

# A Typology of Explanations for Explainability-by-Design

NIKO TSAKALAKIS and SOPHIE STALLA-BOURDILLON, Vrije Universiteit Brussel, Belgium

TRUNG DONG HUYNH, King's College London, United Kingdom

LUC MOREAU, University of Sussex, United Kingdom

As automated decision-making permeates almost all aspects of everyday life, capabilities to generate meaningful explanations for various stakeholders (i.e., decision-makers, addressees of decisions including individuals, auditors, and regulators) should be carefully deployed. This paper presents a typology of explanations intended to support the first pillar of an explainability-by-design strategy. Its production has been achieved by pursuing a responsible innovation approach and introducing a new persona within the research and innovation process, i.e., a legal engineer, whose role is to work at the interface of two teams, the compliance and the engineering teams and to oversee the process of requirement elicitation, which is often opinionated and narrowing. Once explanation requirements have been derived from applicable regulatory requirements, compliance rules or business policies, they have been mapped to the dimensions of the typology to produce fine-grained explanation requirements, forming computable building blocks that can then be translated into system requirements during the technical design phase. The typology has been co-created with industry partners operating in two sectors: finance and education. Two pilot studies have thus been conducted to test both the feasibility of the generation and computation of explanations on the basis of the typology and the usefulness of the outputs in the light of the state of the art. The typology comprises nine hierarchical dimensions. It can be leveraged to operate a stand-alone classifier of explanations that acts as detective controls within a broader partially-automated compliance strategy. A machine-readable format of the typology is provided in the form of a light ontology.

CCS Concepts: • **Social and professional topics** → **Privacy policies**; • **Computing methodologies** → Natural language generation; **Artificial intelligence**; • **Human-centered computing** → **Natural language interfaces**; *Interaction techniques*.

Additional Key Words and Phrases: artificial intelligence, explainability, typology, data protection, automated decisions

## 1 INTRODUCTION

Automated decisions and the ways in which they can be interpreted by and explained to humans have been at the forefront of Artificial Intelligence (AI) research for a while now.<sup>1</sup> Explainability has been described as being a powerful tool for a variety of purposes including “detecting flaws in the model and biases in the data, for verifying predictions, for improving models, and finally for gaining new insights into the problem at hand” [49]. However, there is no consensus on the metrics for evaluating explanation methods, and it is hard to compare them [61]. Kim and Routledge argue that explainable AI cannot limit itself to assessing the simplicity and locality of an explanation and social, ethical, and epistemic values should be taken into account to assert the relevance of a range of explanation types and stakeholders in explaining AI outputs [29]. More recently, it has been argued that “without a proper theory of idealization, it remains difficult to thoroughly diagnose the success of xAI methods” [54].

---

<sup>1</sup>See e.g. for the evolution of interpretable tools and the emergence of Explainable AI (XAI) [3, 31, 47].

---

Authors' addresses: Niko Tsakalakis, nk.ts@icloud.com; Sophie Stalla-Bourdillon, Vrije Universiteit Brussel, Brussels, Belgium, sophie.stalla-bourdillon@vub.be; Trung Dong Huynh, King's College London, London, United Kingdom, dong.huynh@kcl.ac.uk; Luc Moreau, University of Sussex, Falmer, United Kingdom, luc.moreau@sussex.ac.uk.



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

Understanding how and why automated decision-making systems process personal data of individuals became particularly important with the introduction of the General Data Protection Regulation (GDPR) [19] in the European Union (EU). Its Article 22 on ‘Automated individual decision-making, including profiling’ gave rise to a debate about the existence of a right to explanation.<sup>2</sup> National data protection authorities’ and courts’ decisions have confirmed the existence of a right to explanation, although its contours are not clearly delimited yet [35, 56]. In addition, the Court of Justice of the European Union (CJEU) has adopted a broad definition of automated individual decision-making, which includes “the automated establishment, by a credit information agency, of a probability value based on personal data relating to a person and concerning his or her ability to meet payment commitments in the future (...), where a third party, to which that probability value is transmitted, draws strongly on that probability value to establish, implement or terminate a contractual relationship with that person.”<sup>3</sup> Regardless of the contours and source of such a right,<sup>4</sup> the debate provided an opportunity to shift the focus from the inner workings of the ‘black box’, which has been the primary focus of XAI, to its broader environment and purpose, i.e., the production of information meaningful for the exercise of individual rights [13, 50]. Recent discussions on explanations have highlighted the limits of existing approaches for generating meaningful explanations, e.g. the lack of focus upon user needs [59], the difference between explanations that offer a justification (they respond to the ‘why’ question) or produce transparency (they respond to the ‘how’ question) [46], with some indication that end users tend to pay more attention to the ‘why’ question [10], as confirmed by an overview of research on explanations in philosophy, cognitive science, and the social sciences. Mittelstadt et al. thus argue that explanations should be “contrastive, selective, and socially interactive” [38]. At the same time, judges are increasingly stressing the importance of contextualising the personal data processed to ensure that the information received by data subjects is intelligible.<sup>5</sup>

Several commentators focus upon specific groups of consumers of explanations, i.e., data subjects [32], developers [62] without directly addressing differing levels of understanding. Often, proposed solutions have attempted to explain models while ignoring the impact of factors that are external to black boxes [13, 17], or have not taken into account events occurring before the training stage or during model deployment, with a few exceptions like [12, 24]. One type of explanations in particular, i.e., counterfactual explanations [58], has been gaining traction in practice despite its inherent limits and in particular its inability to address the ‘how’ question.<sup>6</sup>

The need to develop a systematic approach to explanations for GDPR compliance has thus been discussed extensively in the literature as a high-level compliance goal<sup>7</sup> or strategy.<sup>8</sup> With this said, rights to explanations are also present in sector-specific regulatory frameworks, such as education and in particular access to schools, as explained in Section 4.2, and the financial sectors, as mentioned in Section 4.1. As regards the financial sector, it is worth mentioning that a new Directive on credit agreements for consumers has been adopted by the EU in 2023, which builds on the GDPR and further specifies the rights to information and explanation in the context of covered credit agreements [1].<sup>9</sup> National

<sup>2</sup>See e.g. [17, 50, 57].

<sup>3</sup>CJEU C-634/21 OQ v Land Hessen 7 December 2023 ECLI:EU:C:2023:957, para. 73

<sup>4</sup>Opinion of Advocate General Richard de la Tour C-203/22 CK 12 September 2024 ECLI:EU:C:2024:745, para. 67 (“In short, ‘meaningful information’, as required by Article 15(1)(h) of the GDPR, must not only be clear and accessible, but must also be accompanied by explanations to ensure that it is properly understood”).

<sup>5</sup>CJEU Case C-487/21 FF v Österreichische Datenschutzbehörde and CRIF GmbH, 4 May 2023 ECLI:EU:C:2023:369, para. 41), and CJEU C-307/22 FT v DW 26 October 2023 ECLI:EU:C:2023:811, para. 74.

<sup>6</sup>By way of example, counterfactual explanations are not necessarily helpful in providing useful information to challenge a financial decision, such as a rejection of a loan application. See in general [5].

<sup>7</sup>See, for example, [30, 33].

<sup>8</sup>See, for example, [15, 41].

<sup>9</sup>Article 18(8)(a), in particular, provides that “[w]here the creditworthiness assessment involves the use of automated processing of personal data, Member States shall ensure that the consumer has the right to request and obtain from the creditor human intervention, consisting of the right to request

transpositions will be applicable from 20 November 2026. Horizontal frameworks, such as the EU AI Act [20] can also introduce bespoke rights to explanations, as discussed in Section 5.2. There is now a pressing need to better understand what explanation outputs could look like in practice to advance the state of the art.

Analysing the characteristics of explanations through the prisms of philosophy and data science, Pieters [46] distinguishes between explanations-for-trust and explanations-for-confidence. Essentially, explanations sit on a spectrum. Explanations with very little detail fail, as they do not contain enough information to build recipients' confidence or trust. Explanations that contain too much detail fail as well because they are misaligned with the recipient's level of understanding. Explanations that contain a low level of detail can help build confidence by answering the 'why' question, and explanations that contain a high level of detail can help build trust by answering the 'how' question. The 'how' question is of interest for assessing data quality, e.g. to help users understand how their data contributed to algorithmic intelligence [11]. Pieters also notes that different explanations can be generated to answer the same question [46].

Arguing for a right to 'meaningful and intelligible explanation', Kim and Routledge [29] distinguish between *ex ante* generic explanations to achieve informed consent and *ex post* generic and specific explanations either to inform remedial actions or to produce individualised information about the algorithmic processing. Castets-Renard [8], categorising the impact levels of predictive policing decisions into four classes,<sup>10</sup> distinguishes explanation types by triggering event: for decisions with no impact on the individual, a publicly available (generic) explanation should suffice; for decision with a moderate impact, explanations should be produced upon a request from the user; for decisions with a high or very high impact, explanations should be systematically produced at the point of disclosure of the decision.

The GDPR, and the introduction of data protection by design as a mandatory requirement,<sup>11</sup> has fed a formalisation effort and attempts to make GDPR obligations machine readable, to produce GDPR-specific description languages like Fides,<sup>12</sup> to create a GDPR vocabulary for use in ontologies [6], or to extend a Data Privacy Vocabulary with GPDR concepts.<sup>13</sup> However, these methods have not been applied to the generation or computation of explanations yet. Pieters [46], nonetheless, acknowledges the relevance of explanations programs.

*“Actants can thus be said to have an explanation program, i.e. their action program projected on the domain of explanation. When actants are asked to explain something about a theory or system, they have certain intentions and possibilities for explaining in a certain way.”*

He also observes that responsibilities for explanations can be delegated, either to humans or machines. When responsibilities for explanations are delegated to machines, however, explanation-for-trust and explanation-for-confidence are then needed to backup computed explanations.

The purpose of this paper is to lay the foundations for building partially-automated explanation programs following a responsible innovation approach. Innovating responsibly implies a “continuous commitment to be anticipatory, reflective, inclusively deliberative, and responsive” [42], which can be pursued through a variety of practices, like socio-technical integration research practices.<sup>14</sup> This paper thus intentionally departs from approaches that would leverage imprecise, probabilistic techniques, such as generative AI, to produce explanations [44] and acknowledges

---

and obtain from the creditor a clear and comprehensible explanation of the assessment of creditworthiness, including on the logic and risks involved in the automated processing of personal data as well as its significance and effects on the decision”.

<sup>10</sup>Ranging from Level I 'no impact' to Level IV 'very high impact' according to the Impact Assessment Levels of the Canadian Directive on Automated Decision Making.

<sup>11</sup>GDPR, Article 25.

<sup>12</sup><https://ethyca.com/fides/>

<sup>13</sup><http://www.w3id.org/dpv/dpv-gdpr#>

<sup>14</sup>For an example see <https://cta-toolbox.nl/tools/STIR/>

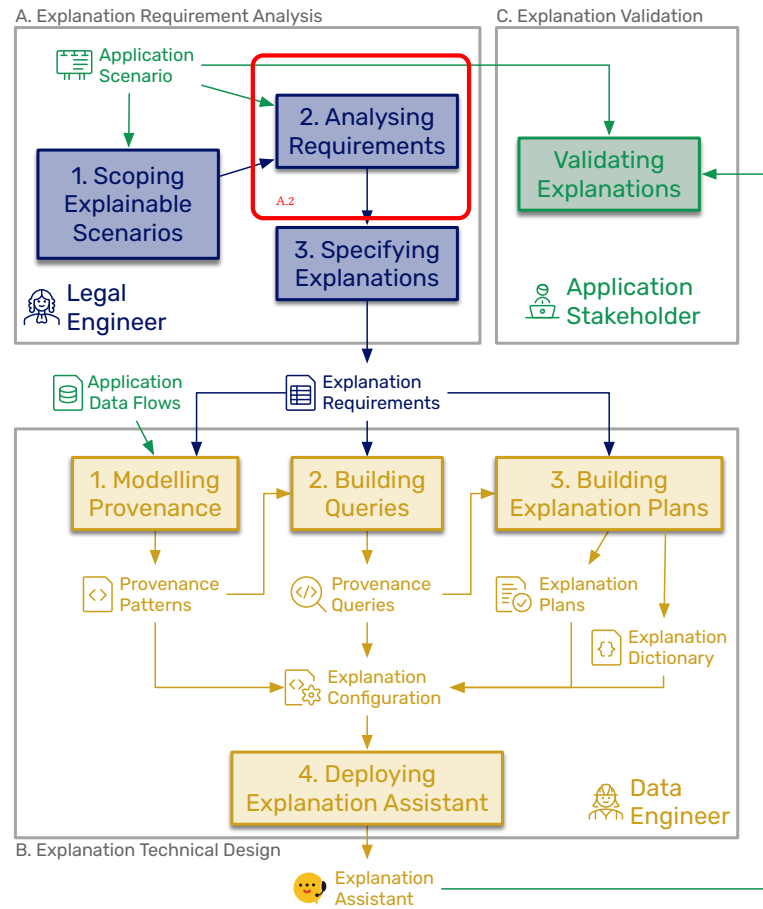


Fig. 1. Example of an Explainability-by-Design methodology

the inherent limits of any formalisation effort when legal requirements are at stake. In fact, computed explanatory statements are conceived as a means to facilitate the task of the entity accountable for the decision or responsible for explaining the decision: they are thus merely supporting statements, which might need to be refined and completed iteratively to meet the needs of a particular addressee of the decision, such as a data subject.

