

Received 25 May 2022, accepted 27 June 2022, date of publication 6 July 2022, date of current version 20 July 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3188967

RESEARCH ARTICLE

Using Knowledge Graphs to Unlock Practical Collection, Integration, and Audit of AI Accountability Information

IMAN NAJA¹, MILAN MARKOVIC¹, PETER EDWARDS¹, WEI PANG²,
CAITLIN COTTRILL³, AND REBECCA WILLIAMS⁴

¹Department of Computing Science, University of Aberdeen, Aberdeen AB24 3UE, U.K.

²School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh EH14 4AS, U.K.

³School of Engineering and Centre for Transport Research, University of Aberdeen, Aberdeen AB24 3UE, U.K. (Deceased)

⁴Faculty of Law and Pembroke College, University of Oxford, Oxford OX1 1DW, U.K.

Corresponding author: Milan Markovic (milan.markovic@abdn.ac.uk)

This work was supported by the Realising Accountable Intelligent Systems (RAInS) Project through the UK Research and Innovation (UKRI) Digital Economy Program under Award EP/R033846/1 and Award EP/R03379X/1.

ABSTRACT To enhance trustworthiness of AI systems, a number of solutions have been proposed to document how such systems are built and used. A key facet of realizing trust in AI is how to make such systems accountable - a challenging task, not least due to the lack of an agreed definition of accountability and differing perspectives on what information should be recorded and how it should be used (e.g., to inform audit). Information originates across the life cycle stages of an AI system and from a variety of sources (individuals, organizations, systems), raising numerous challenges around collection, management, and audit. In our previous work, we argued that semantic Knowledge Graphs (KGs) are ideally suited to address those challenges and we presented an approach utilizing KGs to aid in the tasks of modelling, recording, viewing, and auditing accountability information related to the design stage of AI system development. Moreover, as KGs store data in a structured format understandable by both humans and machines, we argued that this approach provides new opportunities for building intelligent applications that facilitate and automate such tasks. In this paper, we expand our earlier work by reporting additional detailed requirements for knowledge representation and capture in the context of AI accountability; these extend the scope of our work beyond the design stage, to also include system implementation. Furthermore, we present the *RAInS* ontology which has been extended to satisfy these requirements. We evaluate our approach against three popular baseline frameworks, namely, Datasheets, Model Cards, and FactSheets, by comparing the range of information that can be captured by our KGs against these three frameworks. We demonstrate that our approach subsumes and extends the capabilities of the baseline frameworks and discuss how KGs can be used to integrate and enhance accountability information collection processes.

INDEX TERMS Accountability, AI systems, machine learning, ontology, provenance.

I. INTRODUCTION

The increasingly widespread adoption of AI systems has been spurred by their stated advantages and benefits, including automation of labour, reduction in error rates vs. humans

The associate editor coordinating the review of this manuscript and approving it for publication was Chang Choi.

performing the same tasks, and always on availability. However, such systems can also suffer from a range of drawbacks, including biases which result in perpetuating racism and sexism [1]–[8], the production of erroneous or unintended results [9]–[11], and in extreme cases, incidents that result in direct or indirect harm to humans [12], [13]. Such limitations remain the subject of an active academic debate [14]–[17],

and there are emerging efforts to track incidents related to AI-powered technologies, such as the AI Incident Database.¹ A variety of solutions have been proposed to address the perceived limitations of AI systems; these have originated from regulators and policy makers [18]–[21], professional bodies [22], [23], and developers and researchers [24]–[27]. A key facet of realizing trust in AI is how best to make such systems accountable - a challenging task, not least because of the lack of a universally agreed definition of accountability. Despite emerging academic debates about treating AI systems as self-aware entities and therefore as legal agents [28], we base the approach outlined in this paper on the assumption that AI systems, algorithms, or machines should not be treated as legally responsible or liable agents [29], [30], and that as a result, human agents or organizations that created or influenced the behaviors of AI systems should be accountable. As a consequence, we argue that AI accountability must be supported by the ability to inspect, review, or otherwise interrogate such a system with the goals of (i) making processes associated with each of its life cycle stages transparent [19], [22]–[24], [26], [31]; (ii) demonstrating compliance with hard laws (i.e., laws and regulations) and soft laws (i.e., standards and guidelines) [19], [26]; and (iii) aiding investigations into the cause(s) of failure or erroneous decisions and supporting the identification of responsible parties [19], [22], [24], [26], [31].

The various technologies that could enable such comprehensive inspection, revision, or interrogation are often explored by the computer science (CS) community and professional bodies under separate sub-themes [32], such as transparency and traceability [23], [33], [34], intelligibility/interpretability [23], [34], explainability [31], [34], and auditability/reviewability [35]. But due to the socio-technical nature of most AI systems, these themes are also explored by lawyers, regulators, and social sciences researchers [18], [19]. This often leads to a “mismatch” of expectations and proposed solutions. For example, proposed laws and regulations may be perceived as technically unfeasible and lacking in providing clear means for implementation, e.g., GDPR [36] requires the provision of “meaningful information about the logic involved” without specifying what ‘meaningful information’ is. Moreover, technological solutions proposed by the CS community may be difficult to interpret and may lack crucial evidence (e.g., information about involved human actors) in order to be usable in legal proceedings. Further, there is a stance within the CS community that calls to make AI systems more transparent by exposing their source code would, in many cases, hinder rather than help accountability, as such transparency is often not meaningful, because it requires significant efforts and skills to interpret as well as posing other challenges (e.g., those related to the protection of intellectual property rights) [33], [37].

We argue that solutions which aim to support accountability in AI systems which are also enforceable by law must be

comprehensive and cover all life cycle stages of such systems, and at each stage must identify individual human agents that bear responsibility for decisions and outcomes (e.g., implemented system components) influencing the design, implementation, deployment, and operation of these systems. In this paper, we refer to this type of information as *accountability information*. We further argue that this accountability information must be meaningful, semantically annotated and ideally understandable by both humans and machines. The effort needed to collect and to interpret such data in case of an investigation is likely to be a significant burden on many stakeholders, and the ability of machines to understand and process the data is important in order to achieve a certain level of automation.

Our proposed approach is compatible with existing solutions such as Datasheets [38], Model Cards [39], and FactSheets [40]. These frameworks provide guidelines on what information to record about various aspects of AI systems (e.g., descriptions of datasets, machine learning (ML) models, performance evaluations, etc.) to improve transparency of such systems. Datasheets and FactSheets provide questions and checklists that users can refer to in order to record free text and images describing datasets or AI systems. While examples for different types of systems are provided, it is ultimately up to the end users of these frameworks to decide which questions are relevant in the context of their AI systems and the final result is presented in the form of a visual report designed for human inspection. While this simplicity provides certain benefits (e.g., a low barrier to the production and distribution of such reports), the lack of structured, semantically annotated data makes it difficult for computer systems to assist with creating and querying such reports in an automated or semi-automated manner. For example, it would be difficult for software to check if some information that was intended to be collected (e.g., answer to some critical question) was indeed provided. Model Cards do provide means to assist developers with the collection of required information via their Model Card Toolkit,² which produces JSON representations of Model Cards; however, interpretation and querying of such data again relies on visual inspection by humans.

In our approach, the term AI system refers to software comprising ‘core AI’ components (e.g., an ML model) and other supporting functions (e.g., API wrappers) [41] allowing it to function either as a standalone solution, or as a part of a larger system. We consider the development and use of an AI system in terms of four high-level life cycle stages: *Design*, *Implementation*, *Deployment*, and *Operation*; this conforms to the recommendation by Amershi et al. [42] that standard software engineering practices should apply to such systems. The *design* stage covers the aspects associated with designing the AI system; the *implementation* stage includes the processes associated with building and testing the system; the *deployment* stage includes the processes related to installing the system and, if applicable,

¹<https://incidentdatabase.ai/>

²<https://github.com/tensorflow/model-card-toolkit>

configuring and integrating it with other systems, producing documentation, and training users. Finally, the *operation* stage consists of the actual use of the system and (routine) monitoring. This provides a broader scope for recording accountability information, which extends the current state of the art. Modelling accountability information is supported through two ontologies, *SAO*³ and *RAInS*.⁴ These ontologies are built around three main concepts: *accountability plans*, *accountability traces* and *accountable agents*. Accountability plans specify information that should be captured throughout the system's life cycle and guide the recording of corresponding accountability traces. We argue that the ability to plan the collection of accountability information is crucial for scalability of such data capture and to ensure that the required information is not lost (e.g., in order to save sensor observations in an autonomous vehicle this functionality must be planned for and implemented prior to an incident). The actual descriptions of activities that occurred during the different life cycle stages, their inputs and outputs and relationships to accountable human agents are captured in corresponding accountability traces. *SAO* and *RAInS* extend the W3C PROV-O [43] standard for representing provenance as causal graphs consisting of three core concepts: *entities*, which describe any real or imaginary thing; *activities*, which represent processes that use and generate entities; and *agents*, which represent human, organizational, or software actors that can bear some responsibility for an activity. The ontologies are also based on the EP-Plan [44] ontology, which extends PROV-O with the ability to describe abstract acyclic plans represented as a series of steps and input and output variables. The capture and audit of accountability information is supported by a prototype implementation of the *Accountability Fabric*, a suite of tools demonstrating the feasibility of a computational solution to support planning, collection, and inspection of accountability information using *SAO* and *RAInS*. For technical details describing the ontologies and the *Accountability Fabric*, we refer the reader to [45], [46], which demonstrate the application of this approach to documenting the design stage of an AI system.

In this paper we present the following contributions:

- 1) A set of competency questions describing types of accountability information relevant to the design and implementation stages of an AI system, and which goes beyond the information recorded by the current Datasheet, Model Card, and FactSheet frameworks.
- 2) The extended *RAInS* ontology for modelling accountability information as knowledge graphs (KGs) which covers the design and implementation stages of the AI system life cycle.
- 3) An evaluation of our approach by comparing its capabilities with those of Datasheets, Model Cards, and FactSheets.

³<https://w3id.org/sao>

⁴<https://w3id.org/rains>

The remainder of this paper is organized as follows: Section II describes a set of competency questions which define what accountability information needs to be recorded; Section III describes our approach for semantically modelling such accountability information; Section IV describes results from a comparison of our solution to the three popular non-semantic frameworks; Section V discusses the results of the comparison and outlines the benefits of our approach; and Section VI concludes the paper and discusses future work.

II. ACCOUNTABILITY INFORMATION COVERAGE

Competency questions (CQs) are expressed in natural language and are commonly used in semantic modelling to identify individual ontological concepts (e.g., types of objects and their relationships) and to evaluate the resulting ontology (i.e., the ontology must describe information that can sufficiently answer CQs) [47]. In this section, we summarize our approach for creating the CQs relating to accountability of AI systems, and provide an overview of their thematic structure. It should be noted that in *RAInS* we focus on AI systems built around ML technologies. The content presented here includes questions relating to the system design stage (previously discussed in [45]) and those relating to the implementation stage of an AI system. While these two stages are closely coupled and hence many CQs overlap, one of the main differences is that the design stage includes descriptions of the design specifications of the system components, whereas the implementation stage includes descriptions of the tangible system components based on such specifications, as well as justifications of any deviations from the original design. For example, the design stage includes the specification of how an ML model is to be built and evaluated, whereas the implementation stage includes descriptions of the ML model after it has been realized, and the CQ 'What are the limitations of the ML model?' can be read in each case as: 'What are the limitations of the designed ML model?' and 'What are the limitations of the realized ML model?'

A. METHODOLOGY

We gathered CQs relating to accountability of AI systems from publications by regulatory bodies [19], [48], statements and guidelines released by professional bodies [22], [31], [49], and the academic literature [24], [26], [38]–[40], [50] up to the year 2020. The questions included documentation requirements for AI system components (e.g., ML models) as well as explanations of how automated decisions were made. While the literature offers a comprehensive set of questions, these often focus on technological details of AI systems, their individual components (e.g., ML models) and resources used in the development process (e.g., datasets), often without explicit links to human agents who could be held accountable for their existence. The initial set of CQs was analyzed and similar questions (such as those addressing the same issue but worded differently) were grouped under overarching thematic CQs which represented those questions. For example, the two requirements to indicate "*whether the model was developed*

with general or specific tasks in mind” [39] or “What are the goals, purposes and intended applications of the product or service?” [49] were translated to the CQ: ‘What are the intended use cases of the ML model?’ Each thematic CQ was then analyzed to determine the ontological concepts needed to represent that information. For example, *IntendedUse-Case* was one of the concepts that emerged from this CQ. Having completed the initial CQ analysis exercise, we next turned our attention to a number of AI application use cases (autonomous vehicles, image classifiers for cancer detection, and loan application approval systems), with the objective of identifying any missing knowledge requirements. We created additional CQs in response to these use cases which extended the discussion of AI accountability in the current literature; we then determined the ontological concepts necessary to represent the information in the new CQs.

The CQs identified through this exercise that were not included in the initial set of questions were as follows:

- Whether the implementation or evaluation followed the design specification, and if not, what the justification was for the deviation from the specification.
- Who performed each accountable action and was therefore responsible for the generated accountable result(s).
- What were the compliance requirements for the components of the AI system.
- Who checked that the components of the AI system comply with relevant hard and soft laws.
- Who checked that the components of the AI system are fit for purpose, i.e., its intended use case(s).
- Who approved, i.e., signed off on, the components of the AI system.
- What were the guidelines for deploying and operating the AI systems (to be disseminated to the persons responsible for deploying the AI system).

B. COMPETENCY QUESTIONS TOPICS

The final set of CQs that framed the scope of our knowledge representation models contained 127 questions. We organized and grouped these CQs under a number of high-level topics as shown below. Where a topic relates to the *Design* stage, it is indicated by the letter (D), to the *Implementation* stage by the letter (I), and to both *Design* and *Implementation* stages by the letters (D&I). References are provided if questions relating to particular aspects of a topic came from a third-party resource. Where no reference is provided, these were created by our team as outlined in the previous section.

