TILT: A GDPR-Aligned Transparency Information Language and Toolkit for Practical Privacy Engineering

Elias Grünewald

Technische Universität Berlin
Information Systems Engineering Research Group
Berlin, Germany
gruenewald@tu-berlin.de

Frank Pallas

Technische Universität Berlin
Information Systems Engineering Research Group
Berlin, Germany
frank.pallas@tu-berlin.de

ABSTRACT

In this paper, we present TILT, a transparency information language and toolkit explicitly designed to represent and process transparency information in line with the requirements of the GDPR and allowing for a more automated and adaptive use of such information than established, legalese data protection policies do.

We provide a detailed analysis of transparency obligations from the GDPR to identify the expressiveness required for a formal transparency language intended to meet respective legal requirements. In addition, we identify a set of further, non-functional requirements that need to be met to foster practical adoption in real-world (web) information systems engineering. On this basis, we specify our formal language and present a respective, fully implemented toolkit around it. We then evaluate the practical applicability of our language and toolkit and demonstrate the additional prospects it unlocks through two different use cases: *a)* the inter-organizational analysis of personal data-related practices allowing, for instance, to uncover data sharing networks based on explicitly announced transparency information and *b)* the presentation of formally represented transparency information to users through novel, more comprehensible, and potentially adaptive user interfaces, heightening data subjects' actual informedness about data-related practices and, thus, their sovereignty.

Altogether, our transparency information language and toolkit allow – differently from previous work – to express transparency information in line with actual legal requirements and practices of modern (web) information systems engineering and thereby pave the way for a multitude of novel possibilities to heighten transparency and user sovereignty in practice.

CCS CONCEPTS

• **Applied computing** → **Law**; • **Information systems** → **Information systems applications**; **Web data description languages**; • **Software and its engineering** → **Formal language definitions**; **Context specific languages**; • **Security and privacy** → **Privacy protections**;

This work is licensed under a Creative Commons Attribution-NoDerivs International 4.0 License.

FAccT '21, March 3–10, 2021, Virtual Event, Canada

ACM ISBN 978-1-4503-8309-7/21/03.

<https://doi.org/10.1145/3442188.3445925>

KEYWORDS

Data transparency, GDPR, data protection, privacy, privacy by design, legal tech, privacy engineering, web privacy, privacy law

ACM Reference Format:

Elias Grünewald and Frank Pallas. 2021. TILT: A GDPR-Aligned Transparency Information Language and Toolkit for Practical Privacy Engineering. In *Conference on Fairness, Accountability, and Transparency (FAccT '21)*, March 3–10, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3442188.3445925>

1 INTRODUCTION

Transparency has been a core principle in philosophical, legal, and technical deliberations around privacy¹ for decades. This particularly applies to data privacy in everyday digital life: To be able to act in a sovereign and self-determined way and actually make *informed* choices, individuals need to have sufficient knowledge about the actual facts and givens regarding the processing of personal data, including, e.g., which party collects what personal data for which purposes, how long this data is (going to be) stored, et cetera. Data protection laws and regulations around the world – such as the European General Data Protection Regulation (GDPR) [8] or the California Consumer Privacy Act (CCPA) [38] – therefore explicitly include transparency rules for data processing parties, obligating them to reveal respective information to the data subjects.

These obligations are today typically fulfilled through respective transparency parts of written privacy policies. Such privacy policies do, however, exhibit several shortcomings that severely limit their actual reception and comprehension on the side of data subjects: First of all, privacy policies are often long, complex, and written in legalese language, making it hard for data subjects to locate transparency-related information and actually understand them correctly [25, 35]. Second, different privacy policies employ different logical structures and vocabularies for factually similar statements, causing significant reading and decoding efforts for every new policy to be understood [26]. Such practice structurally discriminates people who are less privacy-literate [39]. Finally, data transfers towards other parties are now a broadly established practice. Data subjects wanting to see “the whole picture” in such contexts therefore have to read and understand multiple policies and establish the logical interdependencies between them.

Together with the ever-increasing number of services used (and, thus, of privacy policies to be read), this leads to a state where privacy policies, including respective transparency statements, are not read anymore before using a particular service and/or consenting to

¹Being well aware of the slightly different notions between “Privacy” and “Data Protection”, we use these terms interchangeably herein.

a certain collection and use of personal data [29]. Under such conditions, transparency statements increasingly degenerate into rather self-serving formal compliance exercises instead of actually fostering data subjects' informedness of decisions and self-sovereign conduct with regard to privacy.

For other areas of privacy, dedicated technologies have repeatedly been demonstrated to be capable of lowering the cognitive and administrative effort required on the side of data subjects and thereby of re-aligning the actual real-world practice to the (still valid) original intentions behind privacy regulations. Technologies for digitally mediated and partially automated consent provision are one example: They lower individuals' need to care about every single act of data sharing while still allowing them to fine-tune their consent to individual preferences in line with legal requirements for specificity, thus counteracting ongoing trends towards overly broad consent provision [40].

Following the paradigm of privacy by design, which basically requires *all* privacy principles to be appropriately reflected in technology², we aim to achieve the same for the principle of transparency herein. Our goal is thus to provide technical artifacts that significantly lower user-side efforts required for gathering and comprehending transparency information while at the same time providing the capabilities necessary to meet regulatory requirements and, thus, to actually be applicable in practice. An indispensable precondition for such artifacts is the capability to represent relevant transparency information in a structured, machine-readable format and the possibility to easily employ this representation in real-world information systems. For this purpose, we particularly provide the following contributions:

- An in-depth analysis of transparency information that needs to be expressible to meet the requirements of the GDPR
- A formal specification of a structured, machine-readable language meeting these expressiveness requirements – the “Transparency Information Language”
- Two fully implemented libraries for widely used programming languages (Python and Java) allowing to easily process and manage respective representations in real-world information systems – the most important parts of the “Toolkit” around our language
- A demonstration of the opportunities arising from the machine-readable representation and technical mediation of transparency information through two exemplary case-studies for *a)* inter-organizational analyses of stated data transfers and *b)* novel, adaptive user interfaces

We thus explicitly follow an engineering-driven approach, providing novel, practically usable technological artifacts that address a currently open socio-technical challenge. Our respective considerations and contributions unfold as follows: Section 2 summarizes related work on transparency enhancing technologies and identifies requirements for the aspired language and toolkit. Afterwards, section 3 distills the necessary expressiveness of a transparency information language based on structured GDPR study. Based on these deliberations, the new transparency information language is formally defined and technically implemented (see section 4). In

section 5, the language is practically embedded into a toolkit with storage and interoperability functionalities. In section 6, we then demonstrate two exemplary applications, which enhance transparency using the language and toolkit components. Eventually possible future work is conceptualized in section 7. Finally, section 8 concludes.

2 RELATED WORK & REQUIREMENTS

Technical approaches to (privacy-related) transparency are commonly discussed under the term of transparency-enhancing technologies (TETs). Such TETs exist in a broad variety of forms addressing transparency from many different angles. In particular, TETs comprise such diverse technologies as generating a location history based on social media entries [20], visualizing data exports [22], phone sensor permission management [43], or in-browser cookie tracking [33]. For categorizing these TETs, different taxonomies exist [13, 18]. Zimmerman [44] particularly distinguishes *ex-ante*, *realtime*, and *ex-post* TETs, with *ex-ante* TETs providing information about a processor's *intended* data collection and processing before the collection. Our aspired goal of making transparency-related statements more accessible and comprehensible for data subjects to heighten the informedness of their decisions clearly falls into this *ex-ante* category. Often mentioned “privacy dashboards” providing an integrated view to current and past possession and processing of data [1, 4, 34], in contrast, would fall into the *ex-post* category and thus clearly differ from our intent. Within the sub-domain of *ex-ante* TETs, in turn, the herein addressed problems associated with written transparency statements not meeting their originally intended goals are particularly subject to two research strands: Automated knowledge extraction from privacy policies and pre-existing transparency languages. Both shall be introduced briefly before distilling requirements for our aspired language and toolkit.