The first question to explore when building an automated explanation program relates to the building of a methodology for generating a comprehensive set of computable explanations at each relevant node of a data processing pipeline. In other words, we need a methodology for pursuing an explainability-by-design strategy, conceived as a form of legal protection by design [22], which should both preserve and strengthen contestability, i.e. the ability to contest not only the output of the data processing pipeline but also the very explanations produced to account for the functioning of the system [23].

An explainability-by-design methodology should follow a socio-technical process whereby explanation requirements are used to produce explanation statements that are then tested and validated with stakeholders before being translated

into system requirements. This led us to revisit socio-technical integration research practices, which consist in embedding social scientists into engineering team. We thus decided to add a new role, that of the legal engineer [21], whose function is to bridge between the compliance team and the engineering team. Indeed, not only is embedding social scientists within engineering teams important to identify social values midstream while engineers make decisions (e.g., transparency, understandability), but translating high-level legal and compliance requirements into fine-grained explanation requirements before generating system requirements is a fundamental intermediate step as well.

As alluded above, explanations can have a variety of goals, including justification and transparency. In this paper, we conceive an explanation as a statement that provides the details or reasons relating to an output or decision. A decision is thus defined as any output of a data processing pipeline, produced after consideration of one or more data points. The output or decision can be followed by an action to bring about a correction or a resolution to a situation.

It is possible to distinguish three phases within an explainability-by-design methodology [55], as illustrated in Fig. 1:

- (1) An *Explanation Requirement Analysis* phase, depicted inside rectangle A, where legal engineers determine the explanation requirements of the processing in question;
- (2) An *Explanation Technical Design* phase, depicted inside rectangle B, consisting of a set of technological steps, initiated by data engineers, to instantiate a service, the ‘Explanation Assistant’, that generates explanations according to phase 1; and
- (3) A *Validation* phase, depicted inside rectangle C, where the generated explanations are evaluated by the stakeholders for their suitability and effectiveness.

This paper focuses on the *Explanation Requirement Analysis* phase depicted in Fig. 1 and specifically on step A2 of Fig. 1 ‘*Analysing Requirements*.’ It broadens the approach to explainability by focusing upon automated decision-making pipelines instead of being purely model-centric [52] and is therefore of relevance in cases in which decision-making pipelines are only partially automated, like in the school allocation use case described below. The paper thus presents a typology developed to enable practitioners to systematically generate explanations of decisions through the formulation of explanation requirements, with the aim of putting the explanation recipient at the centre of the design in a ‘reflective sociotechnical approach’ [18]. The typology does not represent an exhaustive set of categories and, therefore, remains open. It has been built out of a minimum set of distinct dimensions which reflect different concerns, e.g., the timing, form, content, and effect of explanations, that have been evaluated in the context of small-scale pilot studies. These dimensions are non-orthogonal in the sense that the selection of a particular value within one dimension may determine the selection of a particular value within another dimension. This typology is used to produce explanation requirements that are then mapped to relevant data points within an audit trail so that it can be consumed by an explanation assistant. The typology is thus associated with categorisation rules that take as input applicable explanation requirements and output explanation specifications. The typology can be represented in machine-readable formats, such as a light ontology, which is included in the appendix of this paper.

The paper is organised into five sections. Following the introduction in Section 1, Section 2 presents the methodology that was followed to develop the typology of explanations. Section 3 describes each dimension of the typology and associated criteria. Section 4 illustrates its application in the context of small-scale pilot studies comprising two use cases. Section 5 discusses evaluation and Section 6 the limitations of the methodology and typology. Finally, Section 7 highlights the potential of the typology as a classifier of explanations and discusses its limitations.

Methodology Components	Application
1. Meta-Characteristics	<p><b>Bases:</b> Information to meet the explanation requirements arising out of applicable regulatory obligations, business needs or other policies</p> <p><b>Expected User:</b> Organisations determining their processing methods / System vendors designing processing systems for third-parties</p> <p><b>Expected Use:</b> The proactive production of explanations as part of a supportive automated compliance strategy and the design of 'Explainable by Design' systems</p> <p><b>Purpose:</b> The categorisation of information to produce comprehensive explanations for the applicable explanation requirements of the Organisation</p>
2. Ending Conditions	<p><b>Objective Conditions:</b></p> <ol style="list-style-type: none"> <li>2.1. A representative sample of objects has been examined</li> <li>2.2. No new dimensions or characteristics were added in the last iteration</li> <li>2.3. No dimensions or characteristics were merged or split in the last iteration</li> <li>2.4. Every dimension is unique and not repeated (i.e., there is no dimension duplication)</li> <li>2.5. Every characteristic is unique within its dimension (i.e., there is no characteristic duplication within a dimension)</li> </ol> <p><b>Subjective Conditions:</b> Conciseness, Robustness, Comprehensiveness, Extensibility, and Explainability</p>
3. Empirical to Conceptual Approach	<ol style="list-style-type: none"> <li>3.1. The sampling of two practical use cases of (semi-)automated decision making (loan application and school allocations) and their applicable rules</li> <li>3.2. The analysis and coding of 'bases' from the applicable rules using desk research</li> <li>3.3. The categorisation of the 'bases' into characteristics and dimensions</li> </ol>
4. Conceptual to Empirical Approach	<ol style="list-style-type: none"> <li>4.1. The identification of explainability characteristics in the literature</li> <li>4.2. Comparing the explainability characteristics to own analysis</li> <li>4.3. The reorganisation of the typology taking into account the findings of step 4.2.</li> </ol>

Table 1. The Nickerson et al. [40]'s methodology as applied.

## 2 METHODOLOGY

The explanation typology is the product of desk research and co-creation research techniques implemented to generate requirements from interactions with domain experts representing industry partners. The explanation typology is the work-product of a legal engineer who has been embedded within the engineering team after having worked closely with compliance teams of industry partners to translate high-level legal and compliance requirements into fine-grained explanation requirements.

Overall, the typology was built using the method by Nickerson et al. [40]. First, the goals of the typology were set, as depicted in Table 1 '1. Meta-Characteristics'. The goal of the typology is to map the information needed to meet explanation requirements derived from an analysis of applicable legal or regulatory obligations, compliance rules or business policies.

Following an empirical-to-conceptual approach, an experimentation was performed in the context of two use cases. Pilot studies with two groups of stakeholders participating to two different use cases were conducted to generate a list of relevant explanation requirements in these sectors. The first use case concerned an automated credit scoring process leveraged to assess loan applications in the financial sector. The second use case involved a semi-automated process for allocating school places to UK pupils during annual school admission rounds. Each use case involved producers and consumers of explanations, i.e., decision makers from institutions that deploy automated decision pipelines and who are under an obligation to explain the output to end-users.

The pilot studies were conducted as a proof of concept to demonstrate the feasibility and usefulness of categorising explanation requirements and transforming the output of the process into a machine-readable classifier. Domain experts representing industry partners participated to the construction of explanation requirements together with the research

team as well as to the production of explanation outputs, i.e. explanations. Interviews were run with domain experts and outputs were discussed in the context of focus groups. Interview responses were analysed to reveal common themes. The themes, or ‘bases’, of the explanation requirements were grouped under different headings which were generalised to form dimensions of concern. Subsequently, in a conceptual-to-empirical approach, the categories and dimensions produced in the experimentation stage were compared to explanation properties identified in the literature. The typology was refined to ensure alignment with desk research findings and to make it more concise.

The pilot studies confirmed the feasibility and usefulness of computing explanations on the basis of the typology, which was leveraged to produce a light ontology and an explanation assistant. Given the limited scope of the pilot studies, we do not claim that the typology amounts to an exhaustive taxonomy. We however assessed both the methodology and the typology in the light of the newly adopted EU AI Act: we found alignment between the methodology, the typology and the data governance requirements for high-risk AI systems<sup>15</sup> and the newly-introduced right to explanation.<sup>16</sup>

Several iterations were performed, during which the research team progressively added new dimensions until both objective and subjective ending conditions were finally met for the particular use cases. The objective conditions were fulfilled when new iterations no longer split or merged dimensions or categories; when each dimension was unique and did not repeat; and, when each category was unique within a specific dimension. The subjective criteria were met by the last iteration. The typology dimensions were within the  $7 \pm 2$  limit, set by Miller [36], which achieved conciseness. Each category within each dimension was reduced to an elementary level to ensure differentiation of the object under investigation. The comprehensiveness (or relative completeness) of the typology, which implies that all dimensions are included and that new objects can be categorised, was tested through the production of examples based on explanation requirements derived from the GDPR and the School Admissions Code. The typology has been designed to be extensible and cover requirements stemming from a variety of frameworks. Finally, the explainability criterion, which evaluates how well the dimensions explain the object of the typology, was addressed in the context of focus groups with domain experts.

### 3 BUILDING A TYPOLOGY FOR EXPLANATIONS

The overall goal of the explanation typology is to support a categorisation process that is able to produce a comprehensive set of explanation requirements from which it is then possible to identify building blocks for the computation stage.

The sources of explanation requirements can vary. They can either be mandated by law (RQ1); or be grounded upon mere internal requirements, such as company policies (RQ2). Co-creation partners stressed the necessity to deploy explanations in different contexts (RQ3), to be able to scale and reuse explanation outputs easily (RQ4), and to be able to create them proactively before receiving requests from data subjects or addressees (RQ5). Issuers and recipients should also be able to retrieve explanations quickly (RQ6), explanations should empower (RQ7) and educate their recipients (RQ8), and ultimately increase the level of engagement with the service (RQ9). Some partners viewed explanations as a tool to automate their transparency obligations (RQ10), even though they had to strike a trade-off between transparency and the protection of intellectual property including trade secrets (RQ11). Finally, explanations should be fit for purpose and to the point so as not to overwhelm their audience (RQ12).

All these considerations formed our first set of requirements. These requirements were then conflated and codified into ‘bases’ to form typology dimensions. Each dimension describes a facet of an explanation. When needed, dimensions

---

<sup>15</sup>AI Act, Article 10

<sup>16</sup>AI Act, Article 86

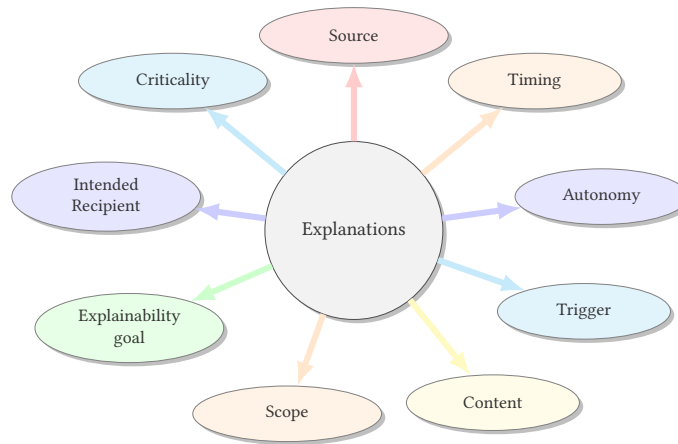


Fig. 2. Dimensions of explanations.

are broken down into sub-categories to provide sufficient granularity to the description. Each dimension is hierarchical. The depth of each dimension can vary.

Fig. 2 shows the nine dimensions.

**Source** defines the origin of the explanation requirement (RQ1).

**Timing** describes the moment at which the explanation is generated in relation to its object (RQ10).

**Autonomy** distinguishes between explanations that are generated without input and explanations that require an input for their generation (RQ5).

**Trigger** defines the events that trigger the generation of an explanation (RQ2,5).

**Content** captures the details about the formulation of an explanation (RQ6,9,11).

**Scope** assesses whether an explanation addresses a limited set of circumstances or has extended applicability (RQ4).

**Explainability goal** defines the functioning purpose of the generated explanation (RQ3,7,8).

**Intended Recipient** defines the target audiences of the generated explanations (RQ3).

**Criticality** specifies whether the explanation is mandatory or recommended (RQ2).

### 3.1 Source

The Source defines the origin of the explanation requirement (Fig. 3). An explanation requirement may stem from three types of sources: (i) applicable laws; (ii) related authoritative guidance and standards; and, (iii) internal compliance or business needs.




Fig. 3. The Source dimension



When an explanation requirement is derived from the law, be it a piece of legislation or case law (i.e. primary sources), they are considered to be primary requirements. Such obligations can arise from a multitude of areas, e.g. public administrative law, consumer law, data protection law. Non-compliance with explicit explanation requirements can result in severe consequences for the organisation, such as reputational damage, fines and litigation costs. Therefore, organisations should have strong incentives to achieve compliance with explicit primary explanation requirements. In other circumstances, explanation requirements are merely implicit. This may be the case when compliance is de facto facilitated by the generation of an explanation. In other words, the generation of an explanation is characterised as good practice for achieving compliance.