- 1) System-level information (D): the intended purpose of the system [19], [39], [40], [50]; the intended users of the system [31], [39], [40], [48]; the compliance specifications which apply to the system, i.e., the hard laws that must be followed and soft laws that should be followed [19], [48], [49]; and who is accountable for the creation of the specifications of the system purpose and the system compliance requirements.
- 2) Dataset information (D&I): characteristics of the dataset (e.g., size, composition of instances, number of features) [19], [31], [38], [39], [48], [50]; collection method [38], [40], [49]; any associated pre-processing (e.g., sampling, aggregation) [19], [38]–[40], [48]; tasks for which it should be used [38], [48] and those for which it should not be used [38]; its known or expected limitations [19], [40], [48], [49] its known or expected risks, including biases [19], [24], [26], [31], [38], [40], [48], [49]; and who is accountable for the creation of the design specification and the subsequent realization of the dataset [38], [51].
- 3) ML model information (D&I): characteristics of the ML model (e.g., decision threshold, excluded dataset features) [24], [39], [40], [51]; details related to its implementation (e.g., algorithm used) [19], [24], [39], [40]; tasks for which it should be used [19], [39], [49] and those for which it should not be used [39], [40]; its known or expected limitations [19], [39], [40], [48] its known or expected risks, including biases [19], [24], [31], [39], [40], [48], [49]; any associated evaluation procedures [19], [24], [39], [40]; and who is accountable for the creation of the design specification and the subsequent realization of the ML model [39] as well as the design of the ML model evaluation and subsequent performance of the evaluation [51].
- 4) Supporting infrastructure information (D&I): the system components which are not ‘core AI’ but may still be the source of erroneous behaviour of an AI system (e.g., user interface, API wrappers); the characteristics of the supporting infrastructure relevant to the accountability of AI systems such as human agency and oversight mechanisms (e.g., human-in-the-loop or human-in-command [19], [40]), audit mechanisms [19], [40], [49], security mechanisms [40], and the level of explanation provided by the system [19], [40], [48], [49]; its known or expected limitations; its known or expected risks, including biases [40]; and who is accountable for the creation of the design specification and the subsequent realization of the supporting infrastructure.
- 5) System evaluation information (D&I): information related to how the AI system as a whole is evaluated [40], including if any third-party evaluations were performed [40]; and who is accountable for the creation of the specification of the system evaluation and subsequent performance of the evaluation.
- 6) Certification information (I): certificates granted to the AI system [40] and by whom.
- 7) Guidelines information (I): deployment guidelines, including what the instructions are for installing and integrating the system, and what the instructions are for training users and updating manuals; operational guidelines including what the instructions are for system usage [40] and system maintenance [40]; and who

is accountable for the creation of the deployment and operational guidelines.

- 8) Human decision making and approvals (D&I): who assessed the fitness of the dataset design and implementation [19], [49], ML model design, implementation, and evaluation, supporting infrastructure design and implementation, and system evaluation against the system's purpose; which hard and soft law requirements were included [19], [48] in the dataset design and implementation, ML model design, implementation, and evaluation, supporting infrastructure design and implementation, system evaluation, and guidelines generation, and who assured the compliance of such specifications against those requirements; and finally, who approved the various specifications, system component implementations and their evaluations, and guidelines that influence subsequent life cycle stages.

III. SEMANTIC MODELLING OF ACCOUNTABILITY INFORMATION

This section provides a high-level overview of our proposed approach for documenting accountability information as part of a single Knowledge Graph (KG). For more technical details on how these ontologies were formalized and how complete KGs described using *SAO* and *RAInS* should be represented (i.e., using the appropriate relationship links between different class instances), we refer the reader to [45] and the ontology documentation. As outlined in Section I, we use *SAO* and *RAInS* ontologies to semantically model information about the AI system life cycle in the form of *accountability plans* and their corresponding *accountability traces*. The ontologies extend the W3C recommendation PROV-O [43] and its extension EP-Plan [44]. The resulting KG described using these ontologies is thus expressed as an acyclic causal provenance graph.

A. ACCOUNTABILITY PLANS

Accountability plans describe information about activities and their outcomes that are expected to occur during a system life cycle and should be recorded for accountability purposes. Fig. 1 illustrates a portion of an accountability plan for the implementation life cycle stage. The plan captures the expected flow of events describing part of a standard machine learning pipeline where a dataset is split into training and evaluation datasets, and then an ML model is trained on the training dataset and evaluated using the evaluation dataset. Activities such as dataset splitting and ML model training are described as accountable actions and the results that they produce are described as accountable results. This high-level vocabulary is defined by the *SAO* ontology which also provides concepts for linking the individual plan components to a particular life cycle stage, identifying which accountable agents may perform certain activities (as described in the plan), linking the information to a particular AI system, etc.

B. ACCOUNTABILITY TRACES

Each accountability plan is linked to one or more corresponding *accountability traces*. These contain the information about the events and outputs described in the plan as they actually happened in the real world. Fig. 2 illustrates part of an accountability trace capturing information about training an ML model (*Training Activity* with start and end times) and who is accountable for this activity (user *Alice*). Accountable results are described in the accountability trace as collections of *information elements*, each describing a different aspect of the result (such as limitations, license). In Fig. 2 *Model Component Information* generated by *Training Activity* has information element members *Performance Limitation* and *Gender Bias*. To draw an analogy with paper-based records, the information realization represents a report describing an outcome of an activity (e.g., an implemented ML model) and the information elements represent individual sections of the report. Note that in our approach we do not attempt to provide a semantic annotation for every concept that may be used to express ML model limitations. Instead, the KG would contain human readable text linked to the information element describing the limitation.

C. EXTENDING THE HIGH-LEVEL ACCOUNTABILITY VOCABULARY

To maximize the applicability of *SAO* to different accountable systems, the ontology only defines high-level concepts for modelling accountability plans and accountability traces. This limits the types of queries that can be answered over a KG and makes interpreting the results of those queries a challenge for human users. For example, referring to the example in Fig. 2, we could construct a query that lists all the information elements describing a particular accountable result (e.g., ML Model), but we would be unable to select only those elements describing limitations. In other words, we can retrieve the sections of the report, but do not know what is described in each section.

As our aim was to model accountability information relating to AI systems built around ML technologies, we extended *SAO* with the *RAInS* ontology to define subclasses (i.e., subtypes) of some of the core classes including accountability plan, accountable action, accountable result, and information element. Fig. 3 illustrates the *SAO* classes (shown in blue-filled rectangles) which *RAInS* extends. Third-party classes which were reused from ML Schema⁵ [52], Dublin Core⁶ (DC), and OntoSoft⁷ [53] have blue borders. For the *SAO* and *RAInS* properties, we refer the reader to the ontology documentation.

Design Stage Accountability Plan and *Implementation Stage Accountability Plan* provide 'containers' for plan definitions and both the design and implementation life cycle stages of an AI system are expected to have one of these

⁵<http://ml-schema.github.io/documentation/ML%20Schema.html>

⁶<https://www.dublincore.org/specifications/dublin-core/dcmi-terms/>

⁷https://ontosoft.org/ontology/software_20150329/




FIGURE 1. Part of an example *accountability plan* for implementing a machine learning system component described using the SAO ontology.

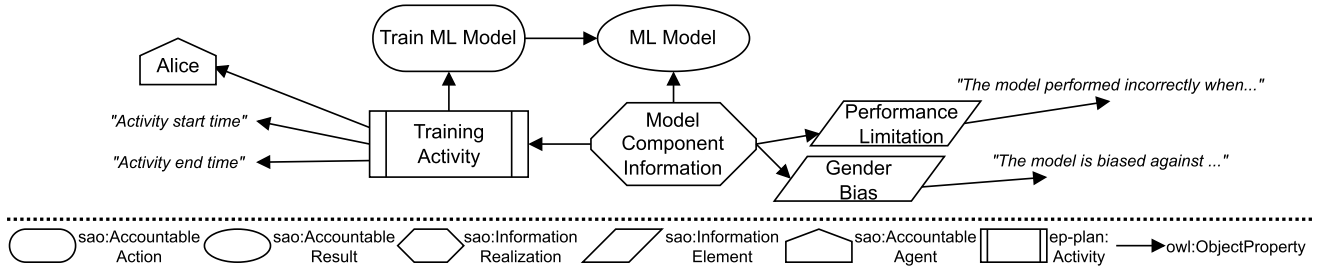


FIGURE 2. Part of an example *accountability trace* for implementing a machine learning system component described using the SAO ontology.

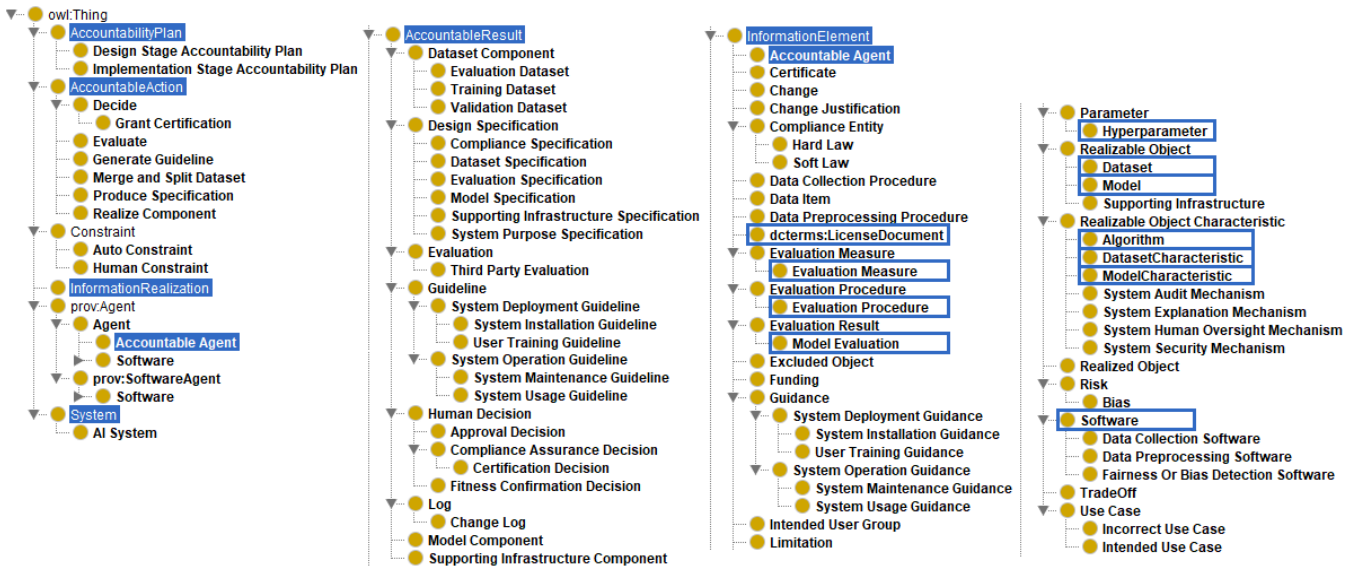


FIGURE 3. *RAINs* classes as subclasses of SAO classes (blue-filled rectangles). Third-party classes reused from ML Schema, Dublin Core, and OntoSoft are shown with blue borders.

plans defined. To provide more detailed descriptions of these accountability plans, subclasses of SAO's *AccountableAction* were defined to describe planned activities such as: (1) designing the components and the system's high-level specifications (see *Produce Specification*); (2) creating new tangible system components/assets (e.g., trained ML models, datasets) following their design specifications (*Realize Component* and *Merge and Split Dataset*); (3) evaluation of the individual components, as well as the AI system as a whole (*Evaluate*); (4) generating guidelines for deployment and operation (*Generate Guideline*); and (5) making decisions (*Decide*). Furthermore, subclasses of SAO's *AccountableResult* describe: (1) design specifications for the AI system and its components (*Design Specification* and its

subclasses); (2) implementations of individual system components (*Dataset Component* - and its subclasses, *Model Component*, and *Supporting Infrastructure Component*); (3) evaluations of the ML model and the AI system as a whole (*Evaluation*); (4) generated guidelines (*Guideline* and its subclasses); (5) human decisions (*Human Decision* and its subclasses); (6) logs of any deviations from design specifications (*ChangeLog*).

RAINs supports richer descriptions of accountable results in accountability traces through subclasses of SAO's *InformationElement*. These represent different types of accountability information such as limitations, references to hard and soft laws, user groups, use cases, risks, ML model and dataset characteristics, guidance, data collection procedure,

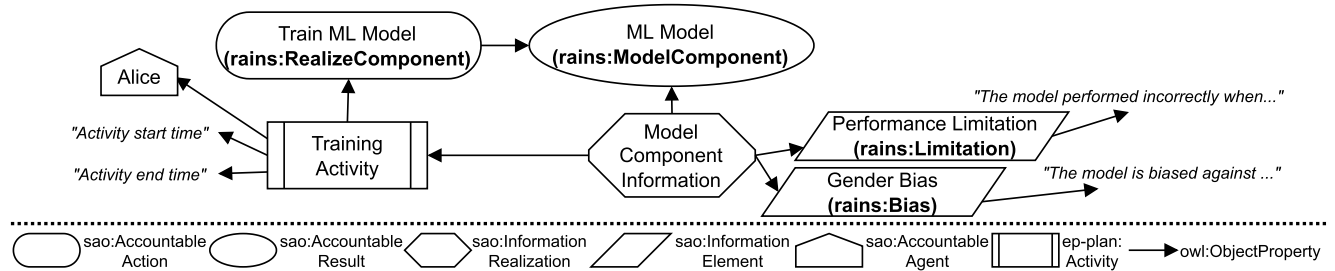


FIGURE 4. Part of the example *accountability trace* depicted in Fig. 2 enhanced with *RAINs* concepts (labelled with the *RAINs* prefix; see Fig. 3). Note, *KG* elements can be associated with multiple types—i.e., *Train ML Model* is of type *sao:AccountableAction* as well as *rains:RealizeComponent*.

etc. Fig. 4 shows a portion of the example accountability trace from Fig. 2 enhanced with *RAINs* concepts.

Finally, the *RAINs* ontology also defines two types of constraints: *AutoConstraint* and *HumanConstraint*. The former represents constraints associated with accountable actions that can be implemented (e.g., in the form of a rule) and automatically evaluated against the accountability trace (e.g., a constraint to check whether a particular information element was created or that the value of an information element is in a particular range). The latter represents constraints that would be difficult to automatically evaluate based on the accountability trace, but instead would need to be evaluated by human agents (e.g., a constraint to ensure that the agent performing an action has relevant expertise). Such constraints were introduced to help manage the quality of accountability traces and to support the discovery of inconsistencies within the accountability KG.

D. INCORPORATING USER FEEDBACK IN THE DESIGN PROCESS

To assess how well our approach addresses the knowledge capture and modelling requirements identified above, as well as to understand the potential benefits of our semantic based approach, two expert workshops were organised to demonstrate and discuss the *Accountability Fabric*. This prototype implementation comprised of a suite of tools demonstrating the feasibility of a computational solution to support planning, collection and inspection of accountability information using *SAO* and *RAINs* [45], [46]. Workshop 1 focused on legal experts (5 in total), whereas Workshop 2 focused on CS researchers and AI experts (10 in total). Both groups highlighted similar advantages of the *Accountability Fabric* and the resulting ontology concepts, including the fact that such a standardized, structured approach would be of use to regulators and investigators who may be less familiar with the mechanics of AI systems but are seeking specific information to determine regulatory compliance and to build a framework of enforcement. It was also felt that our approach could help to overcome knowledge gaps, particularly amongst the legal profession, by introducing a common understanding of accountability in AI that will allow meaningful information to be captured. CS researchers and AI experts also indicated that such a system would be beneficial by bringing structure

to the workflow for new entrants, as such an approach allows for meaningful communication with customers interested in deploying AI systems, again by providing a common, structured language. Participants felt that it could be helpful in troubleshooting where errors may have occurred in the event that an AI system behaves unexpectedly at the deployment stage, particularly given the attention to the life cycle and the potential for different actors to have been involved at different stages. The list of CQs was felt to be comprehensive, but participants suggested that there should be potential to extend the system if gaps were found. While participants agreed that there is some potential for the system to become burdensome, they also felt that the standardized, structured approach involving both automatically and human generated information is one that could be efficiently adopted and would remove some of the subjectivity present in existing accountability frameworks.

