accesses the *Privacy Policy Management*, to store customized policies, and the P-LPL component for compliance and compatibility checks.

3 CONTRIBUTION

First, we present our research methodology, followed by a literature review. To identify already evaluated approaches on privacy policy representations, we conduct a literature review (cf. Section 3.1). As these approaches only partially highlight which factors are the most significant ones for data subjects, we additionally perform a group interview (cf. Section 3.2). We then combine the findings of the literature review and group interview into general requirements for representing a privacy policy (cf. Section 3.3). These requirements help us to develop a data subject empowering user interface for privacy policies. We introduce the most important concepts of our interface in Section 3.4. Afterward, we present the prototype that we implemented (cf. Section 3.5).

3.1 Literature Review

In order to identify related work, requirements, as well as design elements for our privacy policy interface, we perform a rapid literature review. For this review, we defined two search strategies:

- A. “privacy polic* interface” OR “privacy polic* representation” OR “privacy polic* visualization”
- B. “custom* privacy polic*”

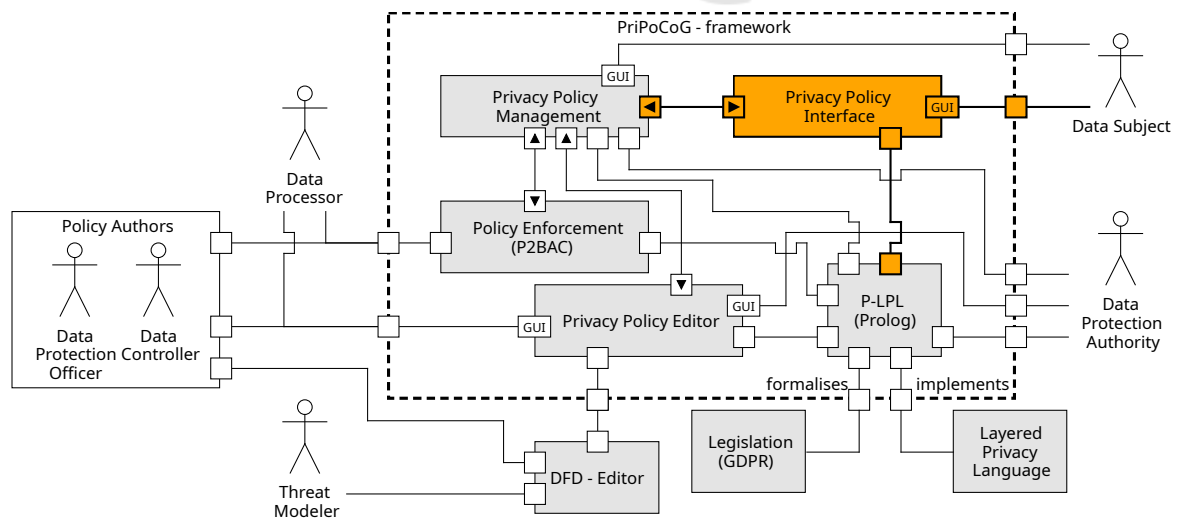


Figure 1: PriPoCoG-framework with our policy interface and its connections; based on (Leicht and Heisel, 2025).

We applied these strategies to the following five research databases: ResearchGate, IEEE Xplore, Springer Link, ACM Digital Library, and Scopus, as well as the search engines Google Scholar, Google, and the library search engine of the University of Duisburg-Essen¹. For the search engines, we slightly adapt the search strategies, in order to receive the most relevant results. The library search engine, for example, does not support placeholders; hence, we replaced asterisks with or-statements containing all different spellings (e.g., policy/policies).

All databases together returned 641 results, and after reviewing these results by title, we selected 165 papers for further review. Before continuing with the review based on abstracts, we removed duplicate results (e.g., returned by different databases/search engines), which resulted in 129 unique papers for further consideration. After looking into each of the 129 papers, we selected 28 papers as relevant for the requirement selection (cf. Section 3.3) and feature generation (cf. Section 3.4). Within these 28 papers, we also identified 19 references to interesting papers. From these secondary sources, we further selected five relevant papers for our work.

An adaptation of requirements from the constructivism learning theory to privacy policy representation inspired two of our requirements and provided a good basis for the development of our concepts (Papaioannou et al., 2022). A study concerning the interaction of data subjects with different cognitive styles with different policy interfaces helped us in the development of our privacy matrix (Tsolakidou et al., 2024). The study revealed that a nutrition label approach increased the time spent in interacting with the policy. Another study analyzed different kinds of policy representation based on the mental model of the data subjects (Paudel et al., 2023) and influenced our identified requirements. Further papers from our literature review are referenced throughout the paper.

3.2 Group Interview

The literature review showed, that the field of optimizing textual privacy policies is very broad. To focus our work on those concepts offering the most clarification and customization for a data subject, we conducted a group interview. In this interview, 15 participants with varying knowledge in the field of privacy and regular usage of internet services discussed the topic in a time frame of 30 minutes. A longer discussion was not considered, as the contributions started repeating. To guide the interview, we provided the participants the following six questions:

- Q1.** What traits should a privacy policy possess?
- Q2.** Which information do you consider important for yourself?
- Q3.** How would you optimize current textual policies?
- Q4.** Which other forms of representation would you like to have?
- Q5.** How much and what kind of assistance would you like to have in understanding privacy policies?
- Q6.** Would you consider a guided approach beneficial for a privacy policy?