2.1 Privacy Policy Knowledge Extraction

One established approach for technically addressing transparency-related problems is to extract respective information from pre-existing policies by means of natural language processing (NLP) and to then provide additional functionalities on top of the so-extracted data. This particularly includes attempts based on extensive static rules and named entity recognition [6] as well as more dynamic ones employing methods of crowd sourcing and machine learning [11]. So far, however, these typically focus on merely identifying blocks from a privacy policy dedicated to different subjects (like “collection”) instead of actually *extracting* information at a sufficient level of detail [5, 37] or achieve accuracies that are insufficient to be considered a reasonable basis for providing functionality with legal relevance on top of them [6, 11]. In addition, respective endeavors do typically not pay explicit regard to the actual and detailed regulatory requirements. A structured, machine-readable and interoperable representation that can be used for a multitude of different purposes is typically also out of scope.

The higher-level functionality aspired and sometimes even implemented in above-mentioned projects may, however, serve as blueprint for what we intend to facilitate with our structured transparency language and toolkit. In particular, this includes enhanced

²See, for instance, Art. 25 GDPR requiring technical measures “designed to implement data protection principles” and, thus, covering all principles mentioned in Art. 5.

graphical presentations intended to improve comprehensibility [11], chatbots providing answers about a company's personal data practices [12], or automatically generated icon representations allowing for a quick and intuitive overview of personal data practices [9].

2.2 Transparency Languages

Given the insufficient prospects of starting with pre-existing, textual transparency information from privacy policies, an alternative approach is to start with a structured and machine-readable representation of the transparency information to be provided. Frameworks proposed in this regard so far – like respective parts of the W3Cs early P3P standard [7] or industry-driven initiatives like IAB's Transparency & Consent Framework [16] – do, however lack the expressiveness actually required by privacy regulation. For instance, P3P provided a severely limited vocabulary of pre-defined purposes which does not suffice to express actually occurring purposes in the required level of specificity [see, e.g., 40] while the IAB Europe Framework completely neglects several information obligations (e.g. rights to access, automated decision making, onward transfer etc.) and claims registration fees for participation. More recent policy languages [e.g., 10, 23] come closer to the legally required expressiveness aspired herein. However, these typically try to cover *all* aspects of a privacy policy from the outset, resulting in a considerable level of complexity (e.g. considering the problem of de-identification mixed with transparency obligations) and/or still resemble the concept of limited, fixed vocabularies known from earlier proposals³. At the same time, they typically lack publicly available and easily adoptable (reference) implementations in the form of re-usable libraries, severely limiting their practical applicability in real-world (web) information systems. The effective adoption of privacy languages also has been detained because of missing content validation and (not even rudimentary) user interfaces.

Altogether, none of the transparency-related technologies and languages proposed so far thus satisfies our intent to represent transparency information as required by the GDPR ex-ante in a machine-readable form and in line with actual legal requirements while at the same time being easily usable in practice.

2.3 Requirements

Based on the above and on other works in the domain of information systems engineering in general and practice-oriented privacy engineering in particular [esp. 31], we can identify a set of requirements our aspired language and toolkit have to fulfill. These comprise specifications regarding *what* functionality is to be implemented (functional requirements, FR) as well as rather non-functional requirements (NFR) describing additional characteristics that the language and toolkit must provide in order to ease and foster practical adoption.

Req. 1 (FR): Sufficient expressiveness to meet legal obligations. The most important requirement is a functional one. Our language must be capable of expressing all transparency information required by the GDPR in a sufficiently specific form. Only when this is the case,

³We believe that fixed vocabularies, provided within the language, either do simplify far too much by using categories that are not very meaningful, or they cannot, by definition, reflect all industries, areas of life, or possible processing activities.

our language can actually form a sustainable basis for novel technical approaches to transparency that still unfold legal relevance and, thus, overcome the limitations of previous approaches like P3P (see above). The details constituting this requirement will be elaborated further in section 3.

Req. 2 (FR): Possible standardization of commonly used transparency information items. Another goal regards the standardization of as many transparency information items as possible. Through the introduction of structured-types, nomenclature and format conventions, the content becomes more accessible and interoperable. Besides fostering the independent development of different components on top of commonly shared semantics, this would also allow to create, for instance, templates for transparency declarations, standardized mappings to textual representations, etc.

Req. 3 (FR): Support for inter-system communication and versioned persistence. Our language is explicitly intended to be used in online contexts, with transparency information being exchanged between multiple sub-components of an overall system, external services, user-side applications, etc. Capabilities for communicating and storing respective representations as well as for referring to them across component- and system boundaries in a standardized manner are therefore indispensable. In particular, this should be possible in a version-aware form, allowing to update transparency information while still keeping previous versions available.

Req. 4 (FR): Application- or Domain-Specific Extensibility. The benefits of standardization laid out in req. 2 notwithstanding, specific application domains, service providers or even single applications might raise additional transparency requirements that cannot reasonably be covered in a generalized manner. This could, for instance, be the case because of sector-specific regulatory obligations going beyond those from the GDPR. To also facilitate the usage of our artifacts in such settings, they must allow to implement respective extensions on top of their underlying, general functionality. In particular, this should be possible in a way that preserves the applicability of extensions in case of our general components being updated (thus speaking against domain-specific forks, for example).

Req. 5 (NFR): Implementation as reusable artifact. A core principle of software engineering is to not implement similar functionality multiple times but to encapsulate respective functionality in re-usable components as far as possible. Besides avoiding inconsistencies between different implementations and allowing to focus development and maintenance efforts on just one target, this also fosters wider adoption and simpler integration into real-world systems. Our technical artifacts should therefore be provided in the form of reusable artifacts that can be integrated into relevant development environments as seamlessly as possible.

Req. 6 (NFR): Developer-Friendliness and low implementation overhead. Besides the provision as reusable artifacts, developer-friendliness and low implementation efforts are also crucial for facilitating adoption in practice. On the one hand, this is the case because developers will try to avoid sophisticated and hard-to-use technical mechanisms while intuitive and easily integratable ones might be adopted out of developers' intrinsic motivation or mere curiosity [31]. On the other hand, developer-friendliness and ease

of implementation are also relevant for a more formal reason: Art. 25 of the GDPR requires “appropriate technical [...] measures” to be implemented for materializing data protection principles (including transparency). The appropriateness, in turn, depends (amongst other factors) on “the costs of implementation”. The lower the implementation efforts (and, thus, costs) for adopting a novel technical mechanism are, the more likely is this adoption therefore to be considered obligatory. Developer-friendliness and low implementation efforts thus also have legal implications.

3 EXPRESSIVENESS REQUIREMENTS

To unfold legal relevance as a possible replacement for written transparency statements and, thus, to provide actual societal and business value, a transparency language must be capable of expressing all legally required transparency information at a sufficient level of detail. We refer to this capability under the term of *expressiveness*. Given its prominent role as a blueprint for privacy and data protection regulations worldwide, we extract respective requirements from the European General Data Protection Regulation (GDPR) [8].