#### Primary explanation requirements

**Explicit:** School admission authorities in England must explain the reasons behind the refusal of a school place to a child when informing the parent (applicant) of the decision, under paragraph 2.32 of [16].

**Implicit:** Data controllers, who must be able to demonstrate compliance with the GDPR under Article 5(2), are under an implicit regulatory requirement to detect whether the processing is performed according to internal policies, which may be achieved through the generation of explanations.

Because primary requirements for explanations come from legal rules, which are often high-level and open-ended, the entities that are called to comply with explanation requirements are often left with a margin of appreciation to determine how best to comply. The translation of high-level rules (e.g., “data subjects should be informed about the logic involved in the processing of their personal data”) into secondary requirements (e.g., “data subjects should be informed about whether a particular data point has contributed to the automated decision”) imply that a decision has been made to interpret the high-level requirement into context. Such translation can be facilitated by the release of guidance and standards, such as by national supervisory authorities and (supra-)national standard bodies.

#### Secondary explanation requirements

The Information Commissioner’s Office, the national supervisory authority for data protection in the UK, has released authoritative guidance on explaining AI decisions in practice in collaboration with the Alan Turing Institute. The Information Commissioner’s Office [27] provides different tasks that organisations need to consider when deciding how to explain their AI systems.

In practice, organisations tend to establish internal compliance functions. The complexity of these functions usually depends on the sector in which the organisation operates and its size. Internal compliance policies consist in translating primary and secondary requirements into actionable and auditable tertiary requirements. Such tertiary requirements are usually more granular than primary and secondary requirements.

#### Tertiary explanation requirements

Internal requirements for explanations may arise out of designated roles and responsibilities (e.g. responsible agents), internal guidelines on the legalities of explanations created by in-house legal teams, standardised processes (e.g. quality control, email templates for correspondence with recipients), monitoring mechanisms and delivery systems for explanations (e.g. technical infrastructure).

### 3.2 Timing

The Timing dimension expresses the moment at which an explanation is generated in relation to the object explained by the explanation (Fig. 4). An explanation can be generated either *ex ante*, i.e. before the object(s) explained within the explanation take(s) place, or *ex post*, i.e. after the object(s) materialise(s).

**Ex ante & ex post explanations**

For example, an explanation about the details of the processing of personal data will be triggered *ex ante*, before the processing takes place, and will explain the intended processing that *will* take place. On the contrary, explanations about a decision will be triggered *ex post*, after the decision has been reached and will explain the decision.

**3.3 Autonomy**

The Autonomy dimension distinguishes between explanations that are generated without any input from the recipient or explanations that are generated only as a response to specific input from the recipient. It is depicted in Fig. 5. A Reactive explanation is responsive. Reactive explanations will typically be triggered by a direct request from an individual or organisation for information about a process, event or decision.

**Reactive explanations**

An individual receives an unexpected or undesired decision (e.g. a lower credit score than expected) and, therefore, decides to contact customer support services for more information. At the point of contact (e.g. on receipt of the call or email), the customer support agent has to provide the reasons for this particular decision to the individual. The explanation is therefore triggered when the individual (subject to the decision) contacts the organisation (who made the decision) to obtain an explanation.

On the other hand, Proactive explanations anticipate the need of the intended recipient for an explanation. Proactive explanations will be triggered as soon as the processing, event, or (a part of the) decision-making process is fulfilled.

**Proactive explanations**

An individual receives an unexpected or undesired decision (e.g. a lower credit score than expected). The decision is accompanied by an explanation, providing the reasons this decision has been taken, the data sources that were taken under consideration during the decision-making process, a brief overview of the decision-making process and the possible effects the decision might have on the individual.

**3.4 Trigger**

The Trigger dimension (Fig. 6) expresses the event(s) that trigger(s) the generation of an explanation. As regards reactive explanations, the triggering event will typically be an action performed by the explanation recipient. However, explanations can also be triggered in the absence of a request or input from the recipient of the explanation, like in the example of an automated decision mentioned above. Triggering events comprise data processing activities such as the collection of an input to a decision-making pipeline, or the generation of an output from a decision-making pipeline.

**Triggering events**

**Action:** a request by a customer to switch telephone providers, triggering an explanation about the necessary time to fulfil the request.

**Processing:** the processing of personal data, triggering the production of explanations to the data subject(s) about the details of processing (even before a decision is reached).

**Decision:** the rejection of a loan application, triggering the production of an explanation about the factors that disqualified the applicant.




Fig. 4. The Timing dimension




Fig. 5. The Autonomy dimension

### 3.5 Content

The Content dimension (Fig. 7) aims to capture details to be included within the formulation of an explanation. The Sensitivity sub-category defines the sensitivity level associated with the content of the explanation. When explanations contain personal data it is important to label it as such. Hence the sub-category of Personal Data. The content of these explanations usually falls within the scope of data protection legislation and, as a result, their use is restricted and made conditional. Personal data can however be subject to a variety of techniques to minimise the data and reduce re-identification risks, such as generalisation, pseudonymisation and anonymisation techniques. Within the category of Personal Data it is often useful from a risk perspective to distinguish between Directly Identifying and Non-Directly Identifying data points. The categories Generalised Values and Personal Data are not mutually exclusive.<sup>17</sup>

**Sensitivity**

**Generalised Values:** An explanation of age-restricted access rights may refer to age brackets (e.g. < 18, 21+ etc.) of user groups instead of the ages of the users.

**Personal Data:** An explanation about the personal data that has been processed after the exercise from a data subject of their right to access will contain all the data that has been associated with an identified or identifiable individual, even if the particular data point is not directly or indirectly identifying.

The Confidentiality category distinguishes between Disclosed and Restricted explanations. The rationale for this category is to indicate whether an explanation is or can be disclosed to the public or whether access control restrictions should be put in place, e.g. because they amount to trade secrets or classified information. The Disclosed and Restricted categories can be further specified to add granularity.<sup>18</sup> The category Disclosed is relevant for all explanations that can be accessed by external actors and Restricted should apply to explanations that are intended for internal use only. For both types of explanations however, access should be determined on the basis of the need-to-know and least privilege principles, unless disclosed explanations are meant to be publicly available.

**Confidentiality**

**Disclosed explanation:** An explanation about a decision to reject an application for a loan because of a low credit score.

**Restricted explanation:** The weights used by a credit scoring model in the applicants' data to calculate the specific credit score that led to the rejection of the application.




Fig. 6. The Trigger dimension




Fig. 7. The Content dimension

<sup>17</sup>In Bridges, the UK Court of Appeal confirmed that the test for determining whether data is personal is whether “the information recorded by [the controller] individuated [the individual] from all others, that is to say it singled him out and distinguished him from all others” (on the application of Edward BRIDGES) [2020] EWCA Civ 1058, para. 46. In IAB Europe, the CJEU confirms that user preferences amount to personal data when the data may “by reasonable means, be associated with an identifier, such as, inter alia, the IP address of that user’s device.” CJEU C-604/22 IAB Europe v Gegevensbeschermingsautoriteit 7 March 2024 ECLI:EU:C:2024:214, para. 51

<sup>18</sup>For example, by breaking down access levels in an ISO:27001 inspired format to confidential for senior management access; restricted for access by most employees; internal for access by all employees; and public information to which everyone has access. See [28] for more details about the levels.

The Minimum Content category contains the types of information that must be included within the explanation to comprehensively address the explanation requirement. The categories of this dimension are intimately related to the domain and effect of the explanation requirement. The Minimum content is thus usually requirement-specific and it is not possible to generalise further at this stage. Explanations must contain information to cover all the necessary bases for the explanation requirement they intend to address. The Minimum Content category correlates with the Source dimension as well as the Sensitivity and Confidentiality dimensions. For example, an explanation containing intellectual property data classified as confidential might need to undergo further processing to obfuscate the data in question before being disclosed to third parties. The Minimum Content is also correlated with the Intended Recipient. Sokol and Flach [52] explains, for example, that “a system designed as a certification tool will usually reveal more information than one providing explanations to customers.”

#### Minimum content

An organisation making automated decisions without human intervention about individuals under Article 22 of the GDPR must provide an explanation about the decision to the subject of the decision. According to [26, pp. 18–19] the explanation shall contain at minimum:

- (1) which information was taken into account;
- (2) the rationale behind the decision;
- (3) the key decision points that formed the basis for the decision;
- (4) any alternative decisions that were considered and why they were not preferred;
- (5) how to ask for a review;
- (6) how to make an appeal and the available appeal grounds.

### 3.6 Scope

The Scope dimension (Fig. 8) assesses whether the explanation only applies to a particular case or whether it can be reused in other contexts. Local explanations are context-specific. For example, when the recipient changes, the content will have to change as well. In contrast, a Universal explanation remains the same each time the decision-making pipeline produces an output.




Fig. 8. The Scope dimension

#### Scope

**Local:** The over-subscription criteria of each of the selected schools in a school admission application will result in an explanation whose scope will be limited to the application in question.

**Universal:** The contact details of the Office of the School Adjudicator remain immutable. An explanation about how to contact the Office to challenge a decision will have applicability across all school allocation decisions.

### 3.7 Explainability Goal

The Explainability Goal dimension (Fig. 9) defines the functional purpose(s) that the generation of a particular explanation intends to achieve. Contrary to what is often assumed, explanations can serve a variety of purposes. The meaningfulness of an explanation, which is dependent upon its wording and format, is intimately related to the specific goal it pursues.

To be both compliant and effective, an explanation must enable the recipient to understand a decision and/or a particular aspect of a decision-making process, and, when necessary, to take action. Explainability goals can thus be divided into two families: Understandability goals and Intervenability goals. The two families reflect the main “cognitive effects” of explanations on recipients, outlined in [34] and [9].<sup>19</sup> By way of example, the goals of justification and transparency are two sub-categories of understandability.

Goals in the Understandability family (Fig. 10) aim at fostering either trust or confidence in the decision output or the decision-making pipeline. Understandability is key in building trust for systems where decision-making can be highly consequential (e.g. because it produces legal effects), and the correctness or validity of the decisions cannot be readily verified [60]. We have identified eight of those goals:

- Performance**, assessing the performance of the system;
- Accountability**, explaining whether and how the system owner is in control;
- Accuracy**, reporting on the accuracy of the input data;
- Consequences**, explaining the processing impact upon the recipient;
- Data minimisation**, explaining how the input data has been tailored to the processing purpose;
- Fairness**, explaining how biases and discrimination have been mitigated;
- Overall Transparency**, explaining the main steps of the processing system;
- Justification**, explaining the reasons for the output.

Goals belonging to the Intervenability family (Fig. 11) aim at empowering recipients to take action,<sup>20</sup> e.g. to exercise corrective individual rights or more simply to monitor the performance of the system (e.g. trouble-shooting). The Intervenability family comprises ten goals:

- Data access**, accessing input data;
- Contestability**, challenging an output;
- Data erasure**, deleting input and/or output data;
- Behaviour modification**, modifying the behavior of the system;
- Human intervention**, performing human review;
- Data portability**, porting data to a third party;
- Data rectification**, rectifying input/output data;
- Additional information**, accessing additional information about the system;
- Restriction**, restricting the use of the data or system;
- Scrutability**, flagging errors.




Fig. 9. The Explainability goal dimension

<sup>19</sup>Maybury [34] identifies “entity knowledge”, “event and state knowledge” and “know-how knowledge” relating to understandability and “change beliefs/evoke action” relating to intervenability, whereas Chander and Srinivasan [9] identify “education” and “trust” on the one hand and “design”, “action” and “trouble-shooting” on the other.

<sup>20</sup>The concept of intervenability is borrowed from the literature on data protection goals [2].

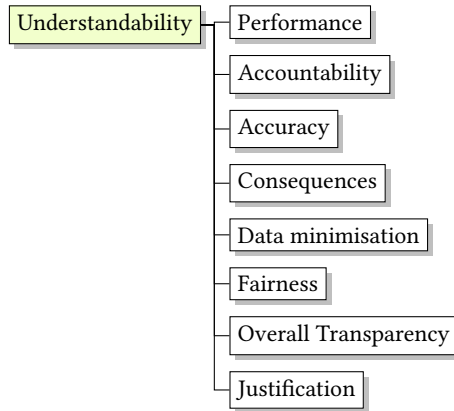


Fig. 10. The Understandability goals

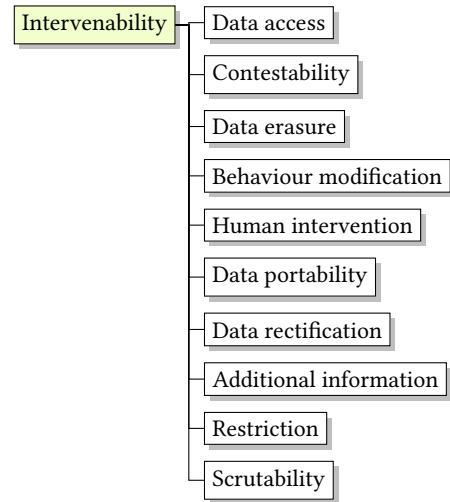


Fig. 11. The Intervenability goals

Explainability goal aims to highlight the purpose of the generated explanation so that its effectiveness and compliance can be assessed against the explainability requirements. An explanation may serve multiple goals at once. For example, an explanation about the use of automated decision-making aims at informing the recipient about the operation of the system (the Transparency goal) but also holding the decision-maker accountable (the Accountability goal). Explanations from different goals can also be combined together into larger explanatory statements: The explanation under Article 22 of the GDPR must explain the reasons and impact of an automated decision (Fairness, Overall Transparency and Consequences), but should also provide instructions on how to request a human review of the decision and appeal it (Human intervention and Contestability).