The answers given by the participants yielded five main areas of importance for data subjects: clarification, simplification, personalization, customizability, and assistance.

Clarification and Simplification mean that data subjects would like to easily and fully understand all information included in the privacy policy, and the resulting implications of their consent, in a short amount of time. They also stated that alternative forms compared to a full text policy, like icons, summaries of the policy, or visual representations, would convey the contained information better.

Personalization could take the form of adapting the policy representation to the type and expertise of a data subject, or introducing preferences to reduce time consumption.

Customizability, in contrast to personalization changing the representation of the policy, customizability refers to actually changing the contents of the policy by accepting or declining individual parts.

Assistance could guide the data subject through the policy by sensibly splitting up the policy or helping to understand the content.

A large percentage of the participants also clarified that they take their own privacy rights very seriously. However, they also stated that they often do not read a full textual policy, but rather give up their rights; having a limited interest in understanding the legalese, limited time to invest in reading policies, or lacking comprehension of the policy. This shows that the participant group confirms the privacy paradox without being experts in the field of privacy or knowing this concept (Norberg et al., 2007).

¹<https://primo.uni-due.de/>

3.3 Requirements

From the literature review and the group interview, we derived the following requirements for privacy policy representations:

- R01.** The privacy policy should align with the concept of transparency stipulated by the GDPR (European Parliament and Council of the European Union, 2016, 5.1.(a)). Hence, all processing actions performed on the personal data should be included in the policy.
Providing detailed information about the data handling can positively affect data subjects' trust in the controller. However, it can also have a negative impact on this trust, as data subjects may become suspicious when they learn about all the data processing that is taking place. (Fischer-Hübner and Karegar, 2024)
- R02.** The policy representation should convey the contents of the privacy policy to the data subject comprehensible, simple, and with minimal extent. Specialized forms of representation should be chosen to best fit the kind of information that should be conveyed. A policy should further avoid any form of legalese.
This requirement is partially in conflict with R01, as a high level of detail increases the complexity of the policies. We base this requirement on the interview (cf. Section 3.2) and different publications (Paudel et al., 2023; Earp et al., 2007; Papaioannou et al., 2022; European Parliament and Council of the European Union, 2016).
- R03.** The content and structure of a privacy policy should align with the mental model of the data subject (Paudel et al., 2023). If both views do not align, the representation of the privacy policy should be personalizable to a more fitting representation for the data subject.
- R04.** The mental effort of the data subject invested in understanding the content of the privacy policy should be at an acceptable level. This investment depends on the amount of information contained in the policy and the time necessary to comprehend the policy. (Tsolakidou et al., 2024)
- R05.** The form of representation of the privacy policy should fit the learning process of the data subject (Papaioannou et al., 2022).
- R06.** A privacy policy should enable the data subject to only accept those parts aligning with their privacy preferences. Configurable consent for separate parts of a privacy policy increases the level

of control perceived by the data subject (Fox et al., 2022).

We base this requirement on the interview (cf. Section 3.2) and (Fox et al., 2022).

- R07.** The consent of a privacy policy should only follow the opt-in principle (European Parliament and Council of the European Union, 2016, 4.11). Explicit consent for each purpose increases the perceived control of the data subject (Fox et al., 2022).
- R08.** The data subject should receive additional assistance and guidance during the understanding and consent phases, if needed.
We base this requirement on the interview (cf. Section 3.2).
- R09.** After accepting a privacy policy, data subjects should receive a receipt containing the contents of the policy and the state of consent (Jesus, 2020).
- R10.** The kind of representation and consent chosen for a privacy policy should be acceptable and usable by the industry (Sailaja and Jones, 2017).

3.4 Concepts

In the following, we present different concepts that address the requirements we identified. Screenshots are taken from our prototype, which we present in more detail in Section 3.5. The screenshots show an exemplary privacy policy, which represents an online shopping scenario.

Icons (R02, R03, R08): We use icons for multiple purposes: 1) icons provide an overview of which data are collected and used for which purpose; 2) icons provide important information about properties of processing purposes, e.g., whether the purpose must be accepted (at least partially) for the service to be usable; 3) country flags visualize where the data will be processed; 4) finally, icons provide feedback about the state of the policy (e.g., whether the accepted parts of the policy suffice for service provision).

Icons convey the contents of the policy in a simple manner and with minimal extent (R02). Data subjects are not required to read lengthy texts to get an overview of the processing of their data (R03). To guide data subjects in understanding the privacy policy icons and, hence, the policy itself, we use mouseovers and a glossary with an overview of all icons and their explanations (R08).

The set of icons we use is based on the Data Protection Icon Set (DaPIS), an approach for standardized icons for the GDPR (Rossi and Palmirani, 2019).

We extended the icon set with 68 additional icons. The new icons are based on existing icons from the set as well as material icons², which are regularly used in UI- and web-design. Icons like (👤) and (🔑), from the original DaPIS set, represent different elements of a privacy policy (i.e., the data controller and the list of purposes). We added icons describing the kind of data that may be processed for a specific purpose, e.g., contact details (📞) or biometric information (🔍). Icons can also give feedback concerning the status of the policy, e.g., if all purposes, required for service provision, have been selected. We provide open access to the extended icon set via GitHub³.