Here, transparency requirements are set out in Articles 12-15. More precisely, Art. 12 requires controllers to provide transparency information to data subjects “in a concise, transparent, intelligible and easily accessible form”. Even though implicitly assuming the provision through traditional, purely textual policies, the GDPR also foresees alternative communicative channels such as standardized and machine-readable icons (see Art. 12 (7)).

The transparency information to be provided to data subjects proactively, in turn, are laid out in Art. 13 and 14 in more detail for different settings (distinguishing cases where personal data is collected from the data subject from those where it is obtained from elsewhere). Beyond these, Art. 15 defines information to be provided to data subjects upon request, particularly also including the personal data itself. Leaving aside the latter (which is overly complex and diverse across different data controllers to be covered by a uniform language as aspired herein), the transparency obligations from Art. 13-15 strongly overlap and shall serve as the basis for determining the required expressiveness herein.

Besides the mentioned ones, Art. 30 – obligating data controllers to maintain “records of processing activities” – also exhibits strong overlaps regarding the information to be kept in such records. Even though serving the principle of *accountability* and not the one of *transparency* (cf. Art. 5(2)), our aspired language may thus also play a role in this regard later. We therefore include the requirements from Art. 30 in our analysis. The resulting summary of all relevant articles is provided in table 1.

Each of these transparency requirements can be translated into a corresponding building block for our aspired transparency information language. The required data fields and respective data types to be included in these blocks result either from the GDPR itself or from technical considerations regarding unambiguousness. In some cases, the exact data fields and types to be provided are not externally given and we had to make reasonable determinations to allow for a well-specified formal representation in subsequent steps. Due to space constraints, the considerations behind each and every field and type cannot be laid out herein. Two building blocks shall, however, exemplarily be examined in some more detail:

Table 1: Transparency information obligations by the GDPR

Reference(s)				Transparency information
13 (1a)	14 (1a)	30 (1a)		Controller
13 (1b)	14 (1b)	30 (1a)		Data protection officer
13 (1c)	14 (1c)	15 (1a)	30 (1b)	Purposes
13 (1c)	14 (1c)			Legal basis
13 (1d)	14 (2b)			Legitimate interests
13 (1e)	14 (1e)	15 (1c)	30 (1d)	Recipient (categories)
13 (1f)	14 (1f)	15 (1c)	30 (1e)	Third country transfer
13 (1f)	14 (1f)	15 (2)	30 (1e)	Adequacy (third country)
13 (1f)	14 (1f)	15 (2)	30 (1e)	Access and Data portability
13 (2a)	14 (2a)	15 (1d)	30 (1f)	Retention or storage criteria
13 (2b)	14 (2c)	15 (1e)		Right to request access
13 (2b)	14 (2c)	15 (1e)		Right to correction or deletion
13 (2b)	14 (2c)	15 (1e)		Right to data portability
13 (2c)	14 (2d)			Right to withdraw consent
13 (2d)	14 (2e)	15 (1f)		Right to complaint
13 (2e)				Necessity/Non-disclosure
13 (2f)	14 (2g)	15 (1h)		Automated decision making
	14 (2f)			Sources
13 (3)				Notification on purpose change
		30 (1c)		Data subjects/Data disclosed

Data Protection Officer: Art. 13 and 14 of the GDPR require “the contact details of the data protection officer” to be provided whereas Art. 30 requires “the *name and* contact details” (emph. added) to be included in the record of processing activities. As no further definition is provided regarding *which* contact details are required, we had to make an educated assumption here and declared a postal address, a country code, and an email address obligatory. In addition, we foresee the phone-number as an optional field. The name of the data protection officer, in turn, is obligatory only for Art. 30. It is therefore not an obligatory element for our language to be defined in subsequent steps. Regarding the formats, standardized representations exist and are therefore used for country codes, phone numbers, and email addresses. The remaining fields, in turn, are considered as simple strings.

Data Disclosed → Legal Basis: The building block “Data Disclosed” integrates multiple transparency obligations related to particular categories of personal data collected and processed. For instance, Art. 13 and 14 both require to inform about “the legal basis for the processing” of personal data. This information has – like those relating to other obligations in this building block – to be provided in relation to specific categories of personal data. It must therefore be possible to repeat the whole “Data Disclosed” block multiple times, once for each category of personal data concerned.⁴ The legitimating legal basis, in turn, can be a provision from the GDPR

⁴See Data Disclosed → Data category.

itself or from any other law or regulation such as a national disease control regulation, which, in some cases, may require additional remarks or explanations. We therefore deduce the need for a bipartite statement here, comprising an obligatory standardized reference to a legal provision and an optional free-text description.

In a similar vein, we analyzed all transparency obligations and specified required blocks for our transparency information language. A summary of the resulting blocks and data fields that need to be expressible is provided in table 2. Notably, this also includes a separate building block with meta-information allowing to uniquely identify a particular record of transparency information, to specify validity dates, or to ensure its integrity. These allow for more advanced functionalities at later stages.

4 LANGUAGE DESIGN AND IMPLEMENTATION

Based on the above analysis of the required expressiveness and the identification of data formats to be used, we designed our transparency information language (TIL), following well-established steps from language design [27] also employed for other privacy-related languages [cf. 40]. In particular, this includes a complete formal grammar in Extended Backus-Naur form (EBNF) [17] and a complete JSON¹⁸ Schema v7 specification [32]. Both are, again due to space constraints, available online.¹⁹ Some core concepts shall, however, be briefly elucidated below.

4.1 Formal definition

Serving as a reference for different implementations, the formal definition of our language translates each of the building blocks

⁵Follows the database-specific implementation; should offer as much entropy for globally unique identifiers.

⁶All language abbreviation codes follow the established ISO 639-1 standard as identifiers for names of languages.

⁷The hash is based on one SHA256 calculation of the document content.

⁸Differentiates internally of a company; relevant for large companies.

⁹All country codes follow the established ones ISO 3166 country abbreviation standard.

¹⁰The legitimate interest only has to be stated if the processing is carried out in accordance with Art. 13 (1d).

¹¹Thus, any third country transfers are also defined.

¹²“A time period is represented in the format [start date]P[YY][MM][WW][TD][T[hH][mM][s.f][S]]. The *P* indicates as a leading information letter that a period follows. Periods that contain a portion of the time are delimited by a *T*, as in the specification of the start time. It is therefore possible to distinguish between the months and minutes (M). The same rules apply for formatting the starting time as for normal dates.” (own translation) Cf. https://de.wikipedia.org/wiki/ISO_8601#Zeitspannen. Note: In most cases, the start date would probably be omitted or implicitly results from the time the personal data were collected.

¹³The TTL (“time to live”) specifies a specific “lifetime” in the form of a time span.

¹⁴The aggregation function describes the calculation basis when specifying several time intervals. For example, if there is storage for 2 weeks for technical reasons (e.g. backup), but there is a legally longer retention period, the maximum aggregation function (max) would be selected (standard case).

¹⁵According to Art. 45. The “suitable guarantees can be available without special authorization from a supervisory authority” (Art. 15) under certain conditions.

¹⁶This duty to provide information is limited to the collection of personal data that does not take place from the data subject (Art. 14).

¹⁷Refers to the necessity and consequences in the event of non-disclosure: According to Art. 13 (2e), this refers to the information whether the provision of the personal data is required by law or contract or is required for the conclusion of a contract, whether the data subject is obliged to provide the personal data and the possible consequences of not providing it.