IV. COMPARISON WITH POPULAR NON-SEMANTIC FRAMEWORKS

In this section we assess the ability of our semantic ontology-based solution to describe accountability information, by comparing it with the popular non-semantic frameworks, Data Sheets, Model Cards, and FactSheets. For the remainder of this paper, we shall refer to these as baseline frameworks.

A. METHODOLOGY

The ability of our ontologies and the baseline frameworks to capture accountability information was evaluated against 96 individual CQs falling under the topics discussed in Section II. Information relating to *Human decisions* was excluded from this comparison, as such information was beyond the original scope of the baseline frameworks and will be discussed separately in Section IV-C.

To decide whether the baseline frameworks are able to capture accountability information required to support individual CQs, we first examined the published papers describing these frameworks ([38]–[40]), including their appendices. We then examined the papers’ supplementary materials, which consisted of two examples on the Model Card website,⁸ examples

⁸<https://modelcards.withgoogle.com/model-reports>

provided with the Model Card Toolkit,⁹ and seven examples listed on the FactSheets website.¹⁰ To decide whether our ontologies are able to capture the information necessary to answer the CQs, we determined whether relevant classes and relationships exist that provide the means to generate part of a corresponding KG (see the Appendix for how our ontologies' classes and relationships capture the required information).

When analyzing the baseline frameworks, we only considered aspects relating to the design and implementation stages of the AI system life cycle - i.e., the current scope of our ontology. For example, information provided to the FactSheet question: "When was the service first released? When was the last release?" relates to the system deployment stage, and the questions "When were the models last updated? How much did the performance change with each update? How often are the models retrained or updated?" relate to the operations stage.

It should be noted that the evaluation of the baseline frameworks' ability to answer individual CQs was based on the understanding of their functionalities by the two authors of this paper who conducted the comparison. This understanding was gained from the associated papers and existing examples as it was not the frameworks' goal to provide semantic realizations. This poses a potential risk in terms of misinterpretation of the materials by the authors and, therefore, we note this as a potential limitation of our approach.

B. RESULTS OF COMPARISONS BETWEEN OUR APPROACH AND THE BASELINE FRAMEWORKS

Tables 1 - 7 report the results of our evaluation. The notation used in each of the tables to indicate whether the frameworks are able to answer the CQs is as follows:

'Y' indicates that the framework does capture the information necessary to answer the CQ; 'S' indicates that the framework captures some of the information needed to answer the CQ; and 'N' indicates that the framework does not capture the necessary information. As previously mentioned, a detailed overview of the semantic concepts introduced in SAO and RAINs is included in the Appendix.

We now discuss the commonalities between the three baseline frameworks and our approach. As shown in Tables 1, 4, 5, 6, and 7 both our approach and FactSheets capture information about the whole system, supporting infrastructure, evaluation of the whole system, certifications, and operational guidelines. Datasheets and Model Cards do not, as they are fundamentally concerned with capturing information about datasets and ML models respectively. Also as expected, Table 2 shows that while Model Cards do capture some information about datasets, the Datasheets framework captures much more, with FactSheets somewhere between the two in terms of coverage. Finally, Table 3 shows that both FactSheets and Model Cards capture a considerable amount of information about ML models and their evaluation.

⁹<https://github.com/tensorflow/model-card-toolkit>

¹⁰<https://aifs360.mybluemix.net/examples>

C. ACCOUNTABILITY INFORMATION COVERAGE BEYOND THE SCOPE OF THE BASELINE FRAMEWORKS

We now summarize how our approach captures accountability information beyond the scope of the baseline frameworks. First, because SAO and RAINs extend PROV, they can model the dates and time frames when the accountable actions were performed and when the accountable results were produced. For example, when the ML model was designed, realized, and evaluated, and also when it was approved. Second, RAINs can model the compliance requirements, i.e., the hard and soft laws, that individual AI system components comply with. This is in addition to modelling the compliance requirements applying to the system as a whole, which none of the baseline frameworks model. Third, RAINs models the different human decisions and approvals (corresponding to item 8 in Section II-B) that are integral to the life cycle of an AI system. This includes: who confirmed the fitness of each of the components to the system's specified purpose; who checked the compliance of each of the components with the system's compliance requirements; and who approved design specifications, system component implementations and their evaluations, and generated guidelines. Fourth, RAINs models the agents within the organizations who were accountable for the different actions. This is important for internal and external auditing. Finally, RAINs models whether the realizations of system components or their evaluation followed the design specification, as well as the justification for any deviations from the design specification. Typically, the accountable agent performing the realization has access to the design specification and is expected to document how it was performed. However, if for some reason, the agent decides to deviate from the design specification, then they must document their actions along with the change log and their justification for the change. This is crucial in associating accountability with the correct human agent. For example, if an agent, *Bob*, followed a faulty design produced by another agent, *Alice*, then the accountability falls on *Alice* or possibly on both *Alice* and *Bob*. If, however, *Bob* deviated from a design specification, which contained no errors, resulting in an AI system that produced unacceptable errors, then they are solely accountable.

V. DISCUSSION

A. BENEFITS OF SEMANTIC REPRESENTATION OF ACCOUNTABILITY INFORMATION

KGs model accountability information in a structured interoperable format that is understandable by both humans and machines. While in our approach a portion of the accountability information is stored in human-readable formats (e.g., text describing information elements, or images depicting evaluation results), many aspects of the accountability KGs are described using semantic concepts following the three core PROV concepts, namely: agents, entities and activities. This approach to modelling data makes it easier to integrate additional accountability information into the KG. For example,

TABLE 1. Information about the AI system (corresponds to item 1 in Section II-B).

Information about the AI system (corresponds to item 1 in Section II-B)

Information relating to the AI system	RAInS	Fact Sheets	Data Sheets	Model Cards
Who is the creator/supplier of this AI system?	Y	Y	N	N
What are the intended use cases or domains of this AI system?	Y	Y		
Who are the intended users of this AI system?	Y	Y	N	N
Who defined the intended use cases and users of this AI system?	Y	S ^a		
What are the compliance requirements of this AI system?	Y	N		
Who defined the compliance requirements of this AI system?	Y	N		

^a The FactSheets paper does not ask who performed which action. It asks about the characteristics of developers and whether domain experts were involved. Therefore, we mark CQs related to who performed the development as ‘S’ and those not related to the development as ‘N’.

TABLE 2. Information about the dataset (corresponds to item 2 in Section II-B).

Information relating to the design and implementation of the dataset	RAInS	Fact Sheets	Data Sheets	Model Cards
Information about the dataset				
Who produced the dataset design specification and on behalf of whom or which organization?	Y	S ^a	N	
Who realized the dataset and on behalf of whom or which organization?	Y	S ^a	Y	N
Was the dataset realization different from its design specification, and if yes then why?	Y	N	N	
What is the dataset used to build the ML model?	Y	Y	Y	Y
What is the version or date of the dataset?	Y	Y	Y	N
Was the dataset created specifically for this AI system or was it reused?	Y	Y	Y	S ^b
If the dataset was reused, who is its creator or manager?	Y	N	Y	
What are the intended use cases of the dataset?	Y	Y	Y	
What are incorrect use cases for the dataset?	Y		Y	
What is the funding associated with the dataset?	Y	N	Y	
What are the licenses associated with the dataset?	Y		S ^c	
What are the limitations of the dataset?	Y	Y	Y	
What are the risks, including biases, of the dataset?	Y	Y	Y	
What risk and bias software was used to analyze the dataset?	Y	Y ^d	Y	
What are the mitigation strategies to address the risks and biases of the dataset?	Y	Y	Y	
What are the risk, including bias, estimates of the dataset before and after mitigation?	Y	S ^c	N	N
What are the trade-offs between risks, including biases, and the effects of using the dataset on humans or others?	Y	S ^f	N	
Who collected the data contained in the dataset?	Y	S ^a	Y	
What were the data collection methods and timeframes associated with the dataset?	Y	Y	Y	
What data collection software was used?	Y	N	Y	
What considerations were applied to the collection of data relating to people?	Y	Y	Y	
Who pre-processed the data in the dataset?	Y	S ^a	Y	
What were the methods for pre-processing the dataset?	Y	Y	Y	Y
What software was used to pre-process the dataset?	Y	N	Y	N
What are the characteristics of the dataset, including description of instances and errors/noise?	Y	Y	Y	Y
What is the link to external sources about the dataset or to the raw dataset?	Y	Y	Y	Y
Information relating to the merging and splitting of the Dataset				
Who performed the merging and/or splitting of the dataset?	Y	S ^a	N	N
How was the dataset split, including the train-validate-evaluate ratio and the class representation percentage in each subset?	Y	Y	Y	Y
What is the training subset of the dataset?	Y	Y	Y	Y
What is the validation subset of the dataset?	Y	Y	Y	N
What is the evaluation subset of the dataset?	Y	Y	Y	Y

^a The FactSheets paper does not ask who performed which action. It asks about the characteristics of developers and whether domain experts were involved. Therefore, we mark CQs related to who performed the development as ‘S’ and those not related to the development as ‘N’.

^b The Model cards paper only recommends that evaluation datasets should include those that are publicly available from third parties.

^c The Datasheets paper asks about the licenses associated with a dataset that was implemented in-house under the theme of ‘Distribution’. It does not ask about licenses if the dataset was reused from a third-party.

^d The FactSheets paper does not explicitly ask about software but because it asks about methods and results, we concluded that this is applicable. One of the FactSheets examples indicates that a fairness tool was used.

^e The FactSheets paper asks about risks and biases regarding the whole service, we split our questions based on individual AI System components. We determined that it is applicable here even though it was not explicitly asked.

^f The FactSheets paper asks about trade-offs regarding the whole service, however, we split our questions based on individual AI System components. We determined that it is applicable here even though it was not explicitly asked.

the descriptions of agents may be expanded using the FOAF¹¹ vocabulary to list their contact information, which can then be retrieved when querying the accountability KG. Furthermore, accountability plans provide the means to guide and assess the collection of accountability trace information that goes beyond simple template documents, where users are only

guided by different questions split over a number of sections as is the case with Datasheets and FactSheets. We do however acknowledge that the SMACTR framework,¹² which builds on Datasheets and Model Cards, does provide document checklists that guide users but split over different sections. As we demonstrate in [45], [46], it is possible to use

¹¹<http://xmlns.com/foaf/spec/>

¹²<https://edwinwenink.github.io/ai-ethics-tool-landscape/tools/smactr/>

TABLE 3. Information about the ML model (corresponds to item 3 in Section II-B).

Information relating to the design, implementation, and evaluation of the ML model	RAInS	Fact Sheets	Data Sheets	Model Cards	
Information about the ML model					
Who produced the ML model design specification and on behalf of whom or which organization?	Y	S ^d	N	N	
Who realized the ML model and on behalf of whom or which organization?	Y	S ^d		Y	
Was the ML model realization different from its design specification, and if yes then why?	Y	N		N	
What is the ML model?	Y	Y		Y	
What is the version or date of the ML model?	Y	Y		Y	
Was the ML model created specifically for this AI system or was it reused?	Y	N		N	
If the ML model was reused, who is its creator or manager?	Y	N		N	
What are the intended use cases of the ML model?	Y	Y		Y	
What are incorrect use cases for the ML model?	Y	Y		Y	
What is the funding associated with the ML model?	Y	N		N	
What are the licenses associated with the ML model?	Y	N		Y	
What are the limitations of the ML model?	Y	Y		N	Y
What are the risks, including biases, of the ML model?	Y	Y		N	Y
What risk and bias software was used to analyze the ML model?	Y	Y ^d			Y
What are the mitigation strategies to address the risks and biases of the ML model?	Y	Y		Y	
What are the risk, including bias, estimates of the ML model before and after mitigation?	Y	S ^e		N	
What are the trade-offs between risks, including biases, and the effects of using the ML model on humans or others?	Y	S ^f		N	
What are the characteristics of the ML model relating to emphasized or excluded features, decision thresholds and basis of prioritization, as well as its inputs and outputs and programming methods?	Y	Y		Y	
What algorithm is used by the ML model?	Y	Y		Y	
What are the ML model's parameters and hyperparameters?	Y	Y		Y	
What are the sources of additional information and citation details describing the ML model?	Y	Y	Y		
Information relating to designing the evaluation of the ML model and performing the evaluation					
Who produced the ML model evaluation specification and on behalf of whom or which organization?	Y	S ^d	N	N	
Who evaluated the ML model and on behalf of whom or which organization?	Y	S ^d		Y	
Was the ML model evaluation different from its design specification, and if yes then why?	Y	N		N	
What was the evaluation procedure, including evaluation factors and scenarios?	Y	Y		Y	
How were the evaluation results measured?	Y	Y		Y	
What are the evaluation results?	Y	Y		Y	
What are the changes in performance metric values before and after mitigation of risks and biases?	Y	Y ^e		N	

^a The FactSheets paper does not ask who performed which action. It asks about the characteristics of developers and whether domain experts were involved. Therefore, we mark CQs related to who performed the development as 'S' and those not related to the development as 'N'.

^d The FactSheets paper does not explicitly ask about software but because it asks about methods and results, we concluded that this is applicable. One of the FactSheets examples indicates that a fairness tool was used.

^e The FactSheets paper asks about risks and biases regarding the whole service, we split our questions based on individual AI System components. We determined that it is applicable here even though it was not explicitly asked.

^f The FactSheets paper asks about trade-offs regarding the whole service, however, we split our questions based on individual AI System components. We determined that it is applicable here even though it was not explicitly asked.

TABLE 4. Information about the supporting infrastructure (corresponds to item 4 in Section II-B).