Tsolakidou et al. examined how well data subjects with different cognitive models can understand privacy policies with different forms of representation (Tsolakidou et al., 2024). They come to the conclusion that policies should align with the cognitive model of the data subject. This is why we provide different forms of representation (cf. Matrix and Map below). Icons can provide a quick way for some data subjects to understand a privacy policy, while they may hinder others. A study by Windl et al. shows that pure graphical representations of privacy policies using icons do not suffice to convey the contents of a privacy policy; icons can be misunderstood (Windl et al., 2022). To overcome this issue of misunderstood icons, we provide mouse-overs and a glossary. Additionally, we provide a structured textual representation of the policy. Providing both, graphical and textual representation, is supported by the studies around the GDPR privacy label (Fox et al., 2022).

Country Flags (R02). We use country flags to quickly indicate in which countries data will be processed. Since not all data subjects may be fluent in vexillology (study of flags), we again make use of mouse-over information about the country indicated by a flag. The use of flags is a specialized form of representation that conveys third country data transfers in a simple manner, hence, addressing R02.

Structure (R02 + R04). To reduce the mental effort of the data subject (R04), we provide a well-structured representation of the policy. Important information is presented in single words or small sentences, instead of a wall of text, which increases comprehensibility (R02). Further details are available when the corresponding view is expanded. Figure 2 shows the expanded purpose overview on the right half of the interface. In the non-expanded state, only

the purpose categories (e.g., *Legal Compliance*) are shown. When selecting a specific purpose, another, more detailed view is opened (cf. Figure 4). By separating essential information from additional information, we expect that data subjects will be more motivated to try to understand the privacy policy. Our approach of simultaneously showing as little text as possible is supported by a study (Ibdah et al., 2021). This study shows that it is beneficial to hide details behind links or buttons and instead show compact overviews at each section of the policy.

UI-Overview: Figure 2 shows the main view of our policy interface. It is split into two columns and top and bottom action bars.

Top Bar. At the top is space for the logo of the service to which the privacy policy applies, followed by the name and version of the privacy policy. On the right-hand side are buttons for the glossary and the help system. Finally, there's the language selection, depending on the different translations of the privacy policy provided by the data controller.

Left Column. The left column consists of four elements: 1) a description of the contents of the policy, e.g., if it only applies to parts of the service, or a description of the service it applies to, 2) general information about the data controller, data protection officer, data subject rights, and supervisory authority (collapsed in Figure 2), 3) the processing overview, which gives an overview of the purposes and data categories using the processing matrix, which we explain in more detail below, and 4) a world map, which highlights the countries, to which data will be transferred for processing; we explain this map in more detail below.

Right Column. An overview of processing purposes, sorted by purpose category, is presented in the right column. When a specific purpose is selected, the purpose overview is replaced with the purpose details, shown in Figure 4.

Bottom Bar. The bottom bar first indicates whether the policy can be saved in its current state, i.e., whether all required purposes have been selected for consent (❌/✅). Next, there are the buttons *Select None*, *Select Required*, and *Select All*, providing a quick way of configuring the policy. We, however, want to motivate data subjects to customize policies according to their preferences instead of using these buttons. Once all required purposes are accepted, the *Next*-button will be enabled, allowing the data subject to proceed to the summary page, which we explain in more detail below. On the right-hand side, we have a link to the *Full Text Policy*.

²<https://fonts.google.com/icons>

³<https://github.com/jensLeicht/DaPIS>

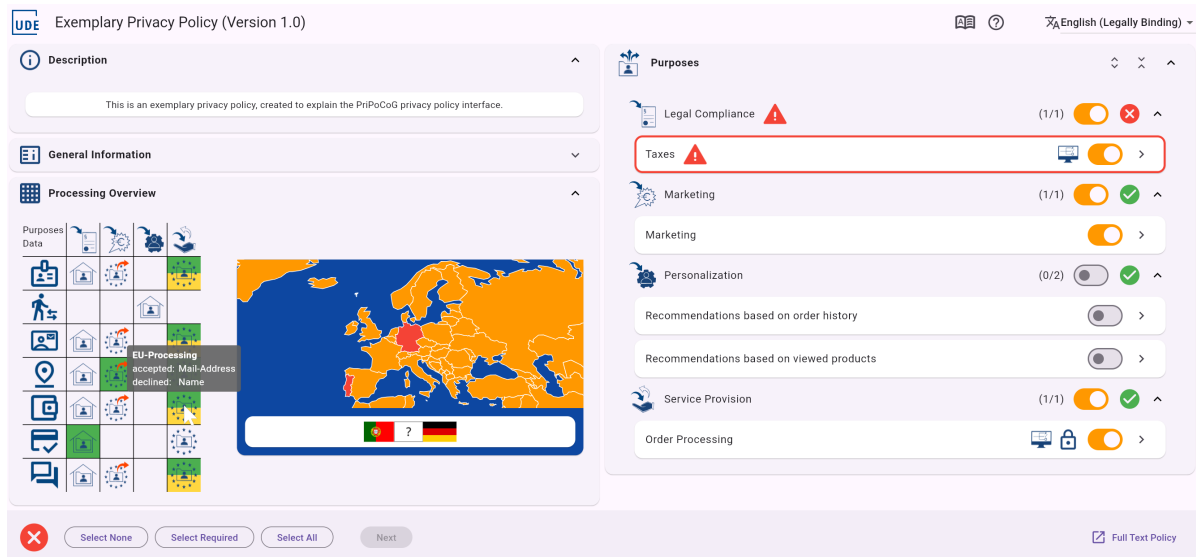


Figure 2: Screenshot of the main view of our privacy policy interface.