¹⁸For the choice of JSON over other formats, see below.

¹⁹<https://github.com/Transparency-Information-Language/schema>

Table 2: Building blocks to be expressed within TILT

Meta	Access/Data portability
1. Identification Number* ⁵ 2. Name* 3. Creation date* 4. Modification date* 5. Version* 6. Language code* ⁶ 7. Status* 8. URL* 9. Hash* ⁷	1. Possibility available?* 2. Description accessibility 3. URL 4. E-mail 5. [Identification evidence] 6. Administrative fee a. Amount b. Currency 7. Data format(s)
Controller	Sources ¹⁶ , per item:
1. Company name* 2. Division ⁸ 3. Address* 4. Country code* ⁹ 5. Name (representative)* 6. Email (representative)* 7. Phone (representative)	1. Data (category) 2. [Sources] a. Description b. URL c. Publicly available?
Data Protection Officer	Rights to information, rectification/deletion, data portability, withdraw consent, for each:
1. Name 2. Address* 3. Country code* 4. Email* 5. Phone	1. Possibility available?* 2. Description 3. URL 4. E-mail 5. [Identification evidence]
Data Disclosed, per item:	Right to complain
1. Data category* 2. [Purpose]* a. Purpose* b. Description* 3. [Legal basis]* a. Reference* b. Description 4. [Legitimate interest]* ¹⁰ a. Exists?* b. Reasoning 5. [Recipient (category)]* ¹¹ a.–d. (as controller) e. Category* 6. [Storage periods]* a. temporal ¹² : TTL ¹³ b. conditional: Purpose c. conditional: Legal basis d. Aggregation function ¹⁴	1.–5.: see above 6. Supervisory authority* 7.–10.: contact details
	Non-disclosure ¹⁷ , per item:
	1. Data (category)* 2. Legal requirement?* 3. Contractual regulation?* 4. Obligation to provide?* 5. Consequences if not provided*
	Automated decision making
	1. In use?* 2. Logic involved 3. Scope and intended effects of processing for the data subject
	Notification on change
	1. Affected categories 2. Date time of change 3. URL of last document

Legend:

*Item**: Mandatory information
[Item]: List (possibly empty)

from table 2 into a finite set of production rules. In anticipation of the JSON Schema implementation, some extra fields are necessary to that end. The complete set of productions creates a context-free grammar.

First (cf. listing 1), the root element `tilt` consists of three associated non-terminal symbols which are `header` and `properties`, and an optional `[addProp]`. The latter allows any customization and extension of the language for their users if needed (see *Req. 4*). After having defined the root element and associated non-terminals, all property non-terminals are resolved to either their relative non-terminal symbols or directly to basic types, which are described at the end of the specification. The meta non-terminal field `id`, for instance, directly resolves to a `string` value which is a basic type. All other non-terminals describe the individual fields of the building blocks listed in table 2.

```
tilt := header, properties, [addProp];

header := schema, id, title, description, examples, [addProp];
schema := 'http://json-schema.org/draft-07/schema';
20 id := 'https://github.com/Transparency-Information-Language/schema';
title := 'Transparency Information Language';
description := string;
examples := { properties };

25 properties := meta, controller, dataProtectionOfficer,
  dataDisclosed, thirdCountryTransfers, accessAndDataPortability,
  sources, rightToInformation, rightToRectificationOrDeletion,
  rightToDataPortability, rightToWithdrawConsent, rightToComplain,
  automatedDecisionMaking, changesOfPurpose, [addProp];
```

Listing 1: Grammar for root element, JSON Schema meta fields, and main properties

An exemplary building block is formally defined in listing 2 and visualized in figure 1: The `dataDisclosed` property definition reflects the cases according to Art. 13(1f) GDPR in which the recipients or categories thereof need to be specified. In many cases, the non-terminal symbols do not directly resolve to terminals as in the case of purposes: The additional production rules allow the recursive definition of a set of (purpose, description) pairs.

```
65 dataDisclosed := {
  id, category,
  purposes, legalBases,
  legitimateInterests,
  (recipients | category),
  storage, nonDisclosure,
70 [addProp]
};

category := string;
purposes :=
75 (purpose, description)
  | purposes;
purpose := string;
...
```

Listing 2: Grammar

Figure 1: Syntax diagrams

All fields are syntactically defined, not necessarily semantically. However, many fields carry implicit semantics in their syntax: The `stringDate` or `stringEmail` (not shown) fields are such examples

for which the exact format does not allow semantically different inputs. These specifications are of particular relevance when technically implementing the language because the meta language at least needs to support data validation via regular expressions.

Finally, the specification comprises all basic types which resolve to terminal symbols. These include `string`, `boolean`, and `number` values. For many of them regular expressions (or type descriptions referring to regular expressions stated in there) are given. Since JSON is used as the meta language of expression, its specification is directly taken and adapted from the official grammar. [19] We provide a visualized version of the grammar including all syntax diagrams online.

4.2 Technical implementation

One of the main contributions of this paper is the full language implementation in the form of a JSON Schema. Alternative technical implementations could also be developed with XML/XSD or RDF/RDFS. For reasons similar to the ones listed in [40] – human readability, robustness, broad programming language and tooling support, and resource efficiency – we opted for JSON/JSON Schema as the default representation format.²⁰ In the following, we first present the core language implementation. In a second step, we then demonstrate how this implementation can be extended.

4.2.1 Core language. Regarding the core language model, in 2692 lines of code we represent every single field of all introduced building blocks. Each one has its own type definition, including a comprehensible description, at least one example value and in most cases further type specifications through regular expressions or other means in JSON schema. The schema also contains a complete document covering most of the possible paths of the syntax tree. Some properties, such as `$schema`, `$id`, `title` etc., are (de-facto²¹) mandatory JSON Schema meta properties which are used to describe the document for validators. Afterwards, the language definition according to the formal definition is carried out. Moreover, all building blocks are marked as required since all of these properties have to be expressed explicitly (addressing *Req. 1* and *Req. 2*). The language allows additional properties as pointed out in the formal definition above, which is relevant for controllers that want to extend the schema by their needs apart from GDPR transparency information (*Req. 4*).

An example for custom type validation is the reference to a legal basis for data processing. Several conventions were introduced (see table 2) that need to be enforced within the technical implementation. In this example, the convention is materialized using a regular expression and its description as depicted in listing 3. Likewise, valid and invalid examples are given in listings 4 and 5. The first example successfully validates against the regular expression, whereas the second does not. Publicly available JSON Schema validators quickly reveal erroneous fields to be corrected.

²⁰Other representations, particularly in XML, might be added later with reasonable effort. First steps in that direction are experimentally documented: <https://github.com/Transparency-Information-Language/tilt-experimental/tree/master/xsd>

²¹The minimal valid JSON Schema document is an empty object type (trivial).

```

"reference": {
  "$id": "#/properties/dataDisclosed/items/anyOf/0/properties/
    legalBases/items/anyOf/0/properties/reference",
  "type": "string",
  "title": "Reference",
  "description": "Refers to the reference in a legal regulation (law,
    directive, declaration etc.). Format follows pattern.",
  "pattern": "^[A-Z]*([-]?[0-9]*|[a-z]*)*$",
  "examples": [ "GDPR-99-1-a" ]
}, ...