Information relating to the design and implementation of the supporting infrastructure	RAInS	Fact Sheets	Data Sheets	Model Cards
Who produced the supporting infrastructure design specification and on behalf of whom or which organization?	Y	S ^a	N	N
Who realized the supporting infrastructure and on behalf of whom or which organization?	Y	S ^a		
Was the supporting infrastructure realization different from its design specification, and if yes then why?	Y	N		
What is the supporting infrastructure, i.e., the non core-AI components of this system?	Y	S ^b		
Was the supporting infrastructure created specifically for this AI system or was it reused?	Y	N		
If the supporting infrastructure was reused, who is its creator or manager?	Y			
What is the version or date of the supporting infrastructure?	Y			
What are the limitations of the supporting infrastructure?	Y			
What are the risks, including biases, of the supporting infrastructure (e.g., sensors tuned to recognize different skin tones)?	Y	S ^c		
What risk and bias software was used to analyze the supporting infrastructure?	Y	S ^d		
What are the mitigation strategies to address the risks and biases of the supporting infrastructure?	Y	Y		
What are the risk, including bias, estimates of the supporting infrastructure before and after mitigation?	Y	S ^e		
What are the human oversight and human-in-control mechanisms of this AI system?	Y	Y		
What are the security mechanisms of this AI system?	Y	Y		
What are the audit mechanisms of this AI system?	Y	Y		
What are the explanation mechanisms provided by this AI system?	Y	Y		

^a The FactSheets paper does not ask who performed which action. It asks about the characteristics of developers and whether domain experts were involved. Therefore, we mark CQs related to who performed the development as 'S' and those not related to the development as 'N'.

^d The FactSheets paper does not explicitly ask about software but because it asks about methods and results, we concluded that this is applicable. One of the FactSheets examples indicates that a fairness tool was used.

^e The FactSheets paper asks about risks and biases regarding the whole service, we split our questions based on individual AI System components. We determined that it is applicable here even though it was not explicitly asked.

^g The FactSheets paper does not explicitly mention non core-AI components, but it asks about the service's inputs and outputs and an appendix example mentions integration of the service with other applications.

^h Not explicitly mentioned in the Model cards paper but mentioned as additional details that may be needed.

TABLE 5. Information about the evaluation of the whole system (corresponds to item 5 in Section II-B).

Information relating to designing the system evaluation and performing the evaluation	RAInS	Fact Sheets	Data Sheets	Model Cards
Who produced the system evaluation specification and on behalf of whom or which organization?	Y	S ^a		
Who evaluated the system and on behalf of whom or which organization?	Y	S ^a		
Was the system evaluation different from its design specification, and if yes then why?	Y	N		
Was the system evaluated by a third-party, and if yes by whom?	Y	Y		
What was the evaluation procedure including evaluation factors, domains and applications, key performance indicators tests and non-functional tests?	Y	Y	N	N
How were the evaluation results measured?	Y	Y		
What were the evaluation results with respect to each evaluation factor?	Y	Y		

^a The FactSheets paper does not ask who performed which action. It asks about the characteristics of developers and whether domain experts were involved. Therefore, we mark CQs related to who performed the development as ‘S’ and those not related to the development as ‘N’.

TABLE 6. Information about the certifications (corresponds to item 6 in Section II-B).

Information relating to certifications	RAInS	Fact Sheets	Data Sheets	Model Cards
Information about certifications				
What was the result of the certification process?	Y	S		
Who granted the certificate?	Y	N	N	N

TABLE 7. Information about deployment and operational guidelines (corresponds to item 7 in Section II-B).

Information relating to deployment and operational guidelines	RAInS	Fact Sheets	Data Sheets	Model Cards
Information about the deployment guidelines				
Who produced the deployment guidelines for installing the AI system and integrating it (e.g., within a larger IT ecosystem) and on behalf of whom or which organization?	Y		N	
What is the guidance for installing this AI system and integrating it (e.g., within a larger IT ecosystem)?	Y	N	N ⁱ	N
Who produced the deployment guidelines for training users to use and maintain this AI system?	Y		N	
What is the guidance for training users to use and maintain this AI system?	Y		N	
Information about the operational guidelines				
Who produced the operational guidelines for how to use this AI system and on behalf of whom or which organization?	Y	S ^a	N	
What is the guidance for using this AI system?	Y	Y	N	N
Who produced the operational guidelines for how to maintain this AI system and on behalf of whom or which organization?	Y	S ^a	N	
What is the guidance for maintenance of this AI system?	Y	Y	N ^j	

^a The FactSheets paper does not ask who performed which action. It asks about the characteristics of developers and whether domain experts were involved. Therefore, we mark CQs related to who performed the development as ‘S’ and those not related to the development as ‘N’.

ⁱ The Datasheets paper includes a section listing questions about how the dataset is to be distributed.

^j The Datasheets paper includes a section listing questions about how the dataset is to be maintained.

accountability plans to build intelligent user interfaces that manage the gathering of accountability information through manual and also semi-automatic processes (e.g., by integrating with programming applications used to implement components of the AI system). Plan constraints can be used to further enrich this process by providing real-time warnings to users about missing accountability information (in the form of information elements) so as to avoid liability by omission. Constraints also allow checking for potential quality issues in the supplied information (e.g., such as when an agent claims to have evaluated an ML model before this model was implemented, or supplies values that fall outside pre-defined ranges). We envision that this planning mechanism could be used by regulators and standardization agencies in future to provide generic plan templates to accompany their AI regulations or standards. Such generic plans would then be employed by the respective organizations involved in AI system development to guide the gathering of accountability information. Accountability traces would be provided by a collection of compatible tools capable of understanding relevant portions of these plans.

TABLE 8. List of ontologies with their namespaces and prefixes.

Ontology	Namespace	Prefix
PROV	http://www.w3.org/ns/prov#	prov
EP-Plan	https://w3id.org/ep-plan#	ep-plan
SAO	https://w3id.org/sao#	sao
RAInS	https://w3id.org/rains#	rains
ML Schema	http://www.w3.org/ns/mls#	mls
DC Terms	http://purl.org/dc/terms/	dcterms
OntoSoft	http://ontosoft.org/software#	software
RDF	http://www.w3.org/1999/02/22-rdf-syntax-ns#	rdf
RDF Schema	http://www.w3.org/2000/01/rdf-schema#	rdfs
XSD	http://www.w3.org/2001/XMLSchema#	xsd

B. INTEGRATION WITH EXISTING FRAMEWORKS

The flexibility of our semantic-based approach does not necessarily mean that users of existing frameworks, such as those discussed in our comparison evaluation, must migrate to an entirely new solution. Indeed, even some of the baseline frameworks such as Model Cards and Datasheets have been designed with an intention to complement each other [30], [39]. While collecting accountability information in a text-based format (such as a Datasheet or FactSheet) would not be supported by our approach, interactive user interfaces

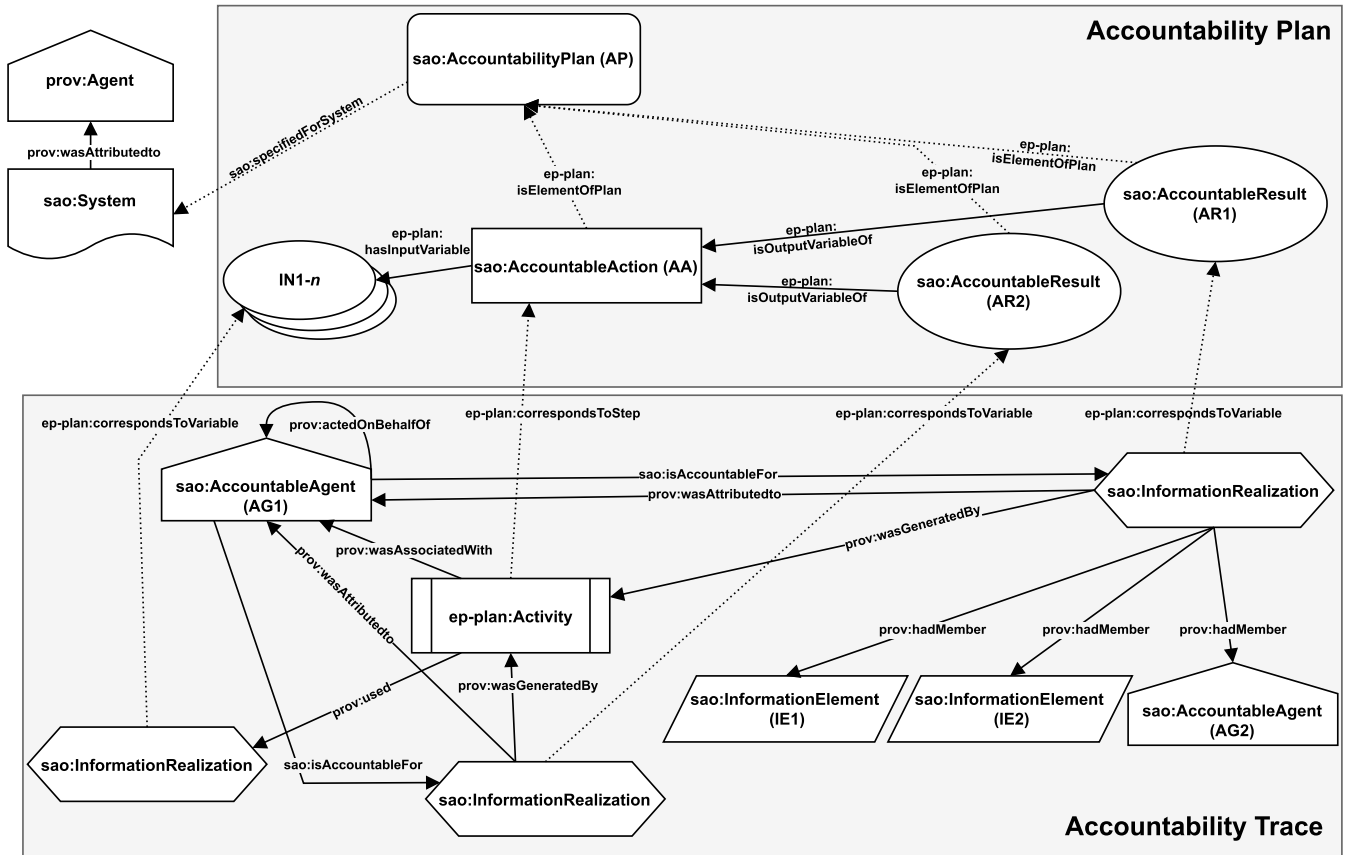


FIGURE 5. How entities, activities, and agents relate to each other.

TABLE 9. Information about the AI system.

Information relating to the AI system	Ontological Modelling
Who is the creator/supplier of this AI system?	<i>prov:Agent</i> (optional instance of <i>prov:Organization</i>) linked using the property <i>prov:wasAttributedTo</i> (with <i>prov:Role</i> being <i>Supplier</i>) to <i>rains:AI_System</i> (treated as <i>prov:Entity</i>)
What are the intended use cases or domains of this AI system?	<i>[rains:DesignStageAccountabilityPlan (AP): rains:ProduceSpecification (AA); rains:SystemPurposeSpecification (AR1); sao:InformationRealization</i> corresponding to <i>rains:SystemPurposeSpecification (AR1); rains:IntendedUseCase^a(IE1)]</i>
Who are the intended users of this AI system?	<i>[rains:DesignStageAccountabilityPlan (AP): rains:ProduceSpecification (AA); rains:SystemPurposeSpecification (AR1); sao:InformationRealization</i> corresponding to <i>rains:SystemPurposeSpecification (AR1); rains:IntendedUserGroup (IE1)]</i>
Who defined the intended use cases and users of this AI system?	<i>[rains:DesignStageAccountabilityPlan (AP): sao:AccountableAgent^b(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to <i>rains:SystemPurposeSpecification (AR1)]</i></i>
What are the compliance requirements (the hard and soft laws) of this AI system?	<i>[rains:DesignStageAccountabilityPlan (AP): rains:ProduceSpecification (AA); rains:ComplianceSpecification (AR1); sao:InformationRealization</i> corresponding to <i>rains:ComplianceSpecification (AR1); rains:HardLaw (IE1); rains:SoftLaw (IE2)]</i>
Who defined the compliance requirements of this AI system?	<i>[rains:DesignStageAccountabilityPlan (AP): sao:AccountableAgent^b(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to <i>rains:ComplianceSpecification (AR1)]</i></i>

^aTasks for which the AI system must not be used can be modelled using *rains:IncorrectUseCase (IE2)*.

^bOur approach requires the URI of the person who performed the action to be recorded, but – unlike FactSheets – does not ask about the characteristics of the development team or whether there were domain experts involved. To enforce capture of whether the agent(s) were domain experts, *rains:HumanConstraint* may be used. Descriptions of agents may be expanded using the FOAF vocabulary to list their contact information and company affiliations and any additional characteristics can be recorded using *rdfs:comment^(dp)* linking *sao:AccountableAgent* to a String literal.

could be created that resemble the design intentions of these existing frameworks. We demonstrate this in [46] where we first integrate the data generated using the Model Card Toolkit into our accountability KG and later provide an option to display relevant parts of the KG in the form of a Model Card. Here, the semantic layer supported by our ontologies serves as middleware enabling the collection of information from heterogeneous sources (e.g., via manual input or through

semi-automated frameworks such as the Model Card Toolkit) and provides a means for the information to be queried and displayed via a range of user interfaces.

C. THE NEED FOR SMARTER INFORMATION COLLECTION AND AUDIT MECHANISMS

There is no doubt that gathering good quality accountability information requires non-trivial effort due to the number of

TABLE 10. Information about the dataset.

Information relating to the design and implementation of the dataset	Ontological Modelling
Information about the dataset	
Who produced the dataset design specification and on behalf of whom or which organization?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf</i> ^(op) .]
Who realized the dataset and on behalf of whom or which organization?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the two <i>sao:InformationRealizations</i> corresponding to the two output variables of <i>rains:RealizeComponent</i> , namely <i>rains:DatasetComponent</i> (AR1) and <i>rains:ChangeLog</i> (AR2). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf</i> ^(op) .]
Was the dataset realization different from its design specification, and if yes then why?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:ChangeLog</i> (AR2); <i>sao:InformationRealization</i> corresponding to <i>rains:ChangeLog</i> (AR2); <i>rains:ChangeJustification</i> (IE1); the instance of the designed but not realized <i>mls:Dataset</i> in the information realization corresponding to <i>rains:DatasetComponent</i> (AR1) becomes an instance of <i>rains:ExcludedObject</i>]
What is the dataset used to build the ML model?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>mls:Dataset</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>mls:Dataset</i> (IE1); the instance of the realized <i>mls:Dataset</i> is modelled as an instance of <i>rains:RealizedObject</i> also].
What is the version or date of the dataset?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:version</i> ^(dp) linking <i>mls:Dataset</i> (IE1) to a String literal; <i>rains:versionDate</i> ^(dp) linking <i>mls:Dataset</i> (IE1) to a date-time literal; <i>rains:versionNote</i> ^(dp) linking <i>mls:Dataset</i> (IE1) to a string literal] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:version</i> ^(dp) linking <i>mls:Dataset</i> (IE1) to a String literal; <i>rains:versionDate</i> ^(dp) linking <i>mls:Dataset</i> (IE1) to a date-time literal; <i>rains:versionNote</i> ^(dp) linking <i>mls:Dataset</i> (IE1) to a string literal]
Was the dataset created specifically for this AI system or was it reused?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:isReusedObject</i> ^(dp) linking <i>mls:Dataset</i> (IE1) to a Boolean literal of a true value if the dataset is reused.] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:isReusedObject</i> ^(dp) linking <i>mls:Dataset</i> (IE1) to a Boolean literal of a true value if the dataset is reused.]
If the dataset was reused, who is its creator or manager?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>sao:AccountableAgent</i> ^b (AG2) linked using <i>sao:isAccountableFor</i> ^(op) to <i>mls:Dataset</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>sao:AccountableAgent</i> ^b (AG2) linked using <i>sao:isAccountableFor</i> ^(op) to <i>mls:Dataset</i> (IE1)]
What are the intended use cases of the dataset?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:IntendedUseCase</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:IntendedUseCase</i> (IE1)]
What are incorrect use cases for the dataset?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:IncorrectUseCase</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:IncorrectUseCase</i> (IE1)]
What is the funding associated with the dataset?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:Funding</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:Funding</i> (IE1)]

TABLE 10. (Continued.) Information about the dataset.