Processing Matrix (R02, R04, R05). To provide a concise overview of the privacy policy, we provide a processing matrix, which is called *Processing Overview* towards the data subject. The matrix is part of the left column of the main view (cf. Figure 2). Columns of the matrix represent the purpose categories, in which the purposes of the policy are classified. In our example, these categories are *Legal Compliance* (📄), *Marketing* (📢), *Personalization* (👤), and *Service Provision* (🛒). The rows of the matrix represent data categories, showing which data is used by which purpose category. Figure 2 shows the following data categories: *Identifying* (👤), *Behavioral* (👤), *Contact* (📞), *Location* (📍), *Account* (🔑), *Credit* (💳), and *Communication* (💬). These categories are part of P-LPL and the PriPoCoG framework.

Cells of the matrix can contain three different icons, or they can be empty. The icon for the internal processing of data (🏠) shows that the data will only be processed by the data controller. The icon for processing inside the European Union (🇪🇺) shows that data is processed by data processors, who are located inside the European Union. The last icon marks processing outside the European Union (🌐). If a cell is empty, no purpose in this purpose category is processing data from this data category.

The icons subsume each other in the following order: 🌐 > 🇪🇺 > 🏠. This means that when processing outside the European Union is indicated, processing can also take place inside the EU, as well as internally. If processing inside the EU is indicated, there will be no processing outside the EU, but internal processing may still be included. The cells always show

the icon that represents the furthest data transfer for a combination of data category and purpose category.

The color of the cells corresponds to the status of the purposes. Uncolored cells are completely disabled, meaning consent will not be provided. When a purpose is completely selected for consent, the cell is highlighted in green (🟢). When a purpose is customized for partial consent, the cell is colored green and yellow (🟡). While the data subject customizes the policy, the cell color changes dynamically, supporting the experiential learning process of the data subjects (R05).

Having a feature that provides an overview of the privacy policy reduces mental effort, as the data subjects do not need to comprehend a long text (R04). This overview also provides information to the data subject in a simplified manner and with minimal extent (R02).

Graphical representations of privacy policies have already been proposed and evaluated in the past. The nutrition label approach has been implemented and positively evaluated in multiple iterations (Kelley et al., 2010). Other forms of overview and labels have also been evaluated positively (Fox et al., 2022, 3.4.1). These graphical representations inspired our processing matrix.

A study concerning privacy policy representation showed that data subjects want “all on one screen” and “see the bigger picture” (Lipford et al., 2010). By combining our visual representation in form of the matrix with the interactivity of customizing the policy, we follow the findings of another study (Reinhardt et al., 2021).




Figure 3: Levels of detail.

Map (R01, R02, R03, R04, R06). To provide immediate transparency about the destinations of data transfers (R01), we implemented the world map, which is part of the left column of the main view (cf. Figure 2). This specialized form of representation (R02) shall align with the mental model of visual learners (R03), and we expect it to reduce the mental effort required to obtain an overview of data transfers (R04). The map is updated dynamically when policy elements get enabled or disabled, giving immediate feedback during customization of the policy (R06).

Levels of Detail (R02 + R04). Closely related to the structured representation, we provide different levels of detail in our interface. Not every little detail of a privacy policy is necessary to make an informed decision; especially with P-LPL, where policy authors are encouraged to provide every little detail of their data handling practices. Studies show that users want long and detailed privacy policies, because they think long policies provide better privacy (Proctor et al., 2008). However, the same users do not want to read long policies. To tackle this privacy paradox, we first provide an overview with as little additional detail as possible (R02). If a user is then interested in finding out more about the details, they can navigate a level deeper.

Figure 3 shows a pyramid, representing different levels of detail we use in our interface. The size of each layer in the pyramid correlates with the amount of information contained in said level. When a data subject is overwhelmed by one of the representations, they can use one of the levels of detail, reducing their mental effort to a minimum (R04). We define the levels as follows:

Icons. The most basic level of information is provided by the iconography (cf. *Icons* above), used in different locations across the interface, e.g., the processing matrix (cf. Figure 2).

Mouse-Over (Tooltips). When the data subject has difficulties in understanding the icons presented to them, they can retrieve additional information by hovering the mouse over the icon. Tooltips describe the meaning of an icon in a single word or provide additional information about the element at hand, e.g., which specific data is processed in a category (cf. Figure 2).

Details Page. When a data subject is interested in learning more about a specific purpose, for which their data shall be processed (e.g., *Order Processing* in Figure 2), they can access the details page of a purpose. Figure 4 shows the details page for the purpose *Order Processing*. The top bar shows the icon for automated decision-making (🤖), as well as the lock (🔒), meaning that this purpose is at least partially required for service provision. The page contains the following information:

Description (i). A general description of the purpose.