### 3.8 Intended Recipient

The Intended Recipient dimension defines the categories of recipients that will be associated with the explanations (Fig. 12). Identifying the target audience of an explanation is essential to be able to assess understandability, as comprehension skills vary across groups of recipients. Besides, as Bohlender and Köhl [7] argue, “the notion of explanation is not absolute but rather relative to a target group”. This dimension draws a distinction between explanations targeting the public and explanations for internal consumption. Outward-facing explanations can target the public, including data subjects and other decision addresses, supervisory authorities responsible for overseeing the organisation’s processing or other third parties that interface with the organisation. Explanations can also be Internal-facing, aiming to assist the employees of an organisation in monitoring and troubleshooting the processes and their compliance with applicable rules. At first glance, five broad categories of internal-facing recipients can be distinguished: administrators of the decision-making system; business analysts making decisions; legal engineers who determine the primary and secondary requirements and contribute to the elicitation of tertiary requirements; data engineers who are responsible for managing the data; and auditors, who are in charge of monitoring compliance. However, this dimension is highly dependent on the nature of the implementing organisation. An organisation could choose to produce explanations for third parties participating to the decision-making pipeline.

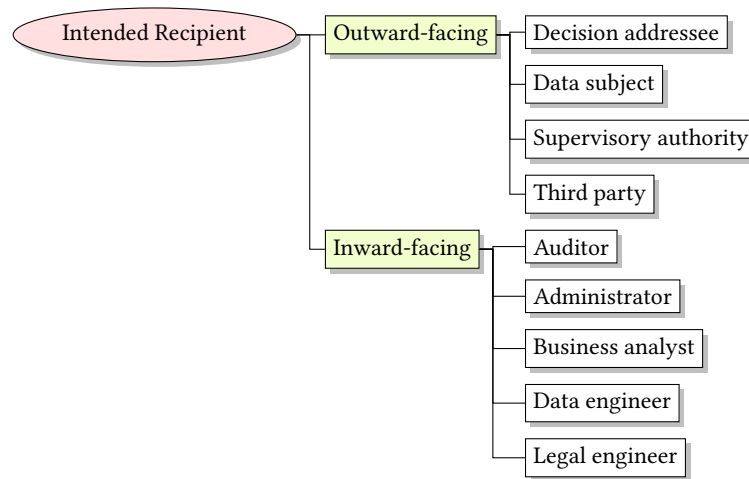


Fig. 12. The Intended Recipient dimension




Fig. 13. The Criticality dimension

**Intended Recipients**

**Outward-facing:** An organisation will need to provide explanations to subjects of automated decisions expressed in layman terms, and may want to disclose explanations to the supervisory authority to demonstrate how its operations comply with the law.

**Inward-facing:** Explanations monitoring the automated operations of the organisation can address software engineers, in-house lawyers or business managers. Since these groups have different types of specialist knowledge, they will each require a type of explanation that takes into account the degree of their related expertise.

**3.9 Criticality**

In practice, there are cases where an obligation to explain an action or omission arises on the basis of an applicable legal rule. However, there are cases where an organisation can decide that providing an explanation is beneficial even in the absence of mandatory requirement. The final dimension, Criticality, depicted in Fig. 13, captures the distinction between a mandatory and a recommended explanation requirement.

An explanation requirement is **Mandatory** where an applicable legal rule and/or a governance framework imposes an obligation to provide an explanation. On the contrary, an explanation is **Recommended** where it is not strictly speaking mandatory to generate an explanation to fulfil an obligation imposed by a legal and/or governance prescription.

**Different criticality levels**

**Mandatory:** Under the School Admissions Code, admission authorities who inform parents of a decision to refuse their child a place at a school are obliged to explain the reasons for the refusal.

**Recommended:** A school admission authority may want to explain to parents that if they accept their second or third preference, their child will be placed on the waiting list for their first and eventually second preference.

#### 4 PRODUCING EXPLANATIONS ON THE BASIS OF THE TYPOLOGY

A first round of explanations was produced by the research team and presented to the pilot study participants, who worked with the research team to fine tune them in an iterative process. Explanations were derived from both primary requirements, i.e., the GDPR, the School Admissions Code, and secondary requirements, i.e., relevant regulatory guidance.<sup>21</sup>

##### 4.1 Loan application

In this scenario, a prospective applicant, Alex, a UK resident, wishes to apply for a loan through the website of a bank located in the UK, the Bank. The bank requests some information about the income and living habits of the applicant. It then uses this information together with a credit score for Alex, received from a collaborating credit reference agency – the CRA, to calculate the creditworthiness of Alex. If Alex’s creditworthiness is above a certain threshold, the application is automatically approved. Below a certain threshold, the application is automatically rejected, whereas if it falls in between the thresholds, it is redirected to a manual assessment by an officer. Because the Bank uses solely automated decision-making for applications that do not go to manual assessment, UK GDPR Article 22 applies. The article prohibits solely automated decision-making ingesting personal data of individuals if the decision produces legal or similarly significant effects for them. Exceptionally, solely automated decision-making can take place if

*the data controller [implements] suitable measures to safeguard the data subject’s rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision.*

Recital 71 specifies that suitable measures “should include specific information to the data subject and the right [...] to obtain an explanation of the decision reached”. Although only persuasive at most, national courts within the EU have on occasions recognised such a right.<sup>22</sup>

Analysing UK GDPR Article 22, Recital 71 and other relevant provisions of the GDPR, we started by distinguishing five distinct explanation requirements. Organisations using solely automated decision-making to produce significant effects must:

1. Provide information about the automated decision-making system to the subjects of the decisions;
2. Provide an explanation about the decision reached;
3. Allow the subject of the decision to express their point of view;
4. Provide a way for the subject of the decision to obtain human intervention; and,
5. Advise how they can challenge the decision.

Under UK GDPR transparency requirements, the information provided to the data subject, i.e. the future subject of the automated decision, must go beyond the mere existence of automated decision-making. Data controllers must also supply “*meaningful information about the logic involved*” as well as the “*significance*” and “*envisaged consequences*” of any automated decision.<sup>23</sup>

<sup>21</sup>Arising out of the Information Commissioner’s guidance for the GDPR, for AI and for credit scoring.

<sup>22</sup>See [48] at para. 4.39 “In that case, the controller must still take appropriate measures, including at least the right to human intervention, the right of the data subject to make his point of view known and the right to challenge the decision (Article 22 (3) GDPR and recital 71 GDPR)”. See also [14, 45] where the Austrian and Irish DPAs held that credit scoring companies were under an obligation to provide meaningful information about the calculation of the credit scores.

<sup>23</sup>GDPR articles 13(2)(f), 14(2)(g) and 15(1)(h).



Source	Timing					Autonomy					Content			Scope		Explainability goal		Intended recipient		Criticality	
	Ex ante	Ex post	Proactive	Reactive	Decision	Process	Action	Sensitivity		Confidentiality		Minimum Content		Local	Universal	Understandability	Intervenability	Outward-facing	Inward-facing		
								Non-Directly Identifying	Directly Identifying	Disclosed	Restricted	Information about the logic & significance of the decision	Meaningful enough to facilitate the data subjects' exercise of their rights								Categories (types) of data processed and their origin
Existence of automated decision-making	Primary	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Accountability	Overall Transparency	Data subject	Data subject	Mandatory	Recommended
Meaningful information about the logic	Primary	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Accountability	Overall Transparency	Data subject	Data subject	•	•
The categories of personal data used in the decision-making process	Primary	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Accuracy	Fairness	Data subject	Data subject	•	•
The relevance of the information for the decision-making	Secondary	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Data Minimisation	Fairness	Data subject	Data subject	•	•
Personal Data points	Primary	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Accuracy	Fairness	Data subject	Data subject	•	•
The significance of envisaged consequences	Primary	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Consequences	Fairness	Data subject	Data subject	•	•
Decision taken solely by automated means	Primary	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Accountability	Overall Transparency	Data subject	Data subject	•	•
Reasons for reaching the decision	Primary/implicit	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Justification	Fairness	Data subject	Data subject	•	•
Key data points	Primary/implicit	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Accuracy	Fairness	Data subject	Data subject	•	•
Key associated weights	Secondary	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Justification	Fairness	Data subject	Data subject	•	•
Process to request a review of the decision	Primary/implicit	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Accountability	Consequence	Data subject	Data subject	•	•
Response to the review request	Primary/implicit	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Fairness	Overall Transparency	Data subject	Data subject	•	•
	Secondary	•	•	•	•	•	•	•	•	•	•	•	•	•	•	Accountability	Human intervention	Data subject	Data subject	•	•

Table 2. Classification of explanation requirements related to GDPR Articles 13, 14, 15 and 22.

Source	Timing			Autonomy			Trigger			Content			Scope		Explainability goal		Intended Recipient		Criticality
	Ex ante	Ex post	Proactive	Reactive	Decision	Process	Action	Sensitivity	Confidentiality	Minimum Content	Universal	Local	Understandability	Intervenability	Outward-facing	Inward-facing	Mandatory	Recommended	
Key decision points									Non-Directly Identifying	Disclosed	Restricted	Key data values	Accuracy Performance	Accountability	Scrutability		Administrator Auditor		
Verification of the results									Directly Identifying			Processed data; data sources; model rules; audit trails	Performance	Accountability	Scrutability		Administrator Auditor		
Business rules applied to the decision												Business rules; processing purposes	Accountability		Scrutability		Administrator Auditor		
Alternatives to the decision & reasons for rejection												Alternative decisions; benefits of decision chosen or drawbacks of alternatives	Performance	Accountability	Modifying a behaviour		Administrator Business analyst Auditor		
Review carried out by qualified & authorised reviewer												Contact info of reviewer; qualifications; track record	Accountability		Scrutability		Administrator Auditor		

Table 2. Classification of explanation requirements related to GDPR Articles 13, 14, 15 and 22 (continued).

The explanation mandate about the way the automated decision-making system works and its consequences comes on top of the obligation to provide information about the nature and purpose of the processing, the identity of the data controller and the available data subject rights.<sup>24</sup> The UK GDPR does not give examples of what it means by “*meaningful*”, but at the very least the information should be meaningful enough to facilitate the exercise of data subjects’ rights [50]. In [26], the UK Information Commissioner specifies additional requirements that are necessary to satisfy Article 22. The meaningful information about the logic must also include the types of information used in making the decisions and why this information is relevant. In explaining a specific decision to a data subject, the organisation should explain how the decision was reached and its rationale. To assist organisations in generating comprehensive explanations, the Information Commissioner recommends that they keep records of the key decision points, the verification methods of the result, the business rules that apply to the decision and whether any alternatives existed. If there were alternatives, organisations should record why they were not preferred over the outcome of the decision.

An explanation is also required to enable the subject of an automated decision to exercise her rights. Article 22 requests clear communication of the right to ask for review, human intervention, express objections, and contest the decision. The Information Commissioner considers that it would make business sense to provide an explanation at the point of delivery of the decision to inform the subject about how to request a review or appeal of the decision. Organisations have to demonstrate that the review is carried out by a qualified reviewer who is authorised to assess the decision-making process. At the end of the review, they have to issue a formal response explaining the outcome of the review process.

Overall, we started by deriving GDPR-related explanation requirements. These requirements are not exhaustive and could be expanded by considering other articles or translating additional parts of the ICO’s guidance. These have been classified in Table 2. The classification allows the identification of necessary components to construct a targeted explanation for each of the GDPR requirements. In addition, it assists in determining potential overlaps between requirements, or eventual complementarity. Working with our partners, we decided to add explanations requirements for inward-facing monitoring and auditing, which are also recommended by the ICO. The remaining dimensions determine the delivery rules that dictate when and to whom the explanation will be presented.<sup>25</sup> For the delivery of explanations, the chosen format follows the design guidelines identified by Simkute et al. [51]:<sup>26</sup>

We thus constructed explanations for the explanation requirements related to UK GDPR Articles 13, 14, 15 and 22, as illustrated in the coloured samples below, where the minimum content is highlighted. The samples are not set in stone and only aim at shedding light on the tension between data subjects’ actual needs and controllers’ own interests.