<p>What are licenses associated with the dataset?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>dterms:LicenseDocument</i> (IE1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>dterms:LicenseDocument</i> (IE1)]</p>
<p>What are the limitations of the dataset?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:Limitation</i> (IE1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:Limitation</i> (IE1)]</p>
<p>What are the risks, including biases, of the dataset?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:Risk</i> (IE1); <i>rains:Bias</i> (IE2)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:Risk</i> (IE1); <i>rains:Bias</i> (IE2)]</p>
<p>What risk and bias software was used to analyze the dataset?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:FairnessOrBiasDetectionSoftware</i> (IE1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:FairnessOrBiasDetectionSoftware</i> (IE1)]</p>
<p>What are the mitigation strategies to address the risks and biases of the dataset?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:hasMitigationStrategy</i>^(op) linking <i>rains:Risk</i> (IE1) or <i>rains:Bias</i> (IE2) to a string literal]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:hasMitigationStrategy</i>^(op) linking <i>rains:Risk</i> (IE1) or <i>rains:Bias</i> (IE2) to a string literal]</p>
<p>What are the risk, including bias, estimates of the dataset before and after mitigation?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:Risk</i> (IE1); <i>rains:Bias</i> (IE2)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:Risk</i> (IE1); <i>rains:Bias</i> (IE2)]</p> <p>-</p> <p>While this CQ is modelled using the same concepts as the CQ ‘What are the risks, including biases, of the dataset?’, new instances of the two information elements <i>rains:Risk</i> and <i>rains:Bias</i> must be created to record the new estimates.</p>
<p>What are the trade-offs between risks, including biases, and the effects of using the dataset on humans or others?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:TradeOff</i> (IE1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:TradeOff</i> (IE1)]</p>
<p>Who collected the data contained in the dataset?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>sao:AccountableAgent</i> (AG2) linked using <i>sao:isAccountableFor</i>^(op) to <i>rains:DatasetCollectionProcedure</i> (IE1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>sao:AccountableAgent</i> (AG2) linked using <i>sao:isAccountableFor</i>^(op) to <i>rains:DatasetCollectionProcedure</i> (IE1)]</p>
<p>What were the data collection methods and time-frames associated with the dataset?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:DataCollectionProcedure</i> (IE1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:DataCollectionProcedure</i> (IE1)]</p>
<p>What data collection software was used?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains>DataCollectionSoftware</i> (IE1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains>DataCollectionSoftware</i> (IE1)]</p>

TABLE 10. (Continued.) Information about the dataset.

What considerations were applied to the collection of data relating to people?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:Risk</i> (IE1); <i>rains:Bias</i> (IE2)]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:Risk</i> (IE1); <i>rains:Bias</i> (IE2)]</p> <p>-</p> <p>While this CQ is modelled using the same concepts as the CQ ‘What are the risks, including biases, of the dataset?’, new instances of the two information elements <i>rains:Risk</i> and <i>rains:Bias</i> may optionally be created to model considerations specific to the collection of the data.</p>
Who pre-processed the data in the dataset?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>sao:AccountableAgent</i> (AG2) linked using <i>sao:isAccountableFor</i>^(op) to <i>rains:DataPreprocessingProcedure</i> (IE1)]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>sao:AccountableAgent</i> (AG2) linked using <i>sao:isAccountableFor</i>^(op) to <i>rains:DataPreprocessingProcedure</i> (IE1)]</p>
What were the methods for pre-processing the dataset?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:DataPreprocessingProcedure</i> (IE1)]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:DataPreprocessingProcedure</i> (IE1)]</p>
What software was used to pre-process the dataset?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:DataPreprocessingSoftware</i> (IE1)]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:DataPreprocessingSoftware</i> (IE1)]</p>
What are the characteristics of the dataset, including description of instances and errors/noise?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:DatasetCharacteristic</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i>^(op) to <i>mls:Dataset</i> (IE1); <i>rains:hasBase64Image</i>^(dp) with a string literal domain]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>mls:DatasetCharacteristic</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i>^(op) to <i>mls:Dataset</i> (IE1); <i>rains:hasBase64Image</i>^(dp) with a string literal domain]</p>
What is the link to external sources about the dataset or to the raw dataset?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i>; (AR1) <i>rdfs:seeAlso</i>^(op) linking <i>mls:Dataset</i> (IE1) to the relevant entity]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rdfs:seeAlso</i>^(op) linking <i>mls:Dataset</i> (IE1) to the relevant entity]</p>
Information relating to the Merging and Splitting of the Dataset	
Who performed the merging and/or splitting of the dataset?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:MergeAndSplitDataset</i> (AA) with <i>rains:DatasetComponent</i> (IN1) as an input; <i>sao:AccountableAgent</i> ^(b) (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the three <i>sao:InformationRealizations</i> corresponding to the three output variables of <i>rains:MergeAndSplitDataset</i> , namely <i>rains:TrainingDataset</i> (AR1), <i>rains:ValidationDataset</i> (AR2), and <i>rains:EvaluationDataset</i> (AR3)]
How was the dataset split, including the train-validate-evaluate ratio and the class representation percentage in each subset? ^c	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:MergeAndSplitDataset</i> (AA) with <i>rains:DatasetComponent</i> (IN1) as an input; <i>sao:InformationRealizations</i> corresponding to <i>rains:TrainingDataset</i> (AR1), <i>rains:ValidationDataset</i> (AR2), and <i>rains:EvaluationDataset</i> (AR3)]
What is the training subset of the dataset?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:MergeAndSplitDataset</i> (AA) with <i>rains:DatasetComponent</i> (IN1) as an input; <i>rains:TrainingDataset</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:TrainingDataset</i> (AR1); <i>mls:Dataset</i> (IE1); <i>mls:DatasetCharacteristics</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i> ^(op) to <i>mls:Dataset</i> (IE1); Plus optional additional <i>sao:InformationElements</i> describing this subset]
What is the validation subset of the dataset?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:MergeAndSplitDataset</i> (AA) with <i>rains:DatasetComponent</i> (IN1) as an input; <i>ValidationDataset</i> (AR2); <i>sao:InformationRealization</i> corresponding to <i>rains:ValidationDataset</i> (AR2); <i>mls:Dataset</i> (IE1); <i>mls:DatasetCharacteristics</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i> ^(op) to <i>mls:Dataset</i> (IE1); Plus optional additional <i>sao:InformationElements</i> describing this subset]
What is the evaluation dataset of the dataset?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:MergeAndSplitDataset</i> (AA) with <i>rains:DatasetComponent</i> (IN1) as an input; <i>EvaluationDataset</i> (AR3); <i>sao:InformationRealization</i> corresponding to <i>rains:EvaluationDataset</i> (AR3); <i>mls:Dataset</i> (IE1); <i>mls:DatasetCharacteristics</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i> ^(op) to <i>mls:Dataset</i> (IE1); Plus optional additional <i>sao:InformationElements</i> describing this subset]

^bOur approach requires the URI of the person who performed the action to be recorded, but – unlike FactSheets – does not ask about the characteristics of the development team or whether there were domain experts involved. To enforce capture of whether the agent(s) were domain experts, *rains:HumanConstraint* may be used. Descriptions of agents may be expanded using the FOAF vocabulary to list their contact information and company affiliations and any additional characteristics can be recorded using *rdfs:comment*^(dp) linking *sao:AccountableAgent* to a String literal.

^cThe dataset’s design specification may include indications on how the realized dataset needs to be split.

TABLE 11. Information about the ML model.

Information relating to the design and implementation of the ML model	Ontological Modelling
Who produced the ML model design specification and on behalf of whom or which organization?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>sao:AccountableAgent^b</i> (AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf^(op)</i> .]
Who realized the ML model and on behalf of whom or which organization?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>sao:AccountableAgent^b</i> (AG1) linked using <i>sao:isAccountableFor^(op)</i> to the two <i>sao:InformationRealizations</i> corresponding to the two output variables of <i>rains:RealizeComponent</i> , namely <i>rains:ModelComponent</i> (AR1) and <i>rains:ChangeLog</i> (AR2). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf^(op)</i> .]
Was the ML model realization different from its design specification, and if yes then why?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ChangeLog</i> (AR2); <i>sao:InformationRealization</i> corresponding to <i>rains:ChangeLog</i> (AR2); <i>rains:ChangeJustification</i> (IE1); the instance of the designed but not realized <i>mls:Model</i> in the information realization corresponding to <i>rains:ModelComponent</i> (AR1) becomes an instance of <i>rains:ExcludedObject</i>]
What is the ML model?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>mls:Model</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>mls:Model</i> (IE1); the instance of the realized <i>mls:Model</i> is modelled as an instance of <i>rains:RealizedObject</i> also.]
What is the version or date of the ML model?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rains:version^(dp)</i> linking <i>mls:Model</i> (IE1) to a String integer literal; <i>rains:versionDate^(dp)</i> linking <i>mls:Model</i> (IE1) to a date-time literal; <i>rains:versionNote^(dp)</i> linking <i>mls:Model</i> (IE1) to a string literal] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rains:version^(dp)</i> linking <i>mls:Model</i> (IE1) to a String literal; <i>rains:versionDate^(dp)</i> linking <i>mls:Model</i> (IE1) to a date-time literal; <i>rains:versionNote^(dp)</i> linking <i>mls:Model</i> (IE1) to a string literal]
Was the ML model created specifically for this AI system or was it reused?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rains:isReusedObject^(dp)</i> linking <i>mls:Model</i> (IE1) to a Boolean literal of a <i>true</i> value if the ML model is reused.] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rains:isReusedObject^(dp)</i> linking <i>mls:Model</i> (IE1) to a Boolean literal of a <i>true</i> value if the ML model is reused.]
If the ML model was reused, who is its creator or manager?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>sao:AccountableAgent^b</i> (AG2) linked using <i>sao:isAccountableFor^(op)</i> to <i>mls:Model</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>sao:AccountableAgent^b</i> (AG2) linked using <i>sao:isAccountableFor^(op)</i> to <i>mls:Model</i> (IE1)]
What are the intended use cases of the ML model?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR2); <i>rains:IntendedUseCase</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rains:IntendedUseCase</i> (IE1)]
What are incorrect use cases for the ML model?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rains:IncorrectUseCase</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> ; <i>rains:IncorrectUseCase</i> (IE1)]
What is the funding associated with the ML model?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rains:Funding</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rains:Funding</i> (IE1)]

TABLE 11. (Continued.) Information about the ML model.

<p>What are the licenses associated with the ML model?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>dterms:LicenseDocument</i> (IE1)]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>dterms:LicenseDocument</i> (IE1)]</p>
<p>What are the limitations of the ML model?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rains:Limitation</i> (IE1)]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rains:Limitation</i> (IE1)]</p>
<p>What are the risks, including biases, of the ML model?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rains:Risk</i> (IE1); <i>rains:Bias</i> (IE2)]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rains:Risk</i> (IE1); <i>rains:Bias</i> (IE2)]</p>
<p>What risk and bias software was used to analyze the ML model?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rains:FairnessOrBiasDetectionSoftware</i> (IE1)]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rains:FairnessOrBiasDetectionSoftware</i> (IE1)]</p>
<p>What are the mitigation strategies to address the risks and biases of the ML model?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA); with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2) <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rains:hasMitigationStrategy</i>^(op) linking <i>rains:Risk</i> (IE1) or <i>rains:Bias</i> (IE2) to a string literal]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rains:hasMitigationStrategy</i>^(op) linking <i>rains:Risk</i> (IE1) or <i>rains:Bias</i> (IE2) to a string literal]</p>
<p>What are the risk, including bias, estimates of the ML model before and after mitigation?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rains:Risk</i> (IE1); <i>rains:Bias</i> (IE2)]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rains:Risk</i> (IE1); <i>rains:Bias</i> (IE2)]</p> <p>-</p> <p>While this CQ is modelled using the same concepts as the CQ ‘What are the risks, including biases, of the ML Model?’, new instances of the two information elements <i>rains:Risk</i> and <i>rains:Bias</i> must be created to record the new estimates.</p>
<p>What are the trade-offs between risks, including biases, and the effects of using the ML model on humans or others?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rains:TradeOff</i> (IE1)]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rains:TradeOff</i> (IE1)]</p>
<p>What are the characteristics of the ML model relating to emphasized or excluded features, decision thresholds and basis of prioritization, as well as its inputs and outputs and programming methods?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>mils:ModelCharacteristic</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i>^(op) to <i>mils:Model</i> (IE1); the input and output formats of the ML model are indicated by linking <i>mils:Model</i> (IE1) to a String literal using <i>rains:hasInputFormat</i>^(dp) and <i>rains:hasOutputFormat</i>^(dp)]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>mils:ModelCharacteristic</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i>^(op) to <i>mils:Model</i> (IE1); the input and output formats of the ML model are indicated by linking <i>mils:Model</i> (IE1) to a String literal using <i>rains:hasInputFormat</i>^(dp) and <i>rains:hasOutputFormat</i>^(dp)]</p>
<p>What algorithm is used by the ML model?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>mils:Algorithm</i> (IE1)]</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>mils:Algorithm</i> (IE1)]</p>

TABLE 11. (Continued.) Information about the ML model.