Retention (🕒). For how long the data will be stored and processed. The icon gives a quick overview of the type of retention: *indefinitely, at point X in time, X amount of time after the purpose concluded* (shown in Figure 4).

Data (👤). The list of data elements that will be collected and processed for this purpose. Icons behind the data name show the data categories, to which this data is linked. The *name*, for example, is linked to the categories *Account, Contact, Communication, and Identifying*. The colored bars at the end show the data severity/riskiness, and the lock shows if some data is required for service provision. These elements can be expanded to see more details about the data collected (cf. *Address* in Figure 4).

Data Recipients (🚩). The list of data recipients (in this example, *Parcel Service*). The flag shows in which country the recipient operates, and the icon behind the flag visualizes the type of data recipient: *Person, Legal Entity* (shown in Figure 4), and *Public Authority*. The lock at the end of the element indicates a requirement for service provision. Data recipients can be expanded, to view more information about the recipient.

Legal Bases (⚖️). The legal bases on which the processing is based.

Automated Decision Making (🤖). More information regarding automated decisions concerning this purpose (in this example, *Automatic Username Generation*).

Subordinate Purposes (📁). The purposes in which the current purpose can be subdivided. Each sub-purpose can be accessed in more detail by clicking on it.

More Information (⋯). Further technical details about the data processing. For example, information regarding privacy models (e.g., k-anonymity) or pseudonymization methods will be applied to the data.

Textual Privacy Policy. For the legal experts and privacy enthusiasts, we provide a link to the underlying textual privacy policy. This link can be found on the bottom right side of the interface (cf. Figure 2). PriPoCoG generates this document from the P-LPL policy, describing every detail in a coherent document.

With these levels, we address the conflict between requirements R01 and R02. Transparency is achieved by providing detailed information about each processing purpose. Whilst the initial view of the policy, in our interface, is limited to the minimal extent.



Figure 4: Screenshot of the purpose details page for the purpose *Order Processing*.

Prolog - Layered Privacy Language (P-LPL) (R01 + R02). To be able to have a well-structured policy representation (R02), as well as different views of the same policy, we make use of an existing privacy policy language: the Prolog - Layered Privacy Language (P-LPL). By providing semantics for the privacy policy defined in P-LPL, the language enables us to extract all necessary information for the different concepts presented in this section. P-LPL is part of a larger framework, called the Privacy Policy Compliance Guidance (PriPoCoG) framework (Leicht et al., 2022). In Section 2, we have shown how our interface integrates into PriPoCoG.

Using P-LPL further enhances the comprehensibility and simplicity of the policy representation, by providing a guided policy editor to policy authors. The editor reduces the use of unnecessary filler text, and recommends the use of simple language, thus reducing the amount of legalese in P-LPL policies (R02).

Current text-based policies often lack many details that might be of interest to data subjects, such as what specific data is being transferred. The PriPoCoG framework encourages data controllers to provide every detail about every data processing. Additionally, it generates parts of the privacy policies from data flow diagrams of the data controller's systems (Leicht and Heisel, 2024; Leicht et al., 2023). While this may increase the overall size and complexity of a privacy policy, using our interface counteracts this complexity. The levels of detail and use of graphical representations of information make the detailed policies more manageable. Putting all this information into the policy enhances the transparency towards data subjects (R01).





Customizability (R04, R05, R06, R07, R08). According to a study about privacy choices, 74% of participants say that they do not have control over their data when they use the internet (Zimmeck et al., 2024). Although legislation forces data controllers to collect consent, data subjects do not feel in control, as many privacy policies follow the take-it-or-leave-it principle.

Over 80% of participants of another study favored explicit consent (e.g., via a check-box) over an *accept-by-sign-up* (Ibdah et al., 2021). With our interface we go a step further and make privacy policies customizable in many details (R06). Data subjects can reject dedicated purposes, data recipients, or data. Additionally, we provide buttons for accepting all purposes completely, accepting only the purposes, data recipients, and data required for service provision, or declining the complete policy.

Due to the default effect, which describes that users tend to click “next” without changing pre-selected options, we expect data subjects to not customize their privacy policies, when everything is enabled by default. To make privacy policies more privacy preserving and encourage data subjects to actively think about the data they share, everything is disabled by default (R07). This is why even required elements have a toggle-switch that is off by default (cf. subordinate purpose *Forums* in Figure 4). This strict opt-in approach may also encourage data subjects to think about positive consequences of sharing data (Knijnenburg et al., 2013). Making data subjects actively consent to processing purposes, by explicitly enabling them, also increases their informedness when providing consent. Starting with an empty policy also reduces the mental effort of the data subjects, as they only need to read and understand the parts of the policy that they want to enable (R04).


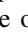
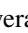

By providing visual feedback for policy changes, e.g., the colored cells in the processing matrix, we support experiential learning (Papaioannou et al., 2022) (R05 + R08). A study concerning visual interactive privacy policies showed, that a combination of visual representation and interactivity leads to data subjects investing more time in understanding the policy (Reinhardt et al., 2021). This improves the informed consent, as policies are not just accepted without comprehension. The same study also identified that interactiveness improves the attractiveness of the privacy policy interface compared to just a visual representation like the nutrition label approach (Kelley et al., 2010).