To provide information about automated decision-making before the processing takes place (*ex ante*) we produced the following explanations, based on the break down in Table 2.

<b>R1: Existence of automated decision-making</b>
E1: To be able to process applications quickly and accurately, the Bank uses a solely automated system.
<b>R2: Meaningful information about the logic</b>

*Continued on next page...*

<sup>24</sup>See GDPR articles 13, 14 and 15.

<sup>25</sup>The details about the technical design and implementation are beyond the scope of this paper. A thorough explanation of the technical steps can be found in [25].

<sup>26</sup>See pp. 11–12, where the authors provide a list of design goals tuned to different expertise levels, e.g. the use of pop-up alerts; the use of structured information and visualisations; the ability to actively question the data; the provision of interactive manipulation; etc.

(...continued)

E2.1: The system will screen your application based on the information you provide in the application form, information we hold on file about you and information we receive about you from our partner the CRA. Your information will be assessed against an affordability threshold, to determine whether you can satisfy the affordability criteria for your requested credit. If a decision is made, we will inform you about the relevant information that has impacted the outcome.

E2.2: The system is calculating the probability of meeting our affordability criteria based on information about your financial standing that you provide in your application, that we already hold on file about you, or that we receive from our partner the CRA. You can experiment with how different financial information can impact our confidence that credit repayment is viable using the slider below.

**Projected loan variations for responsible lending**

Annual Income:	£10,000.00	<input type="range"/>	£100,000.00
Monthly instalment pay- ment:	£100.00	<input type="range"/>	£50,000.00
Previous Credit:	£100.00	<input type="range"/>	£50,000.00
Permanent residency in the UK:	< 1 year	<input type="range"/>	> 5+ years
APR:		<input checked="" type="radio"/> Fixed <input type="radio"/> Variable	

---

For a requested credit amount £  --  a  interest rate is projected to be  %.

Note: This tool is provided for illustration purposes only and does not constitute a decision or a promise of a decision of an actual application.

### R3: The categories of personal data used in making the decision

E3: The data that will be processed includes personal identifiers, such as name, date of birth, and current and previous addresses; information about your financial standing, such as the number of credit and current accounts you hold and their repayment histories, salary data or insolvency-related events; information about your credit rating, such as scores about your credit, your affordability, or the possibility of insolvency or fraud.

### R4: The relevance of the information for decision-making

E4: The personal identifiers are used for matching, so that we know that the processed data relates to you. Information about your finances and credit scores are used to determine the likelihood that you will be able to repay your credit commitments.

### R5: The significance of envisaged consequences

E5.1: The screening process will determine whether your application is successful.

Continued on next page...

(...continued)

E5.2: If your application is successful, the Bank will sign a loan agreement with you and credit you a loan of the amount specified in your application. You will have to repay this loan following the repayment conditions agreed in the loan agreement. Following our screening, the Bank may consider that a loan of a different amount is financially viable to your specific circumstances. You will have the opportunity to accept or decline the Bank's offer. The Bank will provide you with an explanation in case your application is unsuccessful, explaining the reasons why. You will have the chance to challenge that decision if you consider that your application was not processed correctly.

Note that the typology can be used to create various types of explanations. For example, E2.2 uses a counterfactual explanation to illustrate the consequences of the underlying logic. Counterfactual explanations can demonstrate how different behaviour impacts the outcome of a process. However, their usefulness in placing the recipient in a position to challenge a financial decision that has already been made is limited since one is unlikely to be able to alter their financial assets (e.g. their income) at will.

Information given to the data subject after a decision has been made (*ex post*) could expand the first set of explanations in the following way:

<p><b>R1: Decision taken solely by automated means</b></p> <p>E1: The Bank made this decision with an automated scoring system that takes into account the information provided in the borrower application as well as a credit score produced by credit referencing agency the CRA. You can access the personal data about you that were processed for this decision by following <a href="#">this link</a>.</p>
<p><b>R2: Reasons for the decision</b></p> <p>E2: We regret to inform you that the loan application (applications_no/437) was declined. This is because of negative credit history.</p>
<p><b>R3: Key data points for the decision</b></p> <p>E3: The borrower's credit score indicates the following negative information. The late payment (records/70551), late payment (records/70552) and late payment (records/70553). The assessment indicated that the commitment of the credit agreement would have an adverse impact on the borrower's financial situation. Hence the company decision (applications_no/437/decision) is to refuse the borrower's application.</p>
<p><b>R4: Process to request a review of the decision</b></p> <p>E4: If you are not satisfied with this decision, you can get in touch with us in the email address below. If you think some of the data in this decision are incorrect, you can request a review of the decision by one of our officers. You will need to provide any new evidence you consider should be taken into account in assessing your application. To express your point of view or request a review, please send an email to &lt;email@thebank.bank&gt; quoting applications_no/437/decision.</p>
<p><b>R5: Response to the review request</b></p> <p>E6: We have reviewed decision applications_no/437/decision about you. In light of the new evidence you presented about late payment (records/70551) and late payment (records/70552), we are happy to inform you that we are in a position to offer you a credit amount of application_no/437/credit_amount/adjusted. You have 30 days to respond if you would like to accept this offer.</p>

Of note, Article 29 Working Party also includes within its own example, “information to advise the data subject that the credit scoring methods used are regularly tested to ensure they remain fair, effective and unbiased”. [4]. We however held the view that such a comforting statement did not mean much in practice if it was not backed up by additional information, which would allow data subjects to effectively verify the accuracy of such a claim.

To assist the Bank’s staff in monitoring and reviewing the decision-making process, we also produced the following explanations:

<p><b>R1:Key decision points</b></p> <p>E1: The applicant’s credit score is below the acceptance threshold of 750. The applicant’s credit score was impacted by <b>late payment (records/70552)</b>, <b>missed payment (records/70553)</b> and <b>late payment (records/70551)</b>. The source of this information was <b>credit history report credit_history/437</b> provided by <b>credit_agency_3</b>.</p>
<p><b>R2:Verification of the results</b></p> <p>E2: Identity matching was performed using <b>electoral register data received on 2021-01-03T02:00:00</b>. The credit report and credit score from CRA were <b>calculated on 2020-10-03T04:00:00</b>. <b>Late payment (records/70552)</b> of 2020-10-01T01:00:00 was <b>matched to late payment</b> in credit report <b>credit_history/437</b>.</p>
<p><b>R3:Business rules applied to the decision</b></p> <p>E3: <b>The threshold of acceptance</b> for <b>credit_product/322</b> is 750. The decision was balanced against the FCA rule about the Bank’s <b>commitment to responsible lending</b>. The decision took into account the protection of characteristics under the <b>Equality Act 2010</b>.</p>
<p><b>R4:Alternatives to the decision &amp; reasons for rejection</b></p> <p>E4: To overturn the decision, a score above 750 was necessary. The current score of 715 could justify a loan of up to £10,000 with a fixed interest rate of 3/1 with an initial APR of 3.7%. <b>This option was rejected because the amount surpasses the acceptable difference of ±20%</b>.</p>
<p><b>R5:Review carried out by qualified &amp; authorised reviewer</b></p> <p>E5: The decision was <b>reviewed at 2021-02-18T10:25:45.000Z</b> by <b>credit officer (staff/142)</b> who confirmed the credit and risk assessments.</p>

These sample explanations are the result of a co-creation process, which implied signed-off by the industry partners. The explanations produced for internal monitoring and auditing are thus richer (in terms of details) than the explanations addressed to data subjects. It is likely that the range of explanations addressed to data subjects would need to be expanded to include additional information, such as information about the respective weights of the factors taken into account to produce the decision, i.e., when a machine learning model is deployed the respective weights of the relevant model features [35], which we could not capture with the provenance trail developed for the purpose of this paper. In any case, to determine whether the information provided meets an explanation requirement, the information should be assessed in the light of the explainability goal pursued. If contestability is a relevant goal, then it is not sufficient to show that the addressee of the decision is in a position to make sense of the information provided. The decision-maker should also make sure that the information provided enables the addressee of the decision to assess the lawfulness of the decision-making process and its output.

	Timing			Autonomy			Trigger		Content		Scope		Explainability goal			Criticality	
	Source	Ex ante	Ex post	Proactive	Reactive	Decision	Process	Action	Sensitivity		Universal	Local	Understandability	Intervenability	Intended recipient		
									Non-Directly Identifying	Directly Identifying							Confidentiality
Reasons for the decision	Primary	•	•	•	•	•	•	•	•	•	•	•	•	Justification Fairness	Contestability	Decision Addressee	•
Information related to the right to appeal	Primary	•	•	•	•	•	•	•	•	•	•	•	•	Accountability Consequence	Contestability	Decision Addressee	•
Consequences of the decision	Secondary	•	•	•	•	•	•	•	•	•	•	•	•	Consequences Fairness	Contestability	Decision Addressee	•

Table 3. Classification of explanation requirements related to the School Admissions Code

## 4.2 School Allocation

In this scenario, Bob and Mary live in Southampton, UK. They have a son Chris, who will turn 4 years old in May. He should thus start school in September. Bob and Mary apply for a school place at primary school level. As part of the application process, Bob and Mary declare their preferences for three schools: *Bassett Green Primary*, *Oakwood Primary* and *St Denys Primary School*. In June, they receive a letter from the Southampton city council telling them that Chris has been admitted at *Oakwood Primary*.

In England, school admissions are covered by the School Admissions Code 2021 [16]. The Admissions Code is a statutory code of practice issued under the School Standards and Framework Act 1998 (SSFA).<sup>27</sup> All admission authorities and admission appeal panels have a statutory duty to comply with it.<sup>28</sup>

Each school has discretion to choose oversubscription criteria that prioritise applicants if the applications are more than the number of available places. When assessing the application for a child, the local authority must apply the oversubscription criteria of each of the preferred schools. If the child satisfies the oversubscription criteria of more than one schools, a place at the highest ranked of these school is offered. If the child does not satisfy the oversubscription criteria of any of the preferred schools, the Council must offer the child a place at the nearest school with empty places.

In our example, *Basset Green Primary* prioritises applicants according to five criteria:

- (1) Children in public care
- (2) Children subject to a child protection plan
- (3) Children whose sibling is on the roll of the school
- (4) Children who attend one of the feeder nurseries
- (5) Children living closest to the school

*Oakwood Primary* prioritises applicants according to 3 criteria:

- (1) Children in public care
- (2) Children whose sibling is on the roll of the school
- (3) Children living closest to the school

In addition, since admission processes require the processing of personal data, the Data Protection Act 2018 (DPA 2018) and the UK GDPR apply, and in particular UK GDPR Art. 12 and 13. Article 22 is out of scope as long as the process is not merely automated.

For the purposes of this paper, we focus upon two requirements stemming from the School Admissions Code. Para 2.32 of the Code imposes upon admission authorities the following requirement:

*“When an admission authority informs a parent of a decision to refuse their child a place at a school for which they have applied, it must include the reason why admission was refused; information about the right to appeal; the deadline for lodging an appeal and the contact details for making an appeal. Parents must be informed that, if they wish to appeal, they must set out their grounds for appeal in writing. Admission authorities must not limit the grounds on which appeals can be made.”*

Such a requirement can be broken down into two requirements as illustrated in Table 3.

Besides, for the sake of completeness, the Southampton school admission team decided to add a requirement to explain the consequences of accepting the decision.

<sup>27</sup>School Standards and Framework Act 1998 c. 31, ss. 84 and 85.

<sup>28</sup>School Standards and Framework Act 1998 c. 31 s. 84.



To assist the City Council in producing explanations once school-allocation decisions have been made, we produced the following examples of explanations. The first example aims to inform parents/guardians of the reasons for the decision, highlighting the decision criteria per school and disclosing the breakdown of children meeting each criterion. The second example aims to inform parents/guardians of the next steps they can take once they have received the decision and want to appeal the decision. The third examples aims to inform parents/guardians about the consequences of accepting the decision. The minimum content is highlighted for each requirement.

**R1:Reasons for the decision**

E1: Too many applicants applied for a place in **Bassett Green Primary** and as a result we are not able to offer Chris a place in this school. **Bassett Green Primary** prioritises applicants according to five criteria: Crit#1, Crit#2, Crit#3, Crit#4, and Crit#5. Chris satisfied criterion Crit#5, **where he was ranked at Rank#33** within Crit#5 based on the distance of your house to the school. Unfortunately, given the number of applicants satisfying prior criteria, admission of children under criterion Crit#5 was **cut off at Rank#18**. Any further admission to this year group will prejudice the provision of efficient education and the efficient use of resources. A place has been allocated to Chris at **Oakwood Primary**, your 2nd preference school, starting September 2024. A detailed breakdown of admitted children by admission criterion can be found below.