```

Listing 3: dataDisclosed/legalBases/reference schema

```

"legalBases": [
  {
    "reference": "GDPR-6-1-a",
    "description": "The data
      subject has given..."
  }, ...
  {
    "reference": "CCPA 1798",
    "description": "A consumer
      shall..."
  }, ...

```

Listing 4: Valid legal basis**Listing 5: Invalid legal basis**

Similar to the above-mentioned mechanisms, other fields are also validated. For instance, timestamps, language and country codes or international telephone numbers are validated using custom regular expressions according to the conventions stated above. JSON Schema also has built-in types, which are used through the format keyword. For instance, the uri-reference format is used to validate HTTP links to the controller’s web page or the accessAndDataPortability measures.

Besides the validation of singular fields within (sub)schemas of properties, overarching interrelationships are validated by our language implementation as well. As laid out in Art. 22 GDPR and the Recitals 71, 72 and 91, data subjects needs to be informed about the logic involved in case automated decision making is in use. Consequently, there is an interdependency between existing fields within the language schema. Such requirements are solved using logical operators within JSON Schema. An if-then construct enforces the appearance of logicInvolved and scopeAndIntendedEffects, whenever the (automatedDecisionMaking) inUse property is set to true.

Another example for higher-order capabilities of our language is the GDPR requirement to state the “the period for which the personal data will be stored, or [...] the criteria used to determine that period” according to Art. 13 (2a). Taking into account the requirements from real-world scenarios, different modalities to fulfill this requirement are apparent. The formal definition sets up three possibilities to define the time period or respective criteria. In the technical implementation the choice is constructed using an anyOf environment. The logic is implemented as part of the storage item. Consequently, controllers can express temporal as well as purposeConditional and legalBasisConditional options.

Building domain-specific languages is an “incremental, modular, and extensible way from parameterized building blocks” as [27] state. The building blocks as worked out in table 2 might be extended or changed over time, since a new iteration of GDPR is about to come. Therefore, the introduced language model has to be seen as a “living” specification that will have to adapt to new legislation over time.

In order to motivate the usage of our language in practical privacy engineering, we demonstrate several components of a toolkit

that has been specifically designed around the technical implementation at hand in section 5. Before that, we illustrate how the language capabilities can be customized for a specific application domain in the following section.

4.2.2 Extension through vocabularies. In many plausible cases, it should be possible to extend the expressiveness of the language. Reasons for this can be controller-specific requirements or the binding to other language definitions. Therefore we show how to specify single fields by using the example of purpose definitions. In [31] the authors require purpose definitions to be hierarchical and based on flexible vocabularies with allowed and prohibited values. In the context of personal data sharing, different purposes may define “broad” or “narrow” legal consent respectively for which level of transparency should be guaranteed. Examples for such hierarchies are the following:

- Research / Marketing / ... (*broad*)
- Clinical research / Advertising / ... (*rather broad*)
- COVID19 research / Targeted advertising / ... (*rather narrow*)
- Polymerase chain reaction testing / Tracking technologies including mouse movements / ... (*narrow*)

Our transparency information language can be extended at attribute level. Given the purposeConditional field, we can deviate from the simple string array and formulate a custom vocabulary in three steps²²: First, we create a JSON schema for the new purpose field specifying all values and necessary rules as a stand-alone document. Using the allOf, anyOf, and not keywords every combination of purposes can be formulated. There is full flexibility on the purpose vocabulary. An accepted authority could publish a full definition of purposes under an open content license. Then, industry can adopt the vocabulary to incorporate new transparency standards. Adding the enum, const and not meta language elements, there can even be made a distinction between allowed and prohibited purposes. Secondly, we change the items specification within the original schema to the uri-reference where our new purpose vocabulary file can be found (online). Finally, we can validate the extended schema using existing schema validators.

As a consequence, this extension shows how the language definition can become more expressive using externally defined schemata. Therefore even full complementary privacy preference languages can be integrated. If another language is also implemented in JSON schema, data controllers can profit from the capabilities of both languages at the same time (*Reqs. 4, 5, 6*).

5 TOOLKIT

To meet the requirements for re-usability (*Req. 5*) and developer-friendliness (*Req. 6*), we also provide companion libraries for two programming languages widely used in the domain of (web) information systems: Java and Python. These libraries allow to programmatically create, manage, and validate transparency information objects. In addition, we crafted a fully implemented document storage with an easily usable remote API, providing inter-system communication and versioned persistence (*Req. 3*) and heightening re-usability (*Req. 5*) and developer-friendliness (*Req. 6*) even further.

²²For a detailed description, refer to <http://github.com/Transparency-Information-Language/vocabularies>

5.1 General Purpose Programming Language Bindings

With the above language alone, transparency information now can be expressed in a standardized manner using the JSON Schema implementation. The manual creation of these JSON documents can be supported by (visual) editors. However, a manual workflow is error-prone and difficult to distribute among several people or departments in a company. Therefore, we provide two general purpose programming language bindings that allow to express the transparency information building blocks as native language elements.^{23,24} After the creation of such elements, these can be serialized as JSON documents according to the schema. Moreover, already existing JSON documents that successfully validate against the schema can be imported into the programming language at hand. Here, we show two implementations in Java and Python, demonstrating that and how transparency information can also be generated in a developer-centric manner in program code. This approach allows the involvement of the developers who ultimately implement the actual data processing and who know best which data is processed where and for which purpose. Other language bindings are planned to complement the toolkit as future work (candidate languages are JavaScript, Kotlin, Swift etc.) for covering a more diverse field of application development.

The Java implementation of our language allows the dynamic creation of JSON documents according to the language schema. The core consists of 34 classes. For each class, the fields are only accessible through *Getter* and *Setter* methods. Each field is annotated using the *jackson* library for efficient (de)serialization. Additional *toString()* methods help for testing and logging purposes. Listing 6 represents the low-effort use of the implementation (*Req. 6*) in a shortened form: (1) The central *Tilt* object type encapsulates all pieces of transparency information which are empty by default. (2) For the validation of external documents only the URL has to be provided. (3) Existing documents can be de-serialized and manipulated using the *Converter* class. For seamless distribution, the software project is automatically packaged using the Github automation workflow. As a result, the package is available through popular dependency repositories by adding the reference to our implementation to the traditional *pom.xml* file (*Req. 5*). We emphasize that by providing the package, it is now possible to integrate the language into the context of large distributed systems. Java is particularly popular in the context of service-oriented architectures and data-intensive web services [42]. Modern microservice architectures build upon loosely coupled message oriented communication models (*Req. 3*). Through the included JSON adapters, the integration should be realizable with little development efforts (*Req. 6*).

```
// 1 - Generate new Tilt instance from scratch
Tilt tilt = new Tilt();
Controller controller = new Controller();
controller.setName("Example Company SE");
tilt.setController(controller);
// 2 - Validate existing documents
TiltValidator.validateDocumentFromUrl(DOCUMENT_URL);
```

²³<https://github.com/Transparency-Information-Language/java-tilt>

²⁴<https://github.com/Transparency-Information-Language/python-tilt>

```
// 3 - De-serialize and manipulate existing documents
Tilt t = Converter.fromJsonString(instance);
t.getMeta().setHash("42");
```

Listing 6: Java language binding example.

The Python library functions similar to the previously described implementation. It consists of class representations and helper functions for type checking and conversion. With the exception of the *json* library, there are no other dependencies, which makes the full language binding uncomplicated and easy to integrate into existing software projects (*Req. 5*). Along with the core module comes the documentation of the open source project and several example programs that are linked to an interactive online playground. The library, documentation, etc. are provided as a PyPI package.²⁵ Therefore, the tool can be installed by every developer using the standard package sources.²⁶ For the validation tasks, we recommend the *fastjsonschema* library.