Glossary and Help (R08). To support data subjects in understanding our policy interface and the privacy policy presented, we provide a glossary for the icons and terminology used. A help menu supports data subjects in getting started with the interface.

Data Severity/Riskiness (R02, R04, R06, R08). To further support data subjects in learning about the risks of data sharing, we provide data severity indicators (from highest to lowest risk):  >  >  >  - providing these indicators shall reduce the mental effort required from data subjects to make an informed decision (R04). Additionally, riskiness indicators help data subjects to customize the policy according to their preferences (R06), and provide guidance for inexperienced data subjects (R08).

Compliance Check (R04 + R10). Before a policy is presented to the data subject, our interface uses the P-LPL backend of the PriPoCoG framework

to check the policy for GDPR compliance. This ensures that the customized policy does not contain any non-compliant data processing. It also ensures that data controllers have GDPR-compliant privacy policies, which is beneficial for industry (R10). When a compliance issue is detected, the policy will not be opened. The compliance check also reduces the mental effort invested by data subjects, as they do not need to worry about potentially non-compliant data processing (R04).

Compatibility Check (R04, R05, R10). Since we use the opt-in principle and some processing may be required for service provision, we need to assure that data subjects read, understand, and consent to these required elements of the policy. To this end, we mark required elements with a lock symbol (). To reduce the mental effort of checking whether all required elements have been selected for consent (R04), we provide automatic compatibility checks, using the P-LPL backend. We call this *compatibility check*, as it checks whether the customized policy is compatible with the policy of the data controller (R10). For this check, we extended P-LPLs functionality. It not only checks for required purposes, but also checks for other logic errors inside the customized policies. It is, for example, not possible to provide consent for an optional purpose, without allowing the processing of at least one data element in this purpose. The results of the compatibility check are visualized in multiple places throughout the interface. For each purpose category in the right column of the main view (cf. Figure 2), icons highlight the compatibility state of this purpose category (/). The overall compatibility state of the policy is visualized by the same icons in the bottom left corner of the interface. Supporting data subjects in identifying missing required elements, we highlight issues using (). These icons provide further details about the compatibility issues via mouse-overs. By providing this feedback, we support data subjects in learning which data is regularly required by which kind of data controller (e.g., addresses in online shops; R05).

Summary (R02). Before the data subject can submit their consent, we show a summary of all purposes that they selected in the main view of the policy interface. This summary is reached when clicking the *Next* button in the main view (cf. Figure 2), after customizing the privacy policy. The summary is a reduced form of the privacy policy, where all disabled elements have been removed (R02). We added this summary to assure informed consent, when data subjects use the *Select Required*, or *Select All* buttons.

Receipt (R02 + R09). After submitting the consent to the data controller, data subjects will receive a receipt, summarizing the agreed upon privacy policy (R09). The receipt provides a quick overview using the processing matrix (R02). Additionally, it contains the full-text version of the final policy, containing only the agreed upon purposes. Data subjects may use this receipt to prove what processing they consented to, to the data protection authorities (Jesus, 2020). This proof might help data protection authorities investigate potential compliance issues of data controllers.

3.5 Prototype

For our prototype, we only considered policies on websites. Our approach is nevertheless also applicable to other forms of digital consent collection, such as privacy policies in software applications. The prototype needs to be adapted for the different types of underlying software, but the general procedure and architecture can remain identical.

Our prototype consists of two main software systems: the frontend, displaying the user interface to the data subject, and the backend, necessary to execute the compliance and compatibility checks using P-LPL. Both pieces of software are available on GitHub⁴, including a demo. We chose Flutter⁵ as our development framework as it possesses multiplatform functionality and is well capable of fulfilling our requirements. Any other web frameworks or languages should, in general, be equally suitable to implement our proposed user interface. The backend is provided by the PriPoCoG framework and is provided by a Python web server, bridging over to Prolog.

We decided to not run both the frontend and Prolog system locally, on the device of the data subject, as it is not guaranteed that all necessary conditions to run Prolog are given. This approach also reduces the need for data subjects to install new software on their devices. The policy as well as the functionality of the compliance check are provided by the data controller. This may increase mistrust against the data controller, because they might manipulate the compliance check. Thus, data controllers might theoretically be capable of sending wrong compliance results to the data subject. The data subject would have no validation mechanism, meaning that the result of the compliance check is valid and was not manipulated by the data controller or other entities. To counteract this risk, we propose the use of mechanisms like homomorphic encryption and trusted computing platforms to validate the correct execution of the compli-

ance check. Alternatively, a trusted third party could be used, which executes the compliance check in a neutral environment. Such a trusted party could be provided by central public organizations, e.g., by the European Union. The discussion about the correct way of handling compliance and compatibility checks would exceed the scope of this work; therefore, we assume that the data controller will execute these checks truthfully.