What are the ML model's parameters and hyper-parameters?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rains:Parameter</i> (IE1); <i>mls:HyperParameter</i> (IE2)]
What are the sources of additional information and citation details for the ML model?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rains:Parameter</i> (IE1); <i>mls:HyperParameter</i> (IE2)]
What are the sources of additional information and citation details for the ML model?	[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rdfs:seeAlso</i> ^(op) linking <i>mls:Model</i> (IE1) to the relevant entity]
	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rdfs:seeAlso</i> ^(op) linking <i>mls:Model</i> (IE1) to the relevant entity]

^bOur approach requires the URI of the person who performed the action to be recorded, but – unlike FactSheets – does not ask about the characteristics of the development team or whether there were domain experts involved. To enforce capture of whether the agent(s) were domain experts, *rains:HumanConstraint* may be used. Descriptions of agents may be expanded using the FOAF vocabulary to list their contact information and company affiliations and any additional characteristics can be recorded using *rdfs:comment*^(dp) linking *sao:AccountableAgent* to a String literal.

stakeholders involved in the AI system life cycle, and it is further complicated by various (sometimes incompatible) organizational structures and procedures including those preventing the disclosure of sensitive information. While we do not provide a solution to these organizational challenges, our approach does remove some of the technological barriers associated with gathering accountability information. For example, Hind *et al.* [54] identified several barriers relating to the adoption of FactSheets during evaluations with potential users of this framework. These included the need to reduce the time and effort during data collection; the need to address the difference between facts that can be captured automatically and those requiring “human authoring” which sometimes leads to incomplete and poor-quality results; the need to enable data capture at the time of creation of system components rather than having to recall it during later stages; the need for an open collection API that would support an unbounded set of tools that AI developers use to build their systems; and also the ability for the information captured by such systems to meet the different perspectives of stakeholders (e.g., some users would benefit from simplified views).

As we have outlined previously, our semantic approach overcomes all of these barriers as it offers a standard and domain-extendable vocabulary that can be used by external tools to record the collected accountability information as parts of the accountability KG, which can then be submitted to a central semantic middleware layer. In addition, such tools can provide instant feedback to the users who generate accountability traces (e.g., when a required piece of information is missing) based on the constraints defined in the accountability plans as demonstrated in [46]. The KG is also capable of supporting a range of visualizations of the accountability information to users; this is enabled by the ability to perform structured queries over the types of information, chronology of the events, and the identification of missing information (i.e., where a plan expects a piece of information to be provided, but it is absent from the accountability trace). Moreover, our accountability KG associates the

information with individual human agents who can be held accountable for providing false or insufficient information or simply contacted for further clarification. This critical aspect is often limited to organization level (e.g., the company responsible for developing an AI system) in the baseline frameworks.

VI. CONCLUSION AND FUTURE WORK

We have presented an approach for recording accountability information relating to the design and implementation of AI systems using accountability knowledge graphs. We also presented a comparison between the capabilities of our approach and three popular baseline frameworks and have discussed how our approach offers opportunities to integrate and extend the information provided by these frameworks. We have also described how our approach overcomes some of the technological barriers associated with collecting accountability information.

In future work, we will focus on extending the semantic vocabularies to cover the deployment and operation stages of the AI system life cycle. We will explore further opportunities to demonstrate the integration of information provided by external tools, especially for the operation stage, where we expect the greatest benefit of automated accountability information capture. As new stakeholders will be captured in the accountability KG (e.g., subjects of the final AI decisions) we will build on our initial round table discussions to conduct a series of user-based evaluations to evaluate the practicality and usefulness of accountability KGs from both legal and technological perspectives.

IN MEMORIAM

Dr. Caitlin Doyle Cottrill (1977 - 2022) passed away at the age of forty-four following a long battle with cancer. She will always be remembered for her kindness, support, gentle humour, and commitment to creating a better world for all people and for the environment in which we all live.

TABLE 12. Information about the evaluation of the ML model.

Information relating to designing the evaluation of the ML model and performing the evaluation	Ontological Modelling
Who produced the ML model evaluation specification and on behalf of whom or which organization?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with four inputs <i>rains:SystemPurposeSpecification</i> (IN1), <i>rains:ComplianceSpecification</i> (IN2), <i>rains:DatasetSpecification</i> (IN3), and <i>rains:ModelSpecification</i> (IN4); <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the <i>sao:InformationRealization</i> corresponding to <i>rains:EvaluationSpecification</i> (AR1). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf</i> ^(op) .]
Who evaluated the ML model and on behalf of whom or which organization?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Evaluate</i> (AA) with three inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), and <i>rains:EvaluationDataset</i> (IN3); <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the two <i>sao:InformationRealizations</i> corresponding to the two output variables of <i>rains:Evaluate</i> , namely <i>rains:Evaluation</i> (AR1) and <i>rains:ChangeLog</i> (AR2). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf</i> ^(op) .]
Was the ML model evaluation different from its design specification, and if yes then why?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Evaluate</i> (AA) with three inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), and <i>rains:EvaluationDataset</i> (IN3); <i>rains:ChangeLog</i> (AR2); <i>sao:InformationRealization</i> corresponding to <i>rains:ChangeLog</i> ; <i>rains:ChangeJustification</i> (IE1)]
What was the evaluation procedure, including evaluation factors and scenarios?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with four inputs <i>rains:SystemPurposeSpecification</i> (IN1), <i>rains:ComplianceSpecification</i> (IN2), <i>rains:DatasetSpecification</i> (IN3), and <i>rains:ModelSpecification</i> (IN4); <i>rains:EvaluationSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:EvaluationSpecification</i> (AR1); <i>mls:EvaluationProcedure</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Evaluate</i> (AA) with three inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), and <i>rains:EvaluationDataset</i> (IN3); <i>rains:Evaluation</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:Evaluation</i> (AR1); <i>mls:EvaluationProcedure</i> (IE1)]
How were the evaluation results measured? If possible, justifying reasons for choosing performance metrics and describing how accuracy is measured and assessed and how thresholds for failure are defined.	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with four inputs <i>rains:SystemPurposeSpecification</i> (IN1), <i>rains:ComplianceSpecification</i> (IN2), <i>rains:DatasetSpecification</i> (IN3), and <i>rains:ModelSpecification</i> (IN4); <i>rains:EvaluationSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:EvaluationSpecification</i> (AR1); <i>mls:EvaluationMeasure</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Evaluate</i> (AA) with three inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), and <i>rains:EvaluationDataset</i> (IN3); <i>rains:Evaluation</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:Evaluation</i> (AR1); <i>mls:EvaluationMeasure</i> (IE1)]
What were the evaluation results? If possible, with respect to each evaluation factor, and including confidence levels, false positive/negative ratios, and any decision variability under same conditions.	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with four inputs <i>rains:SystemPurposeSpecification</i> (IN1), <i>rains:ComplianceSpecification</i> (IN2), <i>rains:DatasetSpecification</i> (IN3), and <i>rains:ModelSpecification</i> (IN4); <i>rains:EvaluationSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:EvaluationSpecification</i> (AR1); the following are used to model desired results or known previous evaluation results of a reused ML model: <i>mls:ModelEvaluation</i> (IE1); <i>rains:hasEvaluationResult</i> ^(op) linking <i>mls:EvaluationMeasure</i> to <i>mls:ModelEvaluation</i> (IE1); and <i>rains:computedOnSlice</i> ^(dp) , <i>rains:computedOnDecisionThreshold</i> ^(dp) , <i>rains:hasResultUpperBound</i> ^(dp) , <i>rains:hasResultLowerBound</i> ^(dp) , and <i>rains:hasResultValue</i> ^(dp) linking from <i>mls:ModelEvaluation</i> (IE1) to literals.] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Evaluate</i> (AA) with three inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), and <i>rains:EvaluationDataset</i> (IN3); <i>rains:Evaluation</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:Evaluation</i> (AR1); <i>mls:ModelEvaluation</i> (IE1); <i>rains:hasEvaluationResult</i> ^(op) linking <i>mls:EvaluationMeasure</i> to <i>mls:ModelEvaluation</i> (IE1); <i>rains:computedOnSlice</i> ^(dp) , <i>rains:computedOnDecisionThreshold</i> ^(dp) , <i>rains:hasResultUpperBound</i> ^(dp) , <i>rains:hasResultLowerBound</i> ^(dp) , and <i>rains:hasResultValue</i> ^(dp) linking from <i>mls:ModelEvaluation</i> (IE1) to literals.]
What are the changes in performance metric values before and after mitigation of risks and biases?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with four inputs <i>rains:SystemPurposeSpecification</i> (IN1), <i>rains:ComplianceSpecification</i> (IN2), <i>rains:DatasetSpecification</i> (IN3), and <i>rains:ModelSpecification</i> (IN4); <i>rains:EvaluationSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:EvaluationSpecification</i> (AR1); (Optional) <i>mls:ModelEvaluation</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Evaluate</i> (AA) with three inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), and <i>rains:EvaluationDataset</i> (IN3); <i>rains:Evaluation</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:Evaluation</i> (AR1); <i>mls:ModelEvaluation</i> (IE1); <i>rains:hasEvaluationResult</i> ^(op) linking <i>mls:EvaluationMeasure</i> to <i>mls:ModelEvaluation</i> (IE1); <i>rains:computedOnSlice</i> ^(dp) , <i>rains:computedOnDecisionThreshold</i> ^(dp) , <i>rains:hasResultUpperBound</i> ^(dp) , <i>rains:hasResultLowerBound</i> ^(dp) , and <i>rains:hasResultValue</i> ^(dp) linking from <i>mls:ModelEvaluation</i> (IE1) to literals.] ----- While this CQ is modelled using the same concepts as the CQ ‘What were the evaluation results?’, new instances of the information elements must be created to record the new values.

^bOur approach requires the URI of the person who performed the action to be recorded, but – unlike FactSheets – does not ask about the characteristics of the development team or whether there were domain experts involved. To enforce capture of whether the agent(s) were domain experts, *rains:HumanConstraint* may be used. Descriptions of agents may be expanded using the FOAF vocabulary to list their contact information and company affiliations and any additional characteristics can be recorded using *rdfs:comment*^(dp) linking *sao:AccountableAgent* to a String literal.

**APPENDIX
HOW THE ONTOLOGIES ANSWER THE COMPETENCY
QUESTIONS**

In this appendix we present the full list of competency questions (CQs) and illustrate how our approach models

the corresponding accountability information. Table 8 lists the ontologies we use, along with their prefixes and namespaces. Figure 5 showcases how the entities, activities, and agents relate to each other, and where the concepts *sao:AccountabilityPlan*, *sao:AccountableAction*,

TABLE 13. Information about the supporting infrastructure.

Information relating to the design and implementation of the supporting infrastructure	Ontological Modelling
Who produced the supporting infrastructure design specification and on behalf of whom or which organization?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1), <i>rains:ComplianceSpecification</i> (IN2); <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf</i> ^(op) .]
Who realized the supporting infrastructure and on behalf of whom or which organization?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the two <i>sao:InformationRealizations</i> corresponding to the two output variables of <i>rains:RealizeComponent</i> , namely <i>rains:SupportingInfrastructureComponent</i> (AR1) and <i>rains:ChangeLog</i> (AR2). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf</i> ^(op) .]
Was the supporting infrastructure realization different from its design specification, and if yes then why?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:ChangeLog</i> (AR2); <i>sao:InformationRealization</i> corresponding to <i>rains:ChangeLog</i> (AR2); <i>rains:ChangeJustification</i> (IE1); the instance of the designed but not realized <i>rains:SupportingInfrastructure</i> in the information realization corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1) becomes an instance of <i>rains:ExcludedObject</i>]
What is the supporting infrastructure, i.e. the non core-AI components of this system?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>rains:SupportingInfrastructure</i> (IE1). The input and output formats of the AI system are indicated using <i>rains:hasInputFormat</i> ^(dp) and <i>rains:hasOutputFormat</i> ^(dp) linked to a String literal] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>rains:SupportingInfrastructure</i> (IE1); the instance of the realized <i>rains:SupportingInfrastructure</i> is modelled as an instance of <i>rains:RealizedObject</i> also. The input and output formats of the AI system are indicated using <i>rains:hasInputFormat</i> ^(dp) and <i>rains:hasOutputFormat</i> ^(dp) linked to a String literal]
Was the supporting infrastructure created specifically for this AI system or was it reused?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>rains:isReusedObject</i> ^(dp) linking <i>rains:SupportingInfrastructure</i> (IE1) to a Boolean literal of a <i>true</i> value if the supporting infrastructure is reused.] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>rains:isReusedObject</i> ^(dp) linking <i>rains:SupportingInfrastructure</i> (IE1) to a Boolean literal of a <i>true</i> value if the supporting infrastructure is reused.]
If the supporting infrastructure was reused, who is its creator or manager?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:AccountableAgent</i> ^b (AG2) linked using <i>sao:isAccountableFor</i> ^(op) to <i>rains:SupportingInfrastructure</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:AccountableAgent</i> ^b (AG2) linked using <i>sao:isAccountableFor</i> ^(op) to <i>rains:SupportingInfrastructure</i> (IE1)]
What is the version or date of the supporting infrastructure?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>rains:version</i> ^(dp) linking <i>rains:SupportingInfrastructure</i> (IE1) to a String literal; <i>rains:versionDate</i> ^(dp) linking <i>rains:SupportingInfrastructure</i> (IE1) to a date-time literal; <i>rains:versionNote</i> ^(dp) linking <i>rains:SupportingInfrastructure</i> (IE1) to a string literal] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>rains:version</i> ^(dp) linking <i>rains:SupportingInfrastructure</i> (IE1) to a String literal; <i>rains:versionDate</i> ^(dp) linking <i>rains:SupportingInfrastructure</i> (IE1) to a date-time literal; <i>rains:versionNote</i> ^(dp) linking <i>rains:SupportingInfrastructure</i> (IE1) to a string literal]
What are the limitations of the supporting infrastructure?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>rains:Limitation</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>rains:Limitation</i> (IE1)]
What are the risks, including biases, of the supporting infrastructure (e.g., were sensors tuned to recognize different skin tones)?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>rains:Risk</i> (IE1); <i>rains:Bias</i> (IE2)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>rains:Risk</i> (IE1); <i>rains:Bias</i> (IE2)]

TABLE 13. (Continued.) Information about the supporting infrastructure.