Both language bindings have in common that they facilitate the adaptive creation of transparency information documents. Having the tool at hand allows developers to manifest the transparency notice in source code on a par with the relevant data processing logic (*Req. 6*).

5.2 Document storage with Application Programming Interfaces

In order to complement the toolkit and for exploring possible storage and interaction models with regards to machine-readable transparency information, in this section the component *tilt-hub*, an extensible document storage, is introduced.²⁷ Within *tilt-hub*, documents expressed in our language can be stored and versioned.²⁸ Furthermore, two APIs enable 1) queries to retrieve the exact transparency information that is needed for a specific use case and 2) queries to reveal data sharing networks if the documents are describing several data processing entities.

Figure 2: System architecture of the document storage.

²⁵<https://pypi.org/project/tilt/>

²⁶`pip install tilt`

²⁷<https://github.com/Transparency-Information-Language/tilt-hub>

²⁸For this purpose, we introduced the *meta* building block containing modification date, status, hash, and identification number. These fields can express the timeliness, sequence, integrity, and status of documents.

Figure 2 shows the system architecture which relies on the concept of microservices. Each implemented service can be scaled up and down individually using docker-compose. The central component is an instance of MongoDB, a JSON based document database, per default running as a single node cluster. It could be configured to run on horizontally scaled infrastructure to be able to store millions of transparency-related documents. Since MongoDB is published under an open source license, a big ecosystem of various toolkits, APIs and GUIs already has evolved. In particular, this includes the Mongo Express web-based user interface and driver support for almost all popular programming languages. In addition to the aforementioned native drivers for MongoDB, two language-agnostic application programming interfaces were developed for the document storage, namely a GraphQL API service and a REST API service. Both are realized as independently functioning microservices. By the help of these services, clients can for instance perform search queries, integrate the transparency information language into their continuous integration or deployment pipelines, or connect other privacy-related software components such as personal information management systems. When building applications, these techniques and the developed language bindings play hand in hand to lower implementation costs for the efficient storage and queries of transparency documents (*Reqs. 3 and 6*). To the best of our knowledge, these are distinguishing features, not entirely present in related work.

Furthermore, the infrastructure comprises two additional *triggering* microservices. Both of them show the database interaction within the internal *tilt-hub* network. The Python trigger serves as a starting point for developers who want to integrate the document storage into their CI/CD pipelines. An exemplary use case might be the notification in case of changes: A developer could have implemented a new function in one of the data controller's services that introduces a new third country transfer. Consequently, the developer pushes the change to *tilt-hub* e.g. using the REST API. This upsert would trigger the Python microservice to send an email or short-message notification to the law department informing about the successful update. The Python microservice uses the native language binding described above and connects to the database using pymongo. Additionally, a JavaScript microservice using node.js was implemented for the continuous validation of all changes to the database with respect to our JSON schema. This microservice had to be developed since the integrated schema validation capabilities of MongoDB are limited to version 4 of the JSON schema reference implementation [28] while our language relies on the more recent version 7. For the schema validation we employ the node.js library Ajv that has been benchmarked as the fastest validator available for JavaScript at the time. We propose the usage of *tilt-hub* for (i) internal deployments only within the control sphere of a data controller for better internal document management, or (ii) also with external access for interested data subjects and third parties. Moreover, (iii) data protection officers or supervisory authorities could use the interfaces to interact with the transparency document storage for their tasks. Consequently, different deployment models offer several options for business process optimizations because

they replace manual information management and analysis tasks (cf. *Req. 2*).²⁹

6 EXEMPLARY APPLICATIONS

To demonstrate the opportunities for data subjects' informedness and data sovereignty unlocked by the existence of a structured, machine-readable language and respective tools as introduced above, we implemented two exemplary applications utilizing different capabilities and characteristics of these artifacts: A transparency analysis platform allowing to analyze stated data usage and sharing practices across multiple data processing parties and a simple browser extension illustrating the possibilities for novel user-facing representations of transparency information (cf. *Req. 2*).

6.1 Transparency Analysis Platform

We built a transparency analysis platform that serves as a tool for inspection and analysis tasks related to the presented transparency information. With this tool, it is possible to map respective information of a multitude of data controllers onto a graph structure. This allows to reveal potential data sharing networks among multiple data controllers that are hard or impossible to grasp by individual data subjects today as doing so would require scanning thousands of traditional privacy policies and putting them into context. In order to make use of publicly available transparency information from existing transparency and consent frameworks such as IAB Europe, we provide this ready-to-use analysis tool.

An already established software stack for graph databases serves as technical basis (grandstack.io). The prototypical implementation can be subdivided into data extractors and data processing components (Neo4j graph database, GraphQL API, and accompanying libraries). As an input, our platform accepts documents in the two different versions of the IAB Europe framework as well as documents expressed in our proposed language. The output is a visual graph structure or textual query results depending on the interfaces that are used.

With our tool, we were able to combine transparency information documents from different sources and explore complex controller-processor-purpose interrelations. These e.g. include "isolated", "networked" or "linked" controller relationships which, for instance, may point at unforeseen knowledge concentrations or risks of interlinking data from different sources. Future work on the analysis platform includes pre-defined queries that answer specific questions well-known from social network or graph analysis studies [24, 36, 41]. More details on the platform, its capabilities and application for GDPR transparency are envisaged for dedicated publication. For now, we prove the applicability of our language model within a practical privacy engineering software artifact.³⁰

6.2 Internet Browser Extension

At European level, the Art. 29 Working Party (now EDPB) [2] recommended "that layered privacy statements/notices should be used [...] rather than displaying all such information in a single notice

²⁹<https://github.com/Transparency-Information-Language/tilt-sample-analysis>

³⁰<https://github.com/Transparency-Information-Language/transparency-analysis-platform>

on the screen” (cf. *Req. 1*). Hence our toolkit and applicability examples are augmented by the development of an internet browser extension that is able to summarize key aspects of the transparency information expressed in our language.

We target the Google Chrome browser platform to reach the large majority of internet users. Therefore, the architecture of the browser extension follows the developer guidelines provided by Google and is built upon web technologies including HTML, CSS, and JavaScript. The core functionality of the browser extension lies in fetching the language schema and the document expressed in our transparency language from publicly available sources (i.e. Github or the data controller’s web server). Then, these contents are parsed and transformed into a summarizing textual and visual representation. For example, the stated data transfers are counted, the indication regarding automatic decision making is color-coded or the data protection officer is named. Third country transfers are represented using the official flags according to the country codes. The information is updated at each click on the extension’s icon to prevent unnecessary HTTP requests.

For future work, the browser extension should provide users with more diverse and dynamic representations according to their individual preferences and competencies. The inclusion of privacy icons, other visual-based representations or helping comments hold a lot of promise for adaptive representations [14, 15]. Moreover, there is a direct link to the *tilt-hub*, introduced in section 5.2. As the different deployment models of *tilt-hub* allow to provide transparency information independently from every single data controller and potentially even from a public repository, the browser extension could even display a summary of the transparency notice without the original publication by the vendor under study.