In the following, we explain the general process of a data subject interacting with our interface. First, the data subject requests the privacy policy interface for reading and customizing the policy. Before presenting the policy to the data subject, it will be checked for GDPR compliance. This ensures that a data subject is not able to accept a policy that is not GDPR-compliant and forces the data controller to only create compliant policies. In case the policy is not compliant, the process of acceptance will be terminated, and the data subject will be informed. Otherwise, the data subject can read and customize the policy to their preferences by accepting and declining purposes, data recipients, and data (cf. Section 3.4). This process is iterative, and after each change, the compatibility with the initial policy and the information required for service provision is checked. If required elements are disabled and the data controller is not able to provide the service, because the policy disallows it, the policy is not compatible. In this case, the user is informed about the compatibility issues in order to adapt their decisions accordingly or decline the overall policy. If the policy is compatible with the requirements of the data controller, the data subject may submit their consent. The policy is then sent to the backend for storage inside the policy management. Additionally, the data subject receives a receipt of the customized policy.

4 RELATED WORK

Drozd and Kirrane proposed multiple iterations of privacy policy interfaces. In the CURE interface, they provide two customizability options: 1) a slider, providing presets with different levels of privacy, and 2) a list of purposes, which can be customized via check-boxes (Drozd and Kirrane, 2020). Compared to our interface, CURE equally provides customizability, thus empowering data subjects. Their policy representation, limited to the list of purposes with only little detail, is not as comprehensive as ours. We provide diverse representations to support the different mental models of data subjects. Our interface could be extended with a preset system similar to CURE's slider.

⁴<https://github.com/jensLeicht>

⁵<https://flutter.dev/>

Kaili and Kapitsaki also proposed multiple iterations of their privacy policy beautifier. The beautifier is a tool that processes existing textual privacy policies and provides different forms of highlighting and representation of the processed policy (Kaili and Kapitsaki, 2023). While these representations support data subjects in processing a textual privacy policy, as evaluated by Kaili and Kapitsaki, we think that an overall more structured approach may be even more beneficial (cf. Section 5).

European research projects, like Trapeze⁶, also developed different kinds of policy representation. Trapeze developed different privacy dashboards: a data-centric and a consent-centric dashboard (Raschke and Eichinger, 2022). While these dashboards provide many customization options and can also convey many details about a privacy policy, a quick and simple overview is not provided. Our processing matrix can provide this quick overview. Some concepts used in the Trapeze dashboards may be adapted and added to our policy interface in future versions.

5 CONCLUSION & FUTURE WORK

The de facto standard of current privacy policies is the take-it-or-leave-it approach. This could cause data subjects to decline a policy when disagreeing with only some purposes or data recipients. Data subjects might also feel forced to accept a policy, as it might be difficult finding an alternative service. This two-option approach either harms data subjects' privacy and their rights, or the data controllers and their businesses. Both stakeholders could benefit from a more advanced type of consent collection, as we have presented in this paper. According to the reverse privacy paradox, some data subjects protect their privacy more than they say they do (Colnago et al., 2023). By empowering these data subjects with customizable privacy policies, data controllers might animate them to share some of their data. Data subjects with expertise in privacy have more fine-grained control over their data, while data subjects with less knowledge in privacy are supported in understanding their rights and communicating their preferences. While we offer a new form of privacy policy representation and consent, based on existing interdisciplinary research, more research concerning privacy policy representation and consent collection should be performed (Fischer-Hübner and Karegar, 2024). Our proposed

interface may provide input for future research.

To make the most of our presented policy interface concept, it does not suffice to simply put the new interface on existing policies. In our preparation of this work, we tried to convert the Amazon.de privacy policy into P-LPL to be displayed in our privacy policy interface as an example. However, it became clear that this policy did not offer enough detailed information to benefit from the new representation. Some purposes, for example, were written in a way that included the whole world as data transfer destinations. Based on this observation, we argue that simply changing the style of representation cannot achieve transparency on its own. Even if new interfaces are used, policy authors might still be tempted to write legalese instead of using new forms of representing the same information. A new generation of policies is required to be able to increase transparency and empower data subjects. To support data controllers in defining these new policies, additional measures of support could be offered, for example, readability scores. Although the European Union itself suggests the use of a standardized form of policy representation, textual privacy policies are still the default. Regulatory measures might be needed to legally recognize new concepts and convince the data processing industry to adopt frameworks like PriPoCoG.

Although we presented a comprehensive concept and prototypical policy interface, we envision improvements and extensions of this work in the future. Privacy presets, created from opinions of experts or based on the results of privacy studies, may support novice data subjects in protecting their privacy (Papaoannou et al., 2022). We further envision a more guided approach, comparable to a setup assistant. The policy interface could ask the data subject questions regarding their privacy preferences and adjust the privacy policy accordingly. After completing this assistant, our policy interface could be shown with the policy already adjusted to the preferences of the data subject.