<p>What risk and bias software was used to analyze the supporting infrastructure?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>rains:FairnessOrBiasDetectionSoftware</i> (IE1)]</p>
<p>What are the mitigation strategies to address the risks and biases of the supporting infrastructure?</p>	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>rains:FairnessOrBiasDetectionSoftware</i> (IE1)]</p>
<p>What are the risk, including bias, estimates of the supporting infrastructure before and after mitigation?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>rains:hasMitigationStrategy</i>^(op) linking <i>rains:Risk</i> (IE1) or <i>rains:Bias</i> (IE2) to a string literal]</p>
<p>What are the risk, including bias, estimates of the supporting infrastructure before and after mitigation?</p>	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>rains:hasMitigationStrategy</i>^(op) linking <i>rains:Risk</i> (IE1) or <i>rains:Bias</i> (IE2) to a string literal]</p>
<p>What are the human oversight and human-in-control mechanisms of this AI system?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>rains:SystemHumanOversightMechanism</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i>^(op) to <i>rains:SupportingInfrastructure</i> (IE1)]</p>
<p>What are the human oversight and human-in-control mechanisms of this AI system?</p>	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>rains:SystemHumanOversightMechanism</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i>^(op) to <i>rains:SupportingInfrastructure</i> (IE1)]</p>
<p>What are the security mechanisms of this AI system?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>rains:SystemSecurityMechanism</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i>^(op) to <i>rains:SupportingInfrastructure</i> (IE1)]</p>
<p>What are the security mechanisms of this AI system?</p>	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>rains:SystemSecurityMechanism</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i>^(op) to <i>rains:SupportingInfrastructure</i> (IE1)]</p>
<p>What are the audit mechanisms of this AI system?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>rains:SystemAuditMechanism</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i>^(op) to <i>rains:SupportingInfrastructure</i> (IE1)]</p>
<p>What are the audit mechanisms of this AI system?</p>	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>rains:SystemAuditMechanism</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i>^(op) to <i>rains:SupportingInfrastructure</i> (IE1)]</p>
<p>What are the explanation mechanisms provided by this AI system?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>rains:SystemExplanationMechanism</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i>^(op) to <i>rains:SupportingInfrastructure</i> (IE1)]</p>
<p>What are the explanation mechanisms provided by this AI system?</p>	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>rains:SystemExplanationMechanism</i> (IE2) linked using <i>rains:hasRealizableObjectCharacteristic</i>^(op) to <i>rains:SupportingInfrastructure</i> (IE1)]</p>

^b Our approach requires the URI of the person who performed the action to be recorded, but – unlike FactSheets – does not ask about the characteristics of the development team or whether there were domain experts involved. To enforce capture of whether the agent(s) were domain experts, *rains:HumanConstraint* may be used. Descriptions of agents may be expanded using the FOAF vocabulary to list their contact information and company affiliations and any additional characteristics can be recorded using *rdfs:comment*^(op) linking *sao:AccountableAgent* to a String literal.

TABLE 14. Information about the evaluation of the whole system.

Information relating to designing the system evaluation and performing the evaluation	Ontological Modelling
Who produced the system evaluation specification and on behalf of whom or which organization?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with five inputs <i>rains:SystemPurposeSpecification</i> (IN1), <i>rains:ComplianceSpecification</i> (IN2), <i>rains:DatasetSpecification</i> (IN3), <i>rains:ModelSpecification</i> (IN4), and <i>rains:SupportingInfrastructureSpecification</i> (IN5); <i>sao:AccountableAgent</i> ^P (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the <i>sao:InformationRealization</i> corresponding to <i>rains:EvaluationSpecification</i> (AR1). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf</i> ^(op) .]
Who evaluated the system and on behalf of whom or which organization?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Evaluate</i> (AA) with four inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:EvaluationDataset</i> (IN3), and <i>rains:SupportingInfrastructureComponent</i> (IN4); <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the two <i>sao:InformationRealizations</i> corresponding to the two output variables of <i>rains:Evaluate</i> , namely <i>rains:Evaluation</i> (AR1) and <i>rains:ChangeLog</i> (AR2). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf</i> ^(op) .]
Was the system evaluation different from its design specification, and if yes then why?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Evaluate</i> (AA) with four inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:EvaluationDataset</i> (IN3), and <i>rains:SupportingInfrastructureComponent</i> (IN4); <i>rains:ChangeLog</i> (AR2); <i>sao:InformationRealization</i> corresponding to <i>rains:ChangeLog</i> (AR2); <i>rains:ChangeJustification</i> (IE1)]
Was the system evaluated by a third-party, and if yes by whom?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Evaluate</i> (AA) with two inputs <i>rains:ModelComponent</i> (IN1) and <i>rains:SupportingInfrastructureComponent</i> (IN2); <i>rains:ThirdPartyEvaluation</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ThirdPartyEvaluation</i> (AR1); <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the <i>sao:InformationRealization</i> corresponding to <i>rains:ThirdPartyEvaluation</i> (AR1)]
What was the evaluation procedure, including evaluation factors, domains and applications, key performance indicators tests and non-functional tests?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with five inputs <i>rains:SystemPurposeSpecification</i> (IN1), <i>rains:ComplianceSpecification</i> (IN2), <i>rains:DatasetSpecification</i> (IN3), <i>rains:ModelSpecification</i> (IN4), and <i>rains:SupportingInfrastructureSpecification</i> (IN5); <i>rains:EvaluationSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:EvaluationSpecification</i> (AR1); <i>rains:EvaluationProcedure</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Evaluate</i> (AA) with four inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:EvaluationDataset</i> (IN3), and <i>rains:SupportingInfrastructureComponent</i> (IN4); <i>rains:Evaluation</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:Evaluation</i> (AR1); <i>rains:EvaluationProcedure</i> (IE1)]
How were the evaluation results measured? If possible, justify reasons for choosing performance metrics and describe how accuracy is measured and assessed and how thresholds for failure are defined.	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with five inputs <i>rains:SystemPurposeSpecification</i> (IN1), <i>rains:ComplianceSpecification</i> (IN2), <i>rains:DatasetSpecification</i> (IN3), <i>rains:ModelSpecification</i> (IN4), and <i>rains:SupportingInfrastructureSpecification</i> (IN5); <i>rains:EvaluationSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:EvaluationSpecification</i> (AR1); <i>rains:EvaluationMeasure</i> (IE1)] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Evaluate</i> (AA) with four inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:EvaluationDataset</i> (IN3), and <i>rains:SupportingInfrastructureComponent</i> (IN4); <i>rains:Evaluation</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:Evaluation</i> (AR1); <i>rains:EvaluationMeasure</i> (IE1)]
What were the evaluation results, if possible, with respect to each evaluation factor?	[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:ProduceSpecification</i> (AA) with five inputs <i>rains:SystemPurposeSpecification</i> (IN1), <i>rains:ComplianceSpecification</i> (IN2), <i>rains:DatasetSpecification</i> (IN3), <i>rains:ModelSpecification</i> (IN4), and <i>rains:SupportingInfrastructureSpecification</i> (IN5); <i>rains:EvaluationSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:EvaluationSpecification</i> (AR1); the following are used to model desired results: <i>rains:EvaluationResult</i> (IE1), with <i>rains:hasResultUpperBound</i> ^(dp) , <i>rains:hasResultLowerBound</i> ^(dp) , and <i>rains:hasResultValue</i> ^(dp) linking from <i>rains:EvaluationResult</i> (IE1) to literals.] ----- [<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Evaluate</i> (AA) with four inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:EvaluationDataset</i> (IN3), and <i>rains:SupportingInfrastructureComponent</i> (IN4); <i>rains:Evaluation</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:Evaluation</i> (AR1); <i>rains:EvaluationResult</i> (IE2); <i>rains:hasEvaluationResult</i> ^(op) linking <i>rains:EvaluationMeasure</i> (IE1) to <i>rains:EvaluationResult</i> (IE2); <i>rains:hasResultUpperBound</i> ^(dp) , <i>rains:hasResultLowerBound</i> ^(dp) , and <i>rains:hasResultValue</i> ^(dp) linking from <i>rains:EvaluationResult</i> (IE2) to literals.]

^bOur approach requires the URI of the person who performed the action to be recorded, but – unlike FactSheets – does not ask about the characteristics of the development team or whether there were domain experts involved. To enforce capture of whether the agent(s) were domain experts, *rains:HumanConstraint* may be used. Descriptions of agents may be expanded using the FOAF vocabulary to list their contact information and company affiliations and any additional characteristics can be recorded using *rdfs:comment*^(dp) linking *sao:AccountableAgent* to a String literal.

TABLE 15. Information about certifications.

Information relating to certifications	Ontological Modelling
What was the result of the certification process?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:GrantCertification</i> (AA); <i>rains:CertificationDecision</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:CertificationDecision</i> (AR1); <i>rains:Certificate</i> (IE1)]
Who granted the certificate?	<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:GrantCertification</i> (AA); <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:GrantCertification</i> , namely <i>rains:CertificationDecision</i> (AR1)]

^bOur approach requires the URI of the person who performed the action to be recorded, but – unlike FactSheets – does not ask about the characteristics of the development team or whether there were domain experts involved. To enforce capture of whether the agent(s) were domain experts, *rains:HumanConstraint* may be used. Descriptions of agents may be expanded using the FOAF vocabulary to list their contact information and company affiliations and any additional characteristics can be recorded using *rdfs:comment*^(dp) linking *sao:AccountableAgent* to a String literal.

TABLE 16. Information relating to the deployment and operational guidelines.

Information relating to the deployment and operational guidelines	Ontological Modelling
Information about the deployment guidelines	
Who produced the deployment guidelines for installing the AI system and integrating it (e.g., within a larger IT ecosystem) and on behalf of whom or which organization?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:GenerateGuideline</i> (AA) with four inputs <i>rains:DatasetComponent</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:SupportingInfrastructureComponent</i> (IN3), and <i>rains:Evaluation</i> (IN4); <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:GenerateGuideline</i> , namely <i>rains:SystemInstallationGuideline</i> (AR1). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf</i> ^(op) .]
What is the guidance for installing this AI system and integrating it (e.g., within a larger IT ecosystem)?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:GenerateGuideline</i> (AA) with four inputs <i>rains:DatasetComponent</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:SupportingInfrastructureComponent</i> (IN3), and <i>rains:Evaluation</i> (IN4); <i>rains:SystemInstallationGuideline</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SystemInstallationGuideline</i> (AR1); <i>rains:SystemInstallationGuidance</i> (IE1)].
Who produced the deployment guidelines for training users to use and maintain this AI system?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:GenerateGuideline</i> (AA) with four inputs <i>rains:DatasetComponent</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:SupportingInfrastructureComponent</i> (IN3), and <i>rains:Evaluation</i> (IN4); <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:GenerateGuideline</i> , namely <i>rains:UserTrainingGuideline</i> (AR1). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf</i> ^(op) .]
What is the guidance for training users to use and maintain this AI system?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:GenerateGuideline</i> (AA) with four inputs <i>rains:DatasetComponent</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:SupportingInfrastructureComponent</i> (IN3), and <i>rains:Evaluation</i> (IN4); <i>rains:UserTrainingGuideline</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:UserTrainingGuideline</i> (AR1); <i>rains:UserTrainingGuidance</i> (IE1)].
Information about the operational guidelines	
Who produced the operational guidelines for how to use this A system and on behalf of whom or which organization?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:GenerateGuideline</i> (AA) with four inputs <i>rains:DatasetComponent</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:SupportingInfrastructureComponent</i> (IN3), and <i>rains:Evaluation</i> (IN4); <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:GenerateGuideline</i> , namely <i>rains:SystemUsageGuideline</i> (AR1). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf</i> ^(op) .]
What is the guidance for using this AI system?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:GenerateGuideline</i> (AA) with four inputs <i>rains:DatasetComponent</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:SupportingInfrastructureComponent</i> (IN3), and <i>rains:Evaluation</i> (IN4); <i>rains:SystemUsageGuideline</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SystemUsageGuideline</i> (AR1); <i>rains:SystemUsageGuidance</i> (IE1)]
Who produced the operational guidelines for how to maintain this AI system and on behalf of whom or which organization?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:GenerateGuideline</i> (AA) with four inputs <i>rains:DatasetComponent</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:SupportingInfrastructureComponent</i> (IN3), and <i>rains:Evaluation</i> (IN4); <i>sao:AccountableAgent</i> ^b (AG1) linked using <i>sao:isAccountableFor</i> ^(op) to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:GenerateGuideline</i> , namely <i>rains:SystemMaintenanceGuideline</i> (AR1). Modelling that <i>sao:AccountableAgent</i> (AG1) acts on behalf of another agent or organization is done by linking <i>sao:AccountableAgent</i> (AG1) to the other agent or organization using <i>prov:actedOnBehalfOf</i> ^(op) .]
What is the guidance for maintenance of this AI system?	[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:GenerateGuideline</i> (AA) with four inputs <i>rains:DatasetComponent</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:SupportingInfrastructureComponent</i> (IN3), and <i>rains:Evaluation</i> (IN4); <i>rains:SystemMaintenanceGuideline</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SystemMaintenanceGuideline</i> (AR1); <i>rains:SystemMaintenanceGuidance</i> (IE1)]

^bOur approach requires the URI of the person who performed the action to be recorded, but – unlike FactSheets – does not ask about the characteristics of the development team or whether there were domain experts involved. To enforce capture of whether the agent(s) were domain experts, *rains:HumanConstraint* may be used. Descriptions of agents may be expanded using the FOAF vocabulary to list their contact information and company affiliations and any additional characteristics can be recorded using *rdfs:comment*^(dp) linking *sao:AccountableAgent* to a String literal.

sao:AccountableResult, and *sao:InformationElement* are substituted by their appropriate *RAInS* subclasses. This is done as follows:

A *sao:System* is created and attributed to an *sao:Agent*. An associated *sao:AccountabilityPlan* (AP) is created for the system. The created system would be of type *rains:AI_System*. Given the current scope of *RAInS*, the *accountability plan* would be either *rains:DesignStageAccountabilityPlan* or *rains:ImplementationStageAccountabilityPlan*.

A *sao:AccountabilityPlan* contains some *sao:AccountableActions* (AA), each having at least one accountable result output. The first output would be *sao:AccountableResult* (AR1), and if there is a second output, this is shown in Figure 5 as *sao:AccountableResult* (AR2). For example, when the accountable action *rains:RealizeComponent*

is planned to realize a dataset, it would have as its outputs the two accountable results *rains:DatasetComponent* (AR1) and *rains:ChangeLog* (AR2). Where applicable, some *sao:AccountableResults* can be used as inputs to subsequent *sao:AccountableActions*. This is indicated in Figure 5 with IN1-n, where IN1 would be the first input, IN2 the second input, and INn the last input. For example, when the accountable action *rains:Evaluate* is planned to evaluate the ML model, it would have as its inputs *rains:EvaluationSpecification* (IN1), *rains:ModelComponent* (IN2), and *rains:EvaluationDataset* (IN3). Although not shown in Figure 5, when the *accountable actions* and *accountable results* are created in the plan, accountability information to be recorded in the *accountability trace* is indicated in a human-readable plain text format using the data property *rdfs:comment*. Also not shown in the

TABLE 17. Information about the component-specific compliance requirements.