7 FUTURE WORK

The scientific discourse on privacy languages has been ongoing for several years already. Our language has a clear and in-depth focus on current transparency requirements from the GDPR. However, the language design and implementation process is meant to be adjusted to the technical and legal circumstances of tomorrow as well. Consequently, we see our language model as “living” specification that can and shall be extended by experts from various disciplines. With respect to traditional language evaluation criteria, we assume broad conformance [30]. However, to support the application in a diverse field of domains and for different target groups (from controller to data subject), we consider at least the following concepts worth investigating in future work. First, the automated extraction of transparency information from natural language or even source code (annotations) seems promising. In addition, there is a large body of literature on explainable AI which clearly relates to the information obligations with regards to automatic decision making [3]. Moreover, the integration and interplay with other privacy languages, e.g. in the context of legally sufficient consent [40], needs to be examined. Later on, advanced concepts from the GDPR (e.g. joint controllership) or recent judicial decisions (e.g. case law relating EU–US data transfers) need to be taken into consideration. Moreover, we also mentioned overlaps with accountability-related requirements for keeping records of processing activities from Art.

30 of the GDPR. These should also be investigated further in the future. Interpretative bodies (such as the EDPB) also guide us through their debates with regard to the aforementioned topics.

8 CONCLUSION

The current practice of providing privacy-related transparency information is dysfunctional and does not serve its originally intended goals anymore. To pave the way for novel, technically mediated approaches to transparency, we herein presented a structured, machine-readable transparency information language, accompanied by a surrounding toolkit that eases practical adoption. The language specification is based on an extensive analysis of transparency requirements from the GDPR and thus offers the expressiveness required to unfold actual legal relevance. The toolkit, in turn, particularly comprises two libraries for widely used programming languages as well as an easily instantiatable storage and management backend that significantly reduce practical implementation efforts for employing our language. Together, language and toolkit thus provide viable technical means for representing transparency information in a structured, machine-readable form as well as for managing and processing them automatically in real-world (web) information systems with reasonable implementation effort.

Altogether, our language and toolkit are powerful building blocks allowing for the development of a broad variety of novel technical means for materializing the privacy principle of transparency in a way that better aligns with the original goals of respective regulations than current, legalese policy documents do. The transparency analysis platform and the browser extension presented herein vividly demonstrate these capabilities and we envision a broad range of further possible applications and tools to become possible on top of an existing machine-readable representation of transparency information. Presentation was not the core subject of the paper, however, the language can be groundwork for a broad variance of user interfaces.

Finally, by providing our transparency information language and toolkit, we also change the given preconditions for applying regulatory provisions: Art. 25 of the GDPR obligates the implementation of technical measures for *all* privacy principles as soon as they are available as part of the current state of the art and can be implemented with reasonable effort. With the contributions presented herein, both factors come closer to the point where technical transparency mechanisms become obligatory. Insofar, our approach addresses the current structural weaknesses of transparency by taking an interdisciplinary point of view between law and technology and materializing the concepts into practical systems engineering. In so doing, we hope to re-amplify transparency that “might lead consumers to behave differently” and, thus, to contribute to an upcoming “renaissance” in privacy law [21].

ACKNOWLEDGMENTS

The work behind this paper was partially conducted within the project DaSKITA, supported under grant no. 28V2307A19 by funds of the Federal Ministry of Justice and Consumer Protection (BMJV) based on a decision of the Parliament of the Federal Republic of Germany via the Federal Office for Agriculture and Food (BLE) under the innovation support programme.