**Basset Green Primary:**

- Published Admission Number: 175
- Application received: 616

**Crit#1** Children in public care 15  
**Crit#2** Children subject to a child protection plan 17  
**Crit#3** Children whose sibling is on the roll of the school 30  
**Crit#4** Children who attend one of the feeder nurseries 105  
**Crit#5** Children living closest to the school 18

The last child to be admitted at **Basset Green Primary** under Crit#5 was ranked 175 out of 616 on their list of applicants and lived 1.9 miles from the school.

**Oakwood Primary:**

- Published Admission Number: 180
- Application received: 432

**Crit#1** Children in public care 8  
**Crit#2** Children whose sibling is on the roll of the school 85  
**Crit#3** Children living closest to the school 87

The last child to be admitted at **Oakwood Primary** under Crit#3 was ranked 180 out of 432 on their list of applicants and lived 2.8 miles from the school.

**R2:Information about the right to appeal**

E2: If you wish to appeal the Council’s decision, you must lodge your appeal by **30 June 2024** at [appeal@southampton.gov.uk](mailto:appeal@southampton.gov.uk). You can find more information on how to commence the process at <http://www.southampton.gov.uk/schoolappeals>. You should specify your grounds **in writing**.

**R3:Consequences of the decision**

Continued on next page...

(...continued)

E3: If you **accept** the place at Oakwood Primary your child's name **will still go on the waiting list for your higher ranked school**, i.e., Bassett Green Primary.

## 5 EVALUATION

Once explanation requirements were generated in collaboration with domain experts from our industry partners, explanation outputs were assessed by representatives of producers and consumers of explanations gathered by our industry partners. We then revisited our methodology and typology in the light of the EU AI Act adopted later on.

### 5.1 Prior to the AI Act

The pilot studies, conducted prior to the adoption of the EU AI Act, confirmed the feasibility and usefulness of the methodology and typology that was produced for the purpose of computing explanations. By feasibility, we refer to the practicality of translating high-level explanation requirements found in regulatory, compliance or business mandates into specific, computable components. This involved testing whether high-level explanation requirements could be deconstructed into smaller, actionable parts and whether these could be processed by a computational system to provide outputs to users. The study allowed us to explore related operational challenges, such as challenges related to the gathering and translation of explanation requirements and the identification of relevant human and computational resources.

Provenance data models make it possible to traverse and query audit trails of the processing, records that describe the people, institutions, entities, and activities involved in producing, influencing, or delivering the outcome of a processing operation [39]. We thus designed a provenance data model capable of inspecting the data flows generated in the context of the two use cases. In a nutshell, the Minimum Content category was translated into a provenance pattern that retrieves the necessary data from the audit trail, while the Intended Recipient and Explainability Goal were used by a Natural Language Generation engine to construct the explanation after a linguistic syntax tree.

In addition to feasibility, the study also examined the usefulness of these explanations, meaning how well they serve their intended goals. A useful explanation must not only be technically feasible but also meaningful for its audience. This includes testing whether the generated explanations allowed users to get benefits from the output in the light of the state of the art, i.e., existing practices.

The feasibility was demonstrated through the production of a light ontology and an explanation assistant, which was created to address the needs of two use cases: credit scoring and school admissions. The light ontology can be found in Appendix A. Through a machine-readable ontology, we adopted a systematic terminology for concepts, reflected on their inter-relations, and derived hierarchical representation. It also allowed us to systematically encode sets of requirements according to the ontology. These sets of requirements are subsequently used in the explainability by design methodology to generate explanations.

The Explanation Assistant is a software service configured to support a set of explanation requirements generated by a Legal Engineer following the proposed taxonomy. From the application data associated with a decision, it produces narratives in a natural language to explain that decision. Example of computable explanation narratives were presented to two industry partners in the context of interactive focus groups: one CRA (Experian UK) and one school admission team (from Southampton City Council):<sup>29</sup>

<sup>29</sup>For more information, see <https://plead-project.org/research/>.

As regards usefulness, producers (data teams) and consumers (teams in charge of responding to end-user requests for information) of explanations from our industry partners were then asked to assess the explanations, taking into account the solutions that they had already deployed prior to the beginning of the project. They were thus gathered in two focus groups to answer a set of questions in a moderated setting. Each group was chosen due to its prior interactions with explanations produced by each industry partner, and the questions were designed to shed light on the quality of explanation outputs, taking into account the two explainability goals mentioned above: intervenability and understandability. All participants to the focus groups confirmed that the explanations produced on the basis of the typology went further than their existing practices and that if implemented would improve these practices by enhancing the level of understandability, either because they included more details about the decision-making process or because they included more tailored information for the recipients of the explanations. During interactions with producers and consumers of explanations, participants made clear their intention either to become more sophisticated providers of explanations themselves or to delegate explanation-related tasks to third parties with appropriate expertise.

The benefit of the typology was considered to be twofold: First, it forces the legal engineer to systematically identify all building blocks needed to compute requirements for explanation automation. Second, the typology helped the legal engineer to systematically organise explanation requirements and identify duplicates. An organisation may find that multiple explanation requirements map to the same dimensions and sub-categories. These requirements can therefore be addressed by similar explanation(s), allowing the organisation to scale their compliance strategies. The legal engineer was then ready to be embedded within the engineering team to help it develop a software architecture relying upon reusable components and aiming to incorporate the resulting explanation capability into an application. While the fine-grained explanation requirements generated on the basis of the typology have been used as input for the work of the data engineer, an iterative process supported by an on-going dialog between the legal engineer and the data engineer had to be followed to perform a variety of engineering tasks, such as modelling the decision audit trail, creating queries to extract data from those audit trails for each explanation, building explanation plans to transform the queried data in narratives consumable by end-users, and ultimately deploying the Explanation Assistant (see [25] for more information).

Considering these benefits in the light of prior work and in particular prior review of methods and techniques that are being developed to improve explainability in machine learning prediction tasks [61], it is possible to argue that the typology can be used as a means to assess the quality of the explanatory approach when the ML model is integrated within a wider decision-making pipeline and to choose the most appropriate explanation for a given explanation requirement, be it a mere business requirement or a business requirement informed by applicable law. The typology forces the elicitation of explanation requirements, which are then used as benchmarks to tailor the content of the explanation, addressing all dimensions of the typology, in an attempt to reduce subjectivity. This process should be useful to make hidden cognitive biases and social expectations vis-à-vis the explanation process [37] emerge more clearly.

Reflecting upon the research and innovation process implemented during the pilot studies, the main characteristics of the legal engineer role emerged more clearly and are listed below:

- **Role overview:** This interdisciplinary role is responsible for translating legal principles and compliance standards into secondary requirements that are actionable by the engineering team.
- **Key responsibility:** First, this role must gather high-level requirements from the compliance and legal team and deliver a preliminary set of secondary actionable requirements to the engineering team. Second, while embedded

within the engineering team, the role holder must iteratively refine the preliminary set of requirements through interactions with the engineering team on a regular basis to support actively the production of a working prototype.

- **Collaboration and communication:** This role must be in regular communication with two teams, i.e., the legal and compliance team as well as the engineering team.
- **Skill and qualification:** this role must have a legal background and a solid understanding of technology and engineering principles. While not necessarily requiring a formal degree in a technical field, relevant certifications or hands-on experience in data science, computer science, or information technology are advantageous.
- **Reporting structure:** Working at the interface of a legal and compliance team and an engineering team the reporting structure can vary depending on the organisation's size, structure, and priorities. However, one option in particular would ensure that the legal engineer's expertise is appropriately leveraged. It is recommended that the legal engineer report to a Data Governance Officer, whose responsibilities lie in ensuring that data-related activities adhere to legal and regulatory standards and whose role is not merely advisory.
- **Dependencies:** This role is characterised by dependencies on effective communication and cooperation with two teams. First, the legal and compliance team, with whom the legal engineer must work closely to determine the range of applicable legal frameworks and associated internal compliance standards, the interpretative position of the organisation as regards key legal concepts and ethical commitments; and second, with the engineering team, with whom the legal engineer must work closely to understand technical requirements, constraints, and opportunities for embedding compliance measures in data systems or pipelines.

## 5.2 The AI Act

The EU AI Act is a horizontal regulation recently finalised by the EU, which prohibits a certain number of AI practices,<sup>30</sup> governs the placing on the market and the putting into service of high-risk AI systems<sup>31</sup> and general-purpose AI models and Systems,<sup>32</sup> and introduces transparency obligations for certain AI systems.<sup>33</sup> As regards high-risk AI systems, in which category one finds "AI systems intended to be used to evaluate the creditworthiness of natural persons or establish their credit score, with the exception of AI systems used for the purpose of detecting financial fraud"<sup>34</sup> and "AI systems intended to be used to determine access or admission or to assign natural persons to educational and vocational training institutions at all levels,"<sup>35</sup> Chapter III is of particular relevance. Within this chapter, Article 10 imposes a whole set of data governance and management practices, which include for example tracking:

- "data collection processes and the origin of data, and in the case of personal data, the original purpose of the data collection;"<sup>36</sup>
- "relevant data-preparation processing operations, such as annotation, labelling, cleaning, updating, enrichment and aggregation;"<sup>37</sup>
- "appropriate measures to detect, prevent and mitigate possible biases identified."<sup>38</sup>

<sup>30</sup>AI Act, Article 5.

<sup>31</sup>AI Act, Chapter III.

<sup>32</sup>AI Act, Articles 53–56.

<sup>33</sup>AI Act, Article 50.

<sup>34</sup>AI Act, Annex III 5(b)

<sup>35</sup>AI Act, Annex III 3(a)

<sup>36</sup>AI Act, Article 10(2)(b).

<sup>37</sup>AI Act, Article 10(2)(c).

<sup>38</sup>AI Act, Article 10(2)(g).

Source	Timing				Autonomy			Trigger		Content		Scope		Explainability goal		Intended recipient		Criticality
	Ex ante	Ex post	Proactive	Reactive	Decision	Process	Action	Sensitivity		Minimum Content		Local	Universal	Understandability	Interveneability	Outward-facing	Inward-facing	
								Non-Directly Identifying	Directly Identifying	Confidentiality	Disclosed							
Origin of the data	Primary/implicit										Source, e.g. data provider, data subject			Accountability	Scrutability	Administrator	Administrator	
Data collection process	Primary/implicit										E.g., web scraping, API access,...			Accountability	Scrutability	Data Engineer	Auditor	
Data collection purpose	Primary/implicit										Phase of model lifecycle (e.g., training, deployment)			Accountability	Scrutability	Administrator	Auditor	
Data preparation actions	Primary: implicit										Data curation actions			Accountability	Scrutability	Administrator	Auditor	
Bias mitigation actions	Primary/implicit										Mitigation techniques			Accountability	Scrutability	Data engineer	Auditor	
Oversight actions	Primary/implicit										Implementation of explainability techniques et processes			Accountability	Scrutability	Administrator	Auditor	
Role of the AI system	Primary										E.g., production of information summaries, sorting, recommending, access grants			Accountability	Contentability	Decision Addresser	Auditor	
Meaningful elements of the decision	Primary										Elements meaningful to contest the decision (e.g., input, source of input, features, consequence of decision)			Justification	Contentability	Decision Addresser	Auditor	

Table 4. Examples of explanation requirements derived from the AI Act

Article 10 also introduces a data relevancy requirement.<sup>39</sup> Article 12 introduces a record-keeping obligation and Article 15 a requirement for effective human oversight. These provisions thus contain substantive requirements which could be evidenced through the generation of explanations, to the primary benefits of internal roles, but also third-party auditors and eventually regulators. In addition, Article 86 introduces a right to explanation to the benefits of any affected person subject to a decision which is taken by the deployer on the basis of the output from a high-risk AI system when two sets of conditions are met: legal or similar effects plus in-scope adverse effects must be produced. When these conditions are met, the affected person has the right to obtain an explanations of the role of the AI system in the decision-making procedure and the main elements of the decision taken.

Following our methodology and leveraging our typology we could for example produce a set of explanations illustrated in Table 4. The last requirement in the table will have to be broken down in more granular sub-requirements.

## 6 LIMITATIONS

All in all, even if the pilot studies achieved their goal, a comprehensive evaluation process remains needed. The pilot studies included stakeholders from the finance and education sectors only. Although the typology has been designed to be domain-agnostic, evaluation in other sectors would be useful to test the comprehensiveness of the typology and refine or adapt its dimensions and sub-categories for other verticals. More work is also needed to carefully derive explanation requirements from all relevant obligations found in the EU AI Act.

Importantly, despite the highly integrated process that was followed to produce explanation requirements through leveraging the typology, the process remains a manual one, which cannot be fully automated, particularly when the process starts with a translation of primary requirements stemming from applicable laws. Hence, the importance of the role of the Legal Engineer, who will have to interpret high-level requirements and transform them into actionable requirements, while acknowledging possible reductions. Among the different dimensions and sub-dimensions, the Minimum Content one, in particular, requires careful analysis of the high-level requirements, which will then have to be assessed in the light of technical capabilities. Crucially, it is worth noting that computable fine-grained explanation requirements, which can be opinionated, do not necessarily exhaust higher level requirements. By way of example, the provenance trail that may be generated to achieve accuracy is likely to be focused upon integrity, yet data subjects have a right under GDPR Article 15 to access data to ensure that their personal data is correct as well as to verify the lawfulness of the processing.<sup>40</sup> In any case, although this research was constrained by the co-creation process itself, it demonstrated that refusals to explain underlying data processing based on the complexity of the processing are not justified and that a uniform approach can be adopted for both purely and partially automated decision-making pipelines.