REFERENCES

- Colnago, J., Cranor, L. F., and Acquisti, A. (2023). Is there a reverse privacy paradox? an exploratory analysis of gaps between privacy perspectives and privacy-seeking behaviors. In *PoPETs*, volume 2023, pages 455–476.
- Drozdz, O. and Kirrane, S. (2020). Privacy CURE: Consent comprehension made easy. *ICT Systems Security and*

⁶<https://trapeze.ercim.eu/>

- Privacy Protection, pages 124–139. Springer International Publishing.
- Earp, J. B., Vail, M., and Anton, A. I. (2007). Privacy policy representation in web-based healthcare. In *40th Annual Hawaii International Conference on System Sciences (HICSS'07)*, pages 138–138.
- European Parliament and Council of the European Union (2016). Regulation 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation). *Official Journal of the European Union*, pages 1–88.
- Fischer-Hübner, S. and Karegar, F. (2024). *Challenges of Usable Privacy*, chapter 4, pages 103–131. Synthesis Lectures on Information Security, Privacy, and Trust. Springer International Publishing.
- Fox, G., Lynn, T., and Rosati, P. (2022). Enhancing consumer perceptions of privacy and trust: a GDPR label perspective. *Information Technology & People*, 35(8):181–204.
- Gerl, A. (2020). *Modelling of a Privacy Language and Efficient Policy-Based De-Identification*. PhD thesis, Universität Passau.
- Ibdah, D., Lachtar, N., Raparathi, S. M., and Bacha, A. (2021). “Why Should I Read the Privacy Policy, I Just Need the Service”: A study on attitudes and perceptions toward privacy policies. *IEEE Access*, 9:166465–166487.
- Jesus, V. (2020). Towards an accountable web of personal information: The web-of-receipts. *IEEE Access*, 8:25383–25394.
- Kaili, M. and Kapitsaki, G. M. (2023). Improving the representation choices of privacy policies for end-users. In *WEBIST*, pages 42–59. Springer Nature Switzerland.
- Kelley, P. G., Cesca, L., Bresee, J., and Cranor, L. F. (2010). Standardizing privacy notices: an online study of the nutrition label approach. In *SIGCHI*, pages 1573–1582.
- Knijnenburg, B. P., Kobsa, A., and Jin, H. (2013). Counteracting the negative effect of form auto-completion on the privacy calculus. In *Thirty Fourth International Conference on Information Systems*. Citeseer.
- Leicht, J. and Heisel, M. (2023). P2BAC: Privacy policy based access control using P-LPL. In *9th International Conference on Information Systems Security and Privacy*, pages 686–697. SciTePress.
- Leicht, J. and Heisel, M. (2024). Extending PriPoCoG: A Privacy Policy Editor for GDPR-Compliant Privacy Policies. In *ENASE*, pages 307–318.
- Leicht, J. and Heisel, M. (2025). Management of customized privacy policies. In *11th International Conference on Information Systems Security and Privacy*, volume 2, pages 385–396.
- Leicht, J., Heisel, M., and Gerl, A. (2022). PriPoCoG: Guiding policy authors to define GDPR-compliant privacy policies. In *TrustBus 2022*, pages 1–16. Springer.
- Leicht, J., Wagner, M., and Heisel, M. (2023). Creating privacy policies from data-flow diagrams. In *ESORICS 2023 International Workshops*, pages 433–453. Springer Nature Switzerland.
- Lipford, H. R., Watson, J., Whitney, M., Froiland, K., and Reeder, R. W. (2010). Visual vs. compact: A comparison of privacy policy interfaces. In *SIGCHI*, pages 1111–1114.
- Norberg, P. A., Horne, D. R., and Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1):100–126.
- Papaioannou, T., Tsohou, A., Karyda, M., and Karagiannis, S. (2022). Requirements for an information privacy pedagogy based on the constructivism learning theory. In *ARES*, pages 1–8. ACM.
- Paudel, R., Shrestha, A., Dumar, P., and Al-Ameen, M. N. (2023). “it doesn’t just feel like something a lawyer slapped together.”- mental-model-based privacy policy for third-party applications on facebook. In *CSCW*, page 298–306. ACM.
- Proctor, R. W., Ali, M. A., and Vu, K.-P. L. (2008). Examining usability of web privacy policies. *International Journal of Human-Computer Interaction*, 24(3):307–328.
- Raschke, P. and Eichinger, T. (2022). D4.2 - Privacy dashboards. Deliverable, TRAPEZE - TRAnsparency, Privacy and security for European citiZEns. <https://bscw.ercim.eu/pub/bscw.cgi/1274072>.
- Reinhardt, D., Borchard, J., and Hurtienne, J. (2021). Visual interactive privacy policy: The better choice? In *CHI*, pages 1–12.
- Rossi, A. and Palmirani, M. (2019). DaPIS: a data protection icon set to improve information transparency under the GDPR. *Knowledge of the Law in the Big Data Age*, 252:181–195.
- Sailaja, N. and Jones, R. (2017). Industry ideals barriers in using alternative privacy policies. In *HCI*, volume 2017-July. BCS Learning and Development Ltd.
- Tsolakidou, A., Raptis, G. E., Katsini, C., and Katsanos, C. (2024). Exploring the impact of cognitive styles on the visualization of privacy policies. In *PCI*, page 109–115. ACM.
- Windl, M., Ortloff, A.-M., Henze, N., and Schwind, V. (2022). Privacy at a glance: A process to learn modular privacy icons during web browsing. In *CHIIR*, pages 102–112.
- Zimmeck, S., Kuller, E., Ma, C., Tassone, B., and Champeau, J. (2024). Generalizable active privacy choice: Designing a graphical user interface for global privacy control. *Proceedings on Privacy Enhancing Technologies*.