REFERENCES

- [1] Julio Angulo, Simone Fischer-Hübner, Tobias Pulls, and Erik Wästlund. 2015. Usable Transparency with the Data Track: A Tool for Visualizing Data Disclosures. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1803–1808. <https://doi.org/10.1145/2702613.2732701>
- [2] Article 29 Data Protection Working Party. 2017. *Guidelines on Transparency under Regulation 2016/679*. Technical Report. Directive 95/46/EC of the European Parliament.
- [3] Umang Bhatt, Alice Xiang, Shubham Sharma, Adrian Weller, Ankur Taly, Yunhan Jia, Joydeep Ghosh, Ruchir Puri, José M. F. Moura, and Peter Eckersley. 2020. Explainable Machine Learning in Deployment. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 648–657. <https://doi.org/10.1145/3351095.3375624>
- [4] Christoph Bier, Kay Kühne, and Jürgen Beyerer. 2016. PrivacyInsight: The Next Generation Privacy Dashboard. In *Privacy Technologies and Policy*, Stefan Schiffler, Jetzabel Serna, Demosthenes Ikononou, and Kai Rannenberg (Eds.). Springer International Publishing, Cham, 135–152.
- [5] Cheng Chang, Huaxin Li, Yichi Zhang, Suguo Du, Hui Cao, and Haojin Zhu. 2019. Automated and Personalized Privacy Policy Extraction Under GDPR Consideration. In *Wireless Algorithms, Systems, and Applications*, Edoardo S. Biagioni, Yao Zheng, and Siyao Cheng (Eds.). Springer International Publishing, Cham, 43–54.
- [6] Elisa Costante, Jerry den Hartog, and Milan Petković. 2013. What Websites Know About You. In *Data Privacy Management and Autonomous Spontaneous Security*, Roberto Di Pietro, Javier Herranz, Ernesto Damiani, and Radu State (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 146–159.
- [7] Lorrie Faith Cranor. 2003. P3P: Making privacy policies more useful. *IEEE Security & Privacy* 1 (2003), 50–55. Issue 6. <https://doi.org/10.1109/MSECP.2003.1253568>
- [8] European Parliament & Council. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation).
- [9] Keishiro Fukushima, Toru Nakamura, Daisuke Ikeda, and Shinsaku Kiyomoto. 2018. Challenges in Classifying Privacy Policies by Machine Learning with Word-Based Features. In *Proceedings of the 2nd International Conference on Cryptography, Security and Privacy (ICCSPP 2018)*. Association for Computing Machinery, New York, NY, USA, 62–66. <https://doi.org/10.1145/3199478.3199486>
- [10] Armin Gerl and Bianca Meier. 2019. The Layered Privacy Language Art. 12–14 GDPR Extension–Privacy Enhancing User Interfaces. *Datenschutz und Datensicherheit-DuD* 43, 12 (2019), 747–752.
- [11] Hamza Harkous, Kassem Fawaz, Rémi Lebre, Florian Schaub, Kang G. Shin, and Karl Aberer. 2018. Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning. In *27th USENIX Security Symposium (USENIX Security 18)*. USENIX Association, Baltimore, MD, 531–548.
- [12] Hamza Harkous, Kassem Fawaz, Kang G. Shin, and Karl Aberer. 2016. PriBots: Conversational Privacy with Chatbots. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO, 6.
- [13] Hans Hedbom. 2009. A Survey on Transparency Tools for Enhancing Privacy. In *The Future of Identity in the Information Society*, Vashek Matyáš, Simonen Fischer-Hübner, Daniel Cvrček, and Petr Švenda (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 67–82.
- [14] Leif-Erik Holtz, Katharina Nocun, and Marit Hansen. 2011. Towards Displaying Privacy Information with Icons. In *Privacy and Identity Management for Life*, Simone Fischer-Hübner, Penny Duquenoy, Marit Hansen, Ronald Leenes, and Ge Zhang (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 338–348.
- [15] Leif-Erik Holtz, Harald Zwingelberg, and Marit Hansen. 2011. *Privacy Policy Icons*. Springer Berlin Heidelberg, Berlin, Heidelberg, 279–285 pages. https://doi.org/10.1007/978-3-642-20317-6_15
- [16] Interactive Advertising Bureau (IAB). 2020. TCF – Transparency & Consent Framework. <https://iabeurope.eu/transparency-consent-framework/>
- [17] ISO IEC 14977 996 E 1996. *EBNF Syntax Specification*. Standard. International Organization for Standardization, Geneva, CH.
- [18] Milena Janic, Jan Pieter Wijbenga, and Thijs Veugen. 2013. Transparency enhancing tools (TETs): An overview. In *Proceedings of the Workshop on Socio-Technical Aspects in Security and Trust, STAST*. IEEE, New Orleans, LA, USA, 18–25. <https://doi.org/10.1109/STAST.2013.11>
- [19] JSON.org. 2020. ECMA-404 The JSON Data Interchange Standard. <https://www.json.org/json-en.html>
- [20] Ioannis Kakavas. 2016. Creepy. A geolocation OSINT tool. <https://www.geocreepy.com/>
- [21] Margot Kaminski. 2020. Law and Technology. A recent renaissance in privacy law. *Commun. ACM* 63, 9 (2020), 24–27. <https://doi.org/10.1145/3411049>
- [22] Farzaneh Karegar, Tobias Pulls, and Simone Fischer-Hübner. 2016. *Visualizing Exports of Personal Data by Exercising the Right of Data Portability in the Data Track - Are People Ready for This?* Springer International Publishing, Cham, 164–181. https://doi.org/10.1007/978-3-319-55783-0_12
- [23] Sabrina Kirrane, Javier D. Fernández, Piero Bonatti, Iliana Mineva Petrova, Luigi Sauro, and Eva Schlehahn. 2019. The SPECIAL Usage Policy Language. <https://ai.wu.ac.at/policies/policylanguage/>
- [24] David Knoke and Song Yang. 2019. *Social network analysis*. Vol. 154. Sage Publications.
- [25] Thomas Linden, Rishabh Khandelwal, Hamza Harkous, and Kassem Fawaz. 2018. The Privacy Policy Landscape After the GDPR. <http://arxiv.org/abs/1809.08396>
- [26] Aleecia M. McDonald and Lorrie Faith Cranor. 2008. The Cost of Reading Privacy Policies. *Journal of Law and Policy for the Information Society* 4 (2008), 543–568.
- [27] Marjan Mernik, Jan Heering, and Anthony M. Sloane. 2005. When and How to Develop Domain-Specific Languages. *ACM Comput. Surv.* 37, 4 (Dec. 2005), 316–344. <https://doi.org/10.1145/1118890.1118892>
- [28] MongoDB, Inc. 2020. MongoDB Documentation. Data Models. Schema Validation. <https://docs.mongodb.com/manual/core/schema-validation/>
- [29] Jonathan A. Obar and Anne Oeldorf-Hirsch. 2018. The biggest lie on the Internet: ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society* 23, 1 (2018), 128–147. <https://doi.org/10.1080/1369118X.2018.1486870>
- [30] Richard F. Paige, Jonathan S. Ostroff, and Phillip J. Brooke. 2000. Principles for modeling language design. *Information and Software Technology* 42, 10 (2000), 665–675.
- [31] Frank Pallas, Max-R. Ulbricht, Stefan Tai, Thomas Peikert, Marcel Reppenhagen, Daniel Wenzel, Paul Wille, and Karl Wolf. 2020. Towards Application-Layer Purpose-Based Access Control. In *Proceedings of the 35th Annual ACM Symposium on Applied Computing (SAC '20)*. Association for Computing Machinery, New York, NY, USA, 1288–1296. <https://doi.org/10.1145/3341105.3375764>
- [32] Felipe Pezoa, Juan L. Reutter, Fernando Suarez, Martin Ugarte, and Domagoj Vrgoc. 2016. Foundations of JSON Schema. In *Proceedings of the 25th International Conference on World Wide Web (WWW '16)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 263–273. <https://doi.org/10.1145/2872427.2883029>
- [33] Princiya. 2019. Mozilla Firefox Lightbeam. <https://addons.mozilla.org/firefox/addon/lightbeam-3-0/>
- [34] Philip Raschke, Axel Küpper, Olha Drozd, and Sabrina Kirrane. 2018. Designing a GDPR-Compliant and Usable Privacy Dashboard. In *Privacy and Identity Management. The Smart Revolution*, Marit Hansen, Eleni Kosta, Igor Nai-Fovino, and Simone Fischer-Hübner (Eds.). Springer International Publishing, Cham, 221–236. https://doi.org/10.1007/978-3-319-92925-5_14
- [35] Joel R. Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T. Graves, Fei Liu, Aleecia McDonald, Thomas B. Norton, and Rohan Ramanath. 2015. Disagreeable Privacy Policies: Mismatches between Meaning and Users' Understanding. *Berkeley Technology Law Journal* 30 (2015), 39.
- [36] Arianna Rossi and Monica Palmirani. 2020. Can Visual Design Provide Legal Transparency? The Challenges for Successful Implementation of Icons for Data Protection. *Design Issues* 36, 3 (2020), 82–96.
- [37] Norman Sadeh, Ro Acquisti, Travis D. Breaux, Lorrie Faith Cranor, Aleecia M. McDonald, Joel R. Reidenberg, Noah A. Smith, Fei Liu, N. Cameron Russell, Florian Schaub, and Shomir Wilson. 2013. The Usable Privacy Policy Project: Combining Crowdsourcing, Machine Learning and Natural Language Processing to Semi-Automatically Answer Those Privacy Questions Users Care About. <http://ra.adm.cs.cmu.edu/anon/usr0/ftp/home/anon/isr2013/CMU-ISR-13-119.pdf>
- [38] The State of California. 2018. California Consumer Privacy Act of 2018. http://leginfo.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5
- [39] Sabine Trepte, Doris Teutsch, Philipp K. Masur, Carolin Eicher, Mona Fischer, Alisa Hennhöfer, and Fabienne Lind. 2015. *Do People Know About Privacy and Data Protection Strategies? Towards the "Online Privacy Literacy Scale" (OPLIS)*. Springer Netherlands, Dordrecht, 333–365. <https://doi.org/10.1007/978-94-017-9385-8>
- [40] Max-R. Ulbricht and Frank Pallas. 2018. YaPPL - A Lightweight Privacy Preference Language for Legally Sufficient and Automated Consent Provision in IoT Scenarios. In *Proceedings of Data Privacy Management 2018 (LNCS)*, Giovanni Livraga and Ruben Rios (Eds.), Vol. 11025. Springer International Publishing, 329–344.
- [41] Łukasz Warchał. 2012. Using Neo4j graph database in social network analysis. *Studia Informatica* 33, 2A (2012), 271–279.
- [42] Jonathan Wetherbee, Massimo Nardone, Chirag Rathod, and Raghu Kodali. 2018. *EJB, Web Services, and Microservices*. Apress, Berkeley, CA, 265–317. https://doi.org/10.1007/978-1-4842-3573-7_6
- [43] Zhi Xu and Sencun Zhu. 2015. SemaDroid: A Privacy-Aware Sensor Management Framework for Smartphones. In *Proceedings of the 5th ACM Conference on Data and Application Security and Privacy (CODASPY '15)*. Association for Computing Machinery, New York, NY, USA, 61–72. <https://doi.org/10.1145/2699026.2699114>
- [44] Christian Zimmermann. 2015. A Categorization of Transparency-Enhancing Technologies. arXiv. <https://arxiv.org/abs/1507.04914>