More generally, an organisation wanting to operationalise the production of explanations based on the typology introduced in this paper should carefully map applicable explanation requirements to the nine dimensions of the typology. As a result, the quality of the classification will largely depend on the quality of the mapping.

What is more, the quality of the explanations produced on the basis of the typology will also depend upon the quality of the description of the decision-making pipeline and its output. By way of example, the organisation implementing the typology will want to ensure it is able to distinguish between public and restricted explanations. In other words, determining when an explanation can or should be disclosed to third parties is a decision for the implementing organisation to be made on the basis of an internal classification. As mentioned in Section 4, the GDPR allows data

<sup>39</sup>AI Act, Article 10(3).

<sup>40</sup>See e.g., CJEU C-487/21 FF v Österreichische Datenschutzbehörde and CRIF GmbH, para. 34.

controllers to carefully balance the amount of information they disclose to individuals with the protection of their commercial interests as long as they do not undermine data subject rights.<sup>41</sup> In these cases, it is up to the implementing organisation to decide which information should be included in an explanation and how best to use the Confidentiality sub-category of the Content dimension to control disclosure.

Another takeaway relates to the necessity to work with rich input to be able to provide actionable output [53].

Finally, a significant challenge in the context of explanation computation is trustworthiness. It is essential to define metrics that can be used to assess the integrity of these explanations and ensure they have not been tampered with.

## 7 CONCLUSION

This paper presented a typology built to categorise explanation requirements that can be leveraged to compute explanations. The typology comprises nine dimensions, with each dimension forming a set of building blocks for the generation of explanations. This typology is the product of both conceptual and empirical research and has been co-created with industry partners, who are both consumers and producers of explanations. The feasibility of its construction for the purpose of computing explanations has been tested in the context of two small-scale pilot studies. The usefulness of the outputs produced on the basis of the typology for producers and consumers operating within the context of these use cases has also been tested. While Section 3 unpacks the nine dimensions of the typology, Section 4 describes the co-creation process that was followed and includes examples of explanations produced in the context of two scenarios. Section 5 discusses the evaluation of the pilot studies and the importance of a new role, that of the legal engineer, of which main characteristics are broken down and further specified. Driven by a responsible innovation approach, we found that adding an interdisciplinary role, that of the legal engineer, was an effective way to bridge the work of the legal, compliance, and engineering teams and, crucially, acknowledge the limits of automation. The legal engineer was responsible for translating high-level requirements received from legal and compliance teams into fine-grained explanation requirements to be consumed by the engineering team. The legal engineer was subsequently embedded within the engineering team and regularly consulted during the technical design phase to help the data engineer specify the requirements for the software architecture components.

The research led to the development of a light vocabulary based on the typology, included in Appendix A and an explanation assistant, a software service configured to support a set of explanation requirements. The light vocabulary provides a machine-readable format for classifying fine-grained explanation requirements. It can be combined with existing compliance or auditing ontologies, such as, for example, the consent ontology developed by Pandit et al. [43].

Our next steps will be to perform a comprehensive evaluation process and add new use cases to the two use cases mentioned above as well as report more extensively on the technical design phase.

## 8 ACKNOWLEDGMENTS

The work presented in this article has been supported by the UK Engineering and Physical Sciences Research Council (EPSRC) via the Grants [EP/S027238/1] and [EP/S027254/1] for the PLEAD project. The authors have no competing interests to declare that are relevant to the content of this article.

## REFERENCES

- [1] . [n. d.]. Directive (EU) 2023/2225 of the European Parliament and of the Council of 18 October 2023 on credit agreements for consumers and repealing Directive 2008/48/EC OJ L, 2023/2225, 30.10.2023.

---

<sup>41</sup>See GDPR Article 15(4).

- [2] Conference of the Independent Data Protection Supervisory Authorities of the Federation and the Länder. 2022. Standard Data Protection Model v.3. [https://www.datenschutz-mv.de/static/DS/Dateien/Datenschutzmodell/SDM\\_V3\\_en.pdf/](https://www.datenschutz-mv.de/static/DS/Dateien/Datenschutzmodell/SDM_V3_en.pdf/). Accessed: 15 May 2024.
- [3] Alejandro Barredo Arrieta, Natalia Diaz-Rodriguez, Javier Del Ser, Adrien Bernetot, Siham Tabik, Alberto Barbado, Salvador Garcia, Sergio Gil-Lopez, Daniel Molina, Richard Benjamins, Raja Chatila, and Francisco Herrera. 2020. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion* 58 (jun 2020), 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- [4] Article 29 Data Protection Working Party. [n. d.]. <https://ec.europa.eu/newsroom/article29/items/612053>.
- [5] Solon Barocas, Andrew D. Selbst, and Manish Raghavan. 2019. The Hidden Assumptions Behind Counterfactual Explanations and Principal Reasons. (2019). <https://doi.org/10.48550/ARXIV.1912.04930>
- [6] Cesare Bartolini, Robert Muthuri, and Cristiana Santos. 2015. Using Ontologies to Model Data Protection Requirements in Workflows. In *JSAI-isAI Workshops*.
- [7] Dimitri Bohlender and Maximilian A. Köhl. 2019. Towards a Characterization of Explainable Systems. <https://doi.org/10.48550/ARXIV.1902.03096> arXiv:1902.03096
- [8] Céline Castets-Renard. 2021. *Human Rights and Algorithmic Impact Assessment for Predictive Policing*. Cambridge University Press, 93–110. <https://doi.org/10.1017/9781108914857.007>
- [9] Ajay Chander and Ramya Srinivasan. 2018. Evaluating Explanations by Cognitive Value. In *Lecture Notes in Computer Science*. Springer International Publishing, 314–328. [https://doi.org/10.1007/978-3-319-99740-7\\_23](https://doi.org/10.1007/978-3-319-99740-7_23)
- [10] Larissa Chazette, Oliver Karras, and Kurt Schneider. 2019. Do End-Users want explanations? Analyzing the role of explainability as an emerging aspect of non-functional requirements. In *Proceedings of the IEEE International Conference on Requirements Engineering*, Vol. 2019-September. IEEE Computer Society, 223–233. <https://doi.org/10.1109/RE.2019.00032>
- [11] Larissa Chazette and Kurt Schneider. 2020. Explainability as a non-functional requirement: challenges and recommendations. *Requirements Engineering* 25, 4 (dec 2020), 493–514. <https://doi.org/10.1007/s00766-020-00333-1>
- [12] Jennifer Cobbe. 2019. Administrative law and the machines of government: judicial review of automated public-sector decision-making. *Legal Studies* 39, 4 (jul 2019), 636–655. <https://doi.org/10.1017/lst.2019.9>
- [13] Jennifer Cobbe and Jatinder Singh. 2020. Reviewable Automated Decision-Making. *Computer Law & Security Review* 39 (2020), 105475. <https://doi.org/10.1016/j.clsr.2020.105475>
- [14] Datenschutzbehörde. 2020. Case DSB-D124.909, ECLI:AT:DSB:2020:2020.0.436.002.
- [15] Mina Deng, Kim Wuyts, Riccardo Scandariato, Bart Preneel, and Wouter Joosen. 2010. A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. *Requirements Engineering* 16, 1 (nov 2010), 3–32. <https://doi.org/10.1007/s00766-010-0115-7>
- [16] Department of Education. 2021. School Admissions Code. (Mandatory requirements and statutory guidance for admission authorities, governing bodies, local authorities, schools adjudicators and admission appeals panels), <https://www.gov.uk/government/publications/school-admissions-code--2>. Accessed: 15 October 2023.
- [17] Lilian Edwards and Michael Veale. 2017. Slave to the Algorithm? Why a Right to Explanation is Probably Not the Remedy You are Looking for. *SSRN Electronic Journal* (2017). <https://doi.org/10.2139/ssrn.2972855>
- [18] Upol Ehsan and Mark O. Riedl. 2020. Human-Centered Explainable AI: Towards a Reflective Sociotechnical Approach. In *HCI International 2020 - Late Breaking Papers: Multimodality and Intelligence (Lecture Notes in Computer Science, Vol. 12424 LNCS)*, Constantine Stephanidis, Masaaki Kurosu, Helmut Degen, and Lauren Reinerman-Jones (Eds.). Springer International Publishing, 449–466. [https://doi.org/10.1007/978-3-030-60117-1\\_33](https://doi.org/10.1007/978-3-030-60117-1_33)
- [19] European Union. 2016. Regulation 2016/679 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). *Official Journal of the European Union* 59, L 199 (2016), 1–88.
- [20] European Union. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). <http://data.europa.eu/eli/reg/2024/1689/oj>. OJ L, 2024/1689, 12.7.2024.
- [21] Katharine Freeland. 2020. Legal Engineer. *The Law Society Gazette*, <https://www.lawgazette.co.uk/analysis/legal-engineer/5104088.article>. Accessed: 25 October 2023.
- [22] Mireille Hildebrandt. 2017. Saved by Design? The Case of Legal Protection by Design. *NanoEthics* 11, 3 (aug 2017), 307–311. <https://doi.org/10.1007/s11569-017-0299-0>
- [23] Mireille Hildebrandt. 2023. Qualification and Quantification in Machine Learning. From Explanation to Explication. *Sociologica* (2023), Vol. 16 No. 3 (2022). <https://doi.org/10.6092/ISSN.1971-8853/15845>
- [24] Trung Dong Huynh, Niko Tsakalakis, Ayah Helal, Sophie Stalla-Bourdillon, and Luc Moreau. 2021. Addressing Regulatory Requirements on Explanations for Automated Decisions with Provenance—A Case Study. *Digital Government: Research and Practice* 2, 2 (2021), 1–14. <https://doi.org/10.1145/3436897>
- [25] Trung Dong Huynh, Niko Tsakalakis, Ayah Helal, Sophie Stalla-Bourdillon, and Luc Moreau. 2023. A Methodology and Software Architecture to Support Explainability-by-Design. In press. <https://arxiv.org/abs/2206.06251>
- [26] Information Commissioner’s Office. 2018. Automated decision-making and profiling. v. 1.1.62, <https://ico.org.uk/media/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/automated-decision-making-and-profiling-1-1.pdf>. Accessed: 15 October