Information relating to the component-specific compliance requirements	Ontological Modelling
What hard and soft laws does the dataset comply with?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:DatasetSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetSpecification</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:DatasetSpecification</i> (IN1) as an input; <i>rains:DatasetComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:DatasetComponent</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p>
What hard and soft laws does the ML model comply with?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ComplianceSpecification</i> (IN2); <i>rains:ModelSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelSpecification</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with three inputs <i>rains:ModelSpecification</i> (IN1), <i>rains:TrainingDataset</i> (IN2), and <i>rains:ValidationDataset</i> (IN3); <i>rains:ModelComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:ModelComponent</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p>
What hard and soft laws does the ML model evaluation comply with?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with four inputs <i>rains:SystemPurposeSpecification</i> (IN1), <i>rains:ComplianceSpecification</i> (IN2), <i>rains:DatasetSpecification</i> (IN3), and <i>rains:ModelSpecification</i> (IN4); <i>rains:EvaluationSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:EvaluationSpecification</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Evaluate</i> (AA) with three inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), and <i>rains:EvaluationDataset</i> (IN3); <i>rains:Evaluation</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:Evaluation</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p>
What hard and soft laws does the supporting infrastructure comply with?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureSpecification</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:RealizeComponent</i> (AA) with <i>rains:SupportingInfrastructureSpecification</i> (IN1) as an input; <i>rains:ChangeLog</i> (AR2); <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SupportingInfrastructureComponent</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p>
What hard and soft laws does the system evaluation comply with?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:ProduceSpecification</i> (AA) with five inputs <i>rains:SystemPurposeSpecification</i> (IN1), <i>rains:ComplianceSpecification</i> (IN2), <i>rains:DatasetSpecification</i> (IN3), <i>rains:ModelSpecification</i> (IN4), and <i>rains:SupportingInfrastructureSpecification</i> (IN4); <i>rains:EvaluationSpecification</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:EvaluationSpecification</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Evaluate</i> (AA) with four inputs <i>rains:EvaluationSpecification</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:EvaluationDataset</i> (IN3), and <i>rains:SupportingInfrastructureComponent</i> (IN4); <i>rains:Evaluation</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:Evaluation</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p>
What hard and soft laws do the deployment guidelines for installing and integrating this AI system comply with?	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:GenerateGuideline</i> (AA) with four inputs <i>rains:DatasetComponent</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:SupportingInfrastructureComponent</i> (IN3), and <i>rains:Evaluation</i> (IN4); <i>rains:SystemInstallationGuideline</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SystemInstallationGuideline</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p>
What hard and soft laws do the deployment guidelines for training the users and maintainers of this AI system comply with?	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:GenerateGuideline</i> (AA) with four inputs <i>rains:DatasetComponent</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:SupportingInfrastructureComponent</i> (IN3), and <i>rains:Evaluation</i> (IN4); <i>rains:UserTrainingGuideline</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:UserTrainingGuideline</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p>
What hard and soft laws do the operational guidelines for how to use this AI system comply with?	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:GenerateGuideline</i> (AA) with four inputs <i>rains:DatasetComponent</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:SupportingInfrastructureComponent</i> (IN3), and <i>rains:Evaluation</i> (IN4); <i>rains:SystemUsageGuideline</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SystemUsageGuideline</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p>
What hard and soft laws do the operational guidelines for how to maintain this AI system comply with?	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:GenerateGuideline</i> (AA) with four inputs <i>rains:DatasetComponent</i> (IN1), <i>rains:ModelComponent</i> (IN2), <i>rains:SupportingInfrastructureComponent</i> (IN3), and <i>rains:Evaluation</i> (IN4); <i>rains:SystemMaintenanceGuideline</i> (AR1); <i>sao:InformationRealization</i> corresponding to <i>rains:SystemMaintenanceGuideline</i> (AR1); <i>rains:HardLaw</i> (IE1); <i>rains:SoftLaw</i> (IE2)]</p>

^b Our approach requires the URI of the person who performed the action to be recorded, but – unlike FactSheets – does not ask about the characteristics of the development team or whether there were domain experts involved. To enforce capture of whether the agent(s) were domain experts, *rains:HumanConstraint* may be used. Descriptions of agents may be expanded using the FOAF vocabulary to list their contact information and company affiliations and any additional characteristics can be recorded using *rdfs:comment*^(dip) linking *sao:AccountableAgent* to a String literal.

figure, are the constraints to be satisfied in the *accountability trace* which would be of type *rains:AutoConstraint* or *rains:HumanConstraint*.

At the *accountability trace* level, an *ep-plan:Activity* is created to correspond to each *sao:AccountableAction*. This activity is associated with an *sao:AccountableAgent* (AG1) and generates one *sao:InformationRealization* for each planned *sao:AccountableResult* output. Thus, the

sao:AccountableAgent (AG1) which is associated with the *ep-plan:Activity* is also accountable for the generated *sao:InformationRealization*(s). Note that any accountable agent may act on behalf of another person or organization, resulting in them being accountable as well. Accountability information about the *sao:InformationRealization* is recorded in plain text format using the data property *rdfs:comment*. These comments capture the human

TABLE 18. Information about human decision making and approvals.

Information relating to human decision making and approvals	Ontological Modelling
Who confirmed the fitness of the dataset for the use cases indicated by the AI system’s purpose specification?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:DatasetSpecification</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:FitnessConfirmationDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:DatasetComponent</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:FitnessConfirmationDecision</i> (AR1)]</p>
Who assured the compliance of the dataset to the hard and soft laws as indicated by the AI system compliance requirements?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:ComplianceSpecification</i> (IN1) and <i>rains:DatasetSpecification</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ComplianceAssuranceDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:ComplianceSpecification</i> (IN1) and <i>rains:DatasetComponent</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ComplianceAssuranceDecision</i> (AR1)]</p>
Who approved the dataset?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:FitnessConfirmationDecision</i> (IN1) and <i>rains:ComplianceAssuranceDecision</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:FitnessConfirmationDecision</i> (IN1) and <i>rains:ComplianceAssuranceDecision</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p>
Who confirmed the fitness of the ML Model for the use cases indicated by the AI system’s purpose specification?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ModelSpecification</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:FitnessConfirmationDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:ModelComponent</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:FitnessConfirmationDecision</i> (AR1)]</p>
Who assured the compliance of the ML model to the hard and soft laws as indicated by the AI system compliance requirements?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:ComplianceSpecification</i> (IN1) and <i>rains:ModelSpecification</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ComplianceAssuranceDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:ComplianceSpecification</i> (IN1) and <i>rains:ModelComponent</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ComplianceAssuranceDecision</i> (AR1)]</p>
Who approved the ML model?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:FitnessConfirmationDecision</i> (IN1) and <i>rains:ComplianceAssuranceDecision</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:FitnessConfirmationDecision</i> (IN1) and <i>rains:ComplianceAssuranceDecision</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p>
Who confirmed the fitness of the ML Model evaluation for the use cases indicated by the AI system’s purpose specification?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:EvaluationSpecification</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:FitnessConfirmationDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:Evaluation</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:FitnessConfirmationDecision</i> (AR1)]</p>
Who assured the compliance of the ML model evaluation to the hard and soft laws as indicated by the AI system compliance requirements?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:ComplianceSpecification</i> (IN1) and <i>rains:EvaluationSpecification</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ComplianceAssuranceDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:ComplianceSpecification</i> (IN1) and <i>rains:Evaluation</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ComplianceAssuranceDecision</i> (AR1)]</p>
Who approved the ML model evaluation?	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:FitnessConfirmationDecision</i> (IN1) and <i>rains:ComplianceAssuranceDecision</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP): <i>rains:Decide</i> (AA) with two inputs <i>rains:FitnessConfirmationDecision</i> (IN1) and <i>rains:ComplianceAssuranceDecision</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^{op}</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p>

TABLE 18. (Continued.) Information about human decision making and approvals.

<p>Who confirmed the fitness of the supporting infrastructure for the use cases indicated by the AI system’s purpose specification?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:SupportingInfrastructureSpecification</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:FitnessConfirmationDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:SupportingInfrastructureComponent</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:FitnessConfirmationDecision</i> (AR1)]</p>
<p>Who assured the compliance of the supporting infrastructure for the hard and soft laws as indicated by the AI system compliance requirements?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with two inputs <i>rains:ComplianceSpecification</i> (IN1) and <i>rains:SupportingInfrastructureSpecification</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ComplianceAssuranceDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with two inputs <i>rains:ComplianceSpecification</i> (IN1) and <i>rains:SupportingInfrastructureComponent</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ComplianceAssuranceDecision</i> (AR1)]</p>
<p>Who approved the supporting infrastructure?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with two inputs <i>rains:FitnessConfirmationDecision</i> (IN1) and <i>rains:ComplianceAssuranceDecision</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with two inputs <i>rains:FitnessConfirmationDecision</i> (IN1) and <i>rains:ComplianceAssuranceDecision</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p>
<p>Who confirmed the fitness of the system evaluation for the use cases indicated by the AI system’s purpose specification?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:EvaluationSpecification</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:FitnessConfirmationDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with two inputs <i>rains:SystemPurposeSpecification</i> (IN1) and <i>rains:Evaluation</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:FitnessConfirmationDecision</i> (AR1)]</p>
<p>Who assured the compliance of the system evaluation to the hard and soft laws as indicated by the AI system compliance requirements?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with two inputs <i>rains:ComplianceSpecification</i> (IN1) and <i>rains:EvaluationSpecification</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ComplianceAssuranceDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with two inputs <i>rains:ComplianceSpecification</i> (IN1) and <i>rains:Evaluation</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ComplianceAssuranceDecision</i> (AR1)]</p>
<p>Who approved the system evaluation?</p>	<p>[<i>rains:DesignStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with two inputs <i>rains:FitnessConfirmationDecision</i> (IN1) and <i>rains:ComplianceAssuranceDecision</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p> <p>-----</p> <p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with two inputs <i>rains:FitnessConfirmationDecision</i> (IN1) and <i>rains:ComplianceAssuranceDecision</i> (IN2); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p>
<p>Who assured the compliance of the deployment guidelines to the hard and soft laws as indicated by the AI system compliance requirements?</p>	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with three inputs <i>rains:ComplianceSpecification</i> (IN1), <i>rains:SystemInstallationGuideline</i> (IN2), and <i>rains:UserTrainingGuideline</i> (IN3); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ComplianceAssuranceDecision</i> (AR1)]</p>
<p>Who approved the deployment guidelines?</p>	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with <i>rains:ComplianceAssuranceDecision</i> (IN1) as an input; <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p>
<p>Who assured the compliance of the operational guidelines to the hard and soft laws as indicated by the AI system compliance requirements?</p>	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with three inputs <i>rains:ComplianceSpecification</i> (IN1), <i>rains:SystemUsageGuideline</i> (IN2), and <i>rains:SystemMaintenanceGuideline</i> (IN3); <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ComplianceAssuranceDecision</i> (AR1)]</p>
<p>Who approved the operational guidelines?</p>	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with <i>rains:ComplianceAssuranceDecision</i> (IN1) as an input; <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p>
<p>Who approved the implementation of the whole AI system for deployment?</p>	<p>[<i>rains:ImplementationStageAccountabilityPlan</i> (AP); <i>rains:Decide</i> (AA) with previous <i>rains:ApprovalDecisions</i> as inputs; <i>sao:AccountableAgent^b</i>(AG1) linked using <i>sao:isAccountableFor^(op)</i> to the <i>sao:InformationRealization</i> corresponding to the output variable of <i>rains:Decide</i>, namely <i>rains:ApprovalDecision</i> (AR1)]</p>

^bOur approach requires the URI of the person who performed the action to be recorded, but – unlike FactSheets – does not ask about the characteristics of the development team or whether there were domain experts involved. To enforce capture of whether the agent(s) were domain experts, *rains:HumanConstraint* may be used. Descriptions of agents may be expanded using the FOAF vocabulary to list their contact information and company affiliations and any additional characteristics can be recorded using *rdfs:comment^(dp)* linking *sao:AccountableAgent* to a String literal.

readable accountability information, and where appropriate (e.g., with images) also as base64 strings. Each *sao:InformationRealization* is a collection of *sao:Information*

Elements (IE1, IE2, ...). Some *sao:InformationElements* may be attributed to a *sao:AccountableAgent* (AG2) different from the one who performed the *ep-plan:Activity* and thus

is accountable for this *sao:InformationElement*. For example, if the agent who produces a design specification for a dataset chooses a third-party off-the-shelf dataset, then the creator of this third-party dataset is accountable for only the dataset, but the first agent who chose to use that third-party dataset is still accountable for their specified design and their decision to include the dataset. The *sao:AccountableAgent* (AG2) is included as a subtype *sao:InformationElement* in the *sao:InformationRealization* collection.

Tables 9 to 16 show how the identified CQs can be answered using concepts from *SAO* and *RAInS*, as well as those imported from the ontologies listed in Table 8. In order to visualize how the concepts listed in the righthand columns of Tables 9 to 18 relate to each other, the reader should take the concepts listed in the tables and populate the template in Figure 5. Note that not all *sao:InformationRealizations* will have *sao:InformationElement* members. Also, where a reference to a *sao:InformationRealization* corresponding to the same *accountable result* is repeated, it is always the same *InformationRealization*. For example, the repeated mentions of *sao:InformationRealization* corresponding to *rains:DatasetComponent* (AR1) all refer to the same *InformationRealization*.

Note that properties are either object properties (op) or data properties (dp). Also note that the inputs we list in Tables 9 to 18 for accountable actions are what we consider to be the minimum inputs, and the plan creator may indicate additional required inputs as they see fit for their application domain and specific AI system. Moreover, some CQs were generalized and have been included in both the design and implementation stages. In the tables, we show how the modelling is done at the design stage using *rains:DesignStageAccountabilityPlan* (AP) and at the implementation stage using *rains:ImplementationStageAccountabilityPlan* (AP). Where a CQ is to be answered for both stages, both sets of concepts are included and are separated in the tables using a dashed line. Finally, recall that, because the PROV ontology can intrinsically model the dates and times that activities are performed - using the properties *prov:startedAtTime* and *prov:EndedAtTime*, and when entities are generated - using the property *prov:wasGeneratedAtTime*, we did not create CQs about when the different accountable actions are performed or when accountable results are generated.

REFERENCES

- [1] J. Dastin. (Oct. 2018). *Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women*. Reuters. Accessed: Sep. 9, 2021. [Online]. Available: <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/idUSKCN1MK08G>
- [2] T. Simonite. (Nov. 2018). *When it Comes to Gorillas, Google Photos Remains Blind*. Wired. Accessed: Nov. 9, 2021. [Online]. Available: <https://www.wired.com/story/when-it-comes-to-gorillas-google-photos-remains-blind/>
- [3] J. Snow. (Jul. 2018). *Amazon's Face Recognition Falsely Matched 28 Members of Congress With Mugshots*. ACLU. Accessed: Sep. 9, 2021. [Online]. Available: <https://www.aclu.org/blog/privacy-technology/surveillance-technologies/amazons-face