- 2023.
- [27] Information Commissioner's Office. 2022. Explaining decisions made with AI, Part 2: Explaining AI in practice. <https://ico.org.uk/for-organisations/guide-to-data-protection/key-dp-themes/explaining-decisions-made-with-artificial-intelligence/part-2-explaining-ai-in-practice/>. Accessed: 15 October 2023.
- [28] ISO/IEC 27001:2013(E) 2013. *Information technology – Security techniques – Information security management systems – Requirements*. Standard. International Organization for Standardization, Geneva, CH.
- [29] Tae Wan Kim and Bryan R. Routledge. 2022. Why a Right to an Explanation of Algorithmic Decision-Making Should Exist: A Trust-Based Approach. *Business Ethics Quarterly* 32, 1 (2022), 75–102. <https://doi.org/10.1017/beq.2021.3>
- [30] Clément Labadie and Christine Legner. 2019. Understanding Data Protection Regulations from a Data Management Perspective: A Capability-Based Approach to EU-GDPR. In *Wirtschaftsinformatik*.
- [31] Q. Vera Liao and Kush R. Varshney. 2021. Human-Centered Explainable AI (XAI): From Algorithms to User Experiences. <https://doi.org/10.48550/ARXIV.2110.10790> arXiv:2110.10790 [cs.AI]
- [32] Gianclaudio Malgieri and Giovanni Comandé. 2017. Why a Right to Legibility of Automated Decision-Making Exists in the General Data Protection Regulation <https://doi.org/10.1017/ist.2019.9> <https://doi.org/10.1017/ist.2019.9>. *International Data Privacy Law* 7, 4 (nov 2017), 243–265. <https://doi.org/10.1093/idpl/ix019>
- [33] Yod-Samuel Martin and Antonio Kung. 2018. Methods and Tools for GDPR Compliance Through Privacy and Data Protection Engineering. In *2018 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE. <https://doi.org/10.1109/eurospw.2018.00021>
- [34] Mark T. Maybury. 1992. Communicative acts for explanation generation. *International Journal of Man-Machine Studies* 37, 2 (1992), 135–172. [https://doi.org/10.1016/0020-7373\(92\)90083-w](https://doi.org/10.1016/0020-7373(92)90083-w)
- [35] Ljubiša Metikoš and Jef Ausloos. Forthcoming. The Right to an Explanation in Practice: Insights from Case Law for the GDPR and the AI Act. *Law, Innovation, and Technology* 2, 17 (Forthcoming).
- [36] George A. Miller. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review* 63, 2 (1956), 81–97. <https://doi.org/10.1037/h0043158>
- [37] Tim Miller. 2017. Explanation in Artificial Intelligence: Insights from the Social Sciences. <https://doi.org/10.48550/arXiv.1706.072699>
- [38] Brent Mittelstadt, Chris Russell, and Sandra Wachter. 2018. Explaining Explanations in AI. <https://doi.org/10.1145/3287560.3287574>
- [39] Luc Moreau and Paolo Missier. 2013. PROV-DM: The PROV datamodel. World Wide Web Consortium, <https://www.w3.org/TR/2013/REC-prov-dm-20130430/>. Accessed: 15 October 2023.
- [40] Robert Nickerson, Upkar Varshney, and Jan Muntermann. 2013. A method for taxonomy development and its application in information systems. *European Journal of Information Systems* 22, 3 (2013), 336–359. <https://doi.org/10.1057/ejis.2012.26>
- [41] Nicolas Notario, Alberto Crespo, Yod-Samuel Martin, Jose M. Del Alamo, Daniel Le Metayer, Thibaud Antignac, Antonio Kung, Inga Kroener, and David Wright. 2015. PRIPARE: Integrating Privacy Best Practices into a Privacy Engineering Methodology. In *2015 IEEE Security and Privacy Workshops*. IEEE. <https://doi.org/10.1109/spw.2015.22>
- [42] Richard Owen, Jack Stilgoe, Phil Macnaghten, Mike Gorman, Erik Fisher, and Dave Guston. 2013. *Responsible Innovation: Managing the responsible emergence of science and innovation in society*. Wiley, Chapter A Framework for Responsible Innovation, 27–50. <https://doi.org/10.1002/9781118551424.ch2>
- [43] Harshvardhan J. Pandit, Christophe Debruynne, Declan O'Sullivan, and Dave Lewis. 2019. *GConsent - A Consent Ontology Based on the GDPR*. Springer International Publishing, Cham, 270–282.
- [44] Franck Pasquale and Gianclaudio Malgieri. 2024. Generative AI, Explainability, and Score-Based Natural Language Processing in Benefits Administration. *Journal of Cross-disciplinary Research in Computational Law* 2, 2 (2024), 1–13.
- [45] Persónuvernd . 2020. Vinnsla Creditinfo Lánstrausts hf. á persónuupplýsingum í tengslum við gerð lánshæfismats og aðgangs- og upplýsingaréttur vegna gerðar skýrslna um lánshæfismat, Mál nr. 2020010592.
- [46] Wolter Pieters. 2010. Explanation and trust: what to tell the user in security and AI? *Ethics and Information Technology* 13, 1 (nov 2010), 53–64. <https://doi.org/10.1007/s10676-010-9253-3>
- [47] Alun Preece. 2018. Asking 'Why' in AI: Explainability of intelligent systems – perspectives and challenges. *Intelligent Systems in Accounting, Finance and Management* 25, 2 (apr 2018), 63–72. <https://doi.org/10.1002/isaf.1422>
- [48] Rb. Amsterdam. 2021. Ola Driver App. C/13/689705/HA RK 20-258, NL:RBAMS:2021:1019.
- [49] Wojciech Samek, Thomas Wiegand, and Klaus-Robert Müller. 2017. Explainable Artificial Intelligence: Understanding, Visualizing and Interpreting Deep Learning Models. <https://doi.org/10.48550/arXiv.1708.08296>
- [50] Andrew D Selbst and Julia Powles. 2017. Meaningful information and the right to explanation. *International Data Privacy Law* 7, 4 (2017), 233–242. <https://doi.org/10.1093/idpl/ix022> arXiv:<https://academic.oup.com/idpl/article-pdf/7/4/233/22923065/ix022.pdf>
- [51] Auste Simkute, Ewa Luger, Bronwyn Jones, Michael Evans, and Rhianne Jones. 2021. Explainability for experts: A design framework for making algorithms supporting expert decisions more explainable. *Journal of Responsible Technology* 7-8 (2021), 100017. <https://doi.org/10.1016/j.jrt.2021.100017>
- [52] Kacper Sokol and Peter Flach. 2020. Explainability fact sheets: a framework for systematic assessment of explainable approaches. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (Barcelona, Spain) (FAT\* '20)*. Association for Computing Machinery, New York, NY, USA, 56–67. <https://doi.org/10.1145/3351095.3372870>

- [53] Marco Antonio Solano. 2019. *Information Quality: The Nexus of Actionable Intelligence*. Springer International Publishing, Cham, 471–517. [https://doi.org/10.1007/978-3-030-03643-0\\_20](https://doi.org/10.1007/978-3-030-03643-0_20)
- [54] Emily Sullivan. 2024. SIDEs: Separating Idealization from Deceptive Explanations in xAI. <https://doi.org/10.48550/arXiv.2404.16534>
- [55] Niko Tsakalakis, Sophie Stalla-Bourdillon, Laura Carmichael, Dong Huynh, Luc Moreau, and Ayah Helal. 2022. The Dual Function of Explanations: Why Computing Explanations is of Value. In *Data Protection and Privacy*. Hart Publishing. <https://doi.org/10.5040/9781509954544.ch-005>
- [56] Sebastiao Barros Vale and Gabriela Zanfir-Fortuna. 2022. *Automated Decision-Making under the GDPR – A Comprehensive Case-Law Analysis - FPF Report*. Report. <https://fpf.org/blog/fpf-report-automated-decision-making-under-the-gdpr-a-comprehensive-case-law-analysis/>
- [57] Sandra Wachter, Brent Mittelstadt, and Luciano Floridi. 2017. Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation. *International Data Privacy Law* 7, 2 (2017), 76–99. <https://doi.org/10.1093/idpl/ix005>
- [58] Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2017. Counterfactual Explanations without Opening the Black Box: Automated Decisions and the GDPR. (2017). <https://doi.org/10.48550/ARXIV.1711.00399>
- [59] Lingxue Yang, Hongrun Wang, and Léa A. Deleris. 2021. What Does It Mean to Explain? A User-Centered Study on AI Explainability. In *Artificial Intelligence in HCI (Lecture Notes in Computer Science, Vol. 12797 LNAI)*, Helmut Degen and Stavroula Ntoa (Eds.). Springer International Publishing, 107–121. [https://doi.org/10.1007/978-3-030-77772-2\\_8](https://doi.org/10.1007/978-3-030-77772-2_8)
- [60] L. Richard Ye and Paul E. Johnson. 1995. The Impact of Explanation Facilities on User Acceptance of Expert Systems Advice. *MIS Quarterly* 19, 2 (jun 1995), 157. <https://doi.org/10.2307/249686>
- [61] Jianlong Zhou, Amir H. Gandomi, Fang Chen, and Andreas Holzinger. 2021. Evaluating the Quality of Machine Learning Explanations: A Survey on Methods and Metrics. *Electronics* 10, 5 (2021), 336–359. <https://doi.org/10.3390/electronics10050593>
- [62] Jichen Zhu, Antonios Liapis, Sebastian Risi, Rafael Bidarra, and G. Michael Youngblood. 2018. Explainable AI for Designers: A Human-Centered Perspective on Mixed-Initiative Co-Creation. In *2018 IEEE Conference on Computational Intelligence and Games (CIG)*. IEEE, 1–8. <https://doi.org/10.1109/CIG.2018.8490433>

## APPENDIX A

The Explainability-by-Design typology is encoded as a lightweight vocabulary encoded as an OWL2 ontology. Its namespace is <https://openprovenance.org/ns/plead>. It is designed to be as simple as possible, consisting of 10 top-level classes and 9 object properties. The 10 classes model the concept of Explanation and its associated 9 dimensions (Fig. 14). The 9 object properties have Explanation as their domain and the corresponding dimension class as their range (Fig. 15).

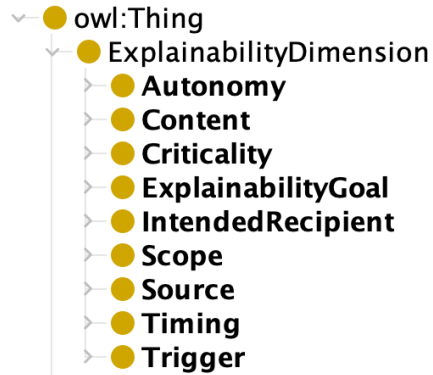


Fig. 14. Top level classes of PLEAD Explainability-by-Design typology

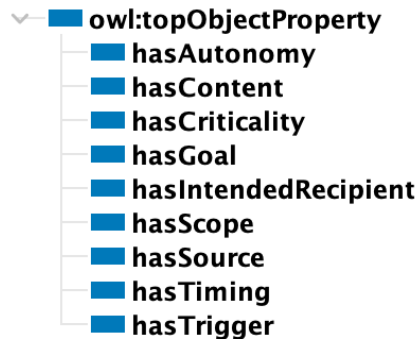


Fig. 15. Object properties of PLEAD Explainability-by-Design typology

An instance of explanation (such as Existence of automated decision-making in Table 2) can be expressed as set of RDF statements, as in Listing 1.

---

```

1  ### https://openprovenance.org/ns/plead#explain1
2  plead:explain1 rdf:type owl:NamedIndividual ,
3    plead:ExistenceAutomatedDecisionMaking ;
4    plead:hasAutonomy plead:proactive ;
5    plead:hasContent plead:generalised ,
6      plead:disclosed ,
7      plead:minimum1 ;
8    plead:hasGoal plead:fairness1 ,
9      plead:accountability1 ,
10     plead:overall_transparency1 ;
11   plead:hasIntendedRecipient plead:consumer1 ;
12   plead:hasTiming plead:ex_ante ;
13   plead:hasCriticality plead:mandatory ;
14   plead:hasScope plead:universal ;
15   plead:hasSource plead:implicit_art_22 ;
16   plead:hasTrigger plead:loan_application ;
17   plead:example "To be able to process applications quickly and
18     accurately, the Bank uses an automated system." .
  
```

---

Listing 1. Explanation instance in RDF

**APPENDIX B**

Table 5 is listing examples out of the pilot studies, where the study participants validated proposed explanations for the two use cases. The excerpts are correlated in the table to their assigned 'base', i.e. identified themes that were then used to form explanation requirements.

Received 11 December 2023; revised 16 September 2024; accepted 12 November 2024

ID	Requirement	Example excerpt(s)
RQ1	Explanations must comply with a legal obligation	“Yeah, well [...] we provide the query process [because] there is effectively law that says the bank has to tell the customer which CRA was used.”
RQ2	Explanations must comply with company policy	“Now that CRA data goes into the bank and they make a decision based on their policy, so surely you need that angle with some nicely worded [explanation]” “For me, you’re quite right the policy piece is important”
RQ3	Explanations must address different contexts	“I think that would be really, really handy for our contact centre as well.”
RQ4	Explanations must be scalable and re-usable	“I think it would minimize the amount of traffic coming through [the call centre] as well”
RQ5	Explanations must be produced regardless of data subject participation	“we kind of don’t have to dig for that information unless obviously you know someone creates [a request]. We wouldn’t actually look into [the explanation] ourselves. So having [the explanation] there would be quite useful.” “... as members of the team [we] wouldn’t actually be able to have that to hand necessarily [w/out the explanation]”
RQ6	Explanations must convey the necessary information fast	“I think it would minimize the amount of traffic coming through [the call centre] as well”
RQ7	Explanations must empower the recipient	“Definitely adds value. Like we said, it’s kind of putting the ownership back into the hands of the consumer so that they can actually see physically.” “It would actually be accessible to them [the data subjects] quite quickly, and so it would probably alleviate some of the calls that we would be getting, and it would explain in more detail exactly how that decision was made.” “it kind of comes back to your tool, which I quite like because then that kind of gives that extra layer to say ‘OK, you’ve seen your credit report from the CRAs. You kind of know what the bank is doing from this tool to say what’s actually going to be influencing and what’s not.’ And then you’ve got clear kind of guidance or text to say that this is actually what’s causing the impact on your credit score or the decision making.” “Yeah, the information you presented it might help in some regards in terms of ‘oh, actually I’ve misstated this!’ ” “That is a lot more education and a lot more transparent.”
RQ9	Explanations must increase the engagement with the product	“I think you get a lot more engagement and as a result I think you could put forward a very plausible case.”
RQ10	Explanations must assist in the transparency obligations	“That is a lot more education and a lot more transparent.”
RQ11	Explanations must respect intellectual property/trade secrets	“I think in detail, no, I think that you know that that would be bad for the lenders [to reveal too much]. [...] you know it is a little bit like asking Coca Cola to print the formula on the outside of every bottle. I mean it just kind of, you know, you kind of break the model.” “I mean [name], that would be fantastic, but I mean, I guess finance houses will be reluctant to share their lending policy rules. Would they?”
RQ12	Explanations must be fit for purpose	“I think that actually even though it’s a lot longer, it gives a lot more information.”

Table 5. Example excerpts from the pilot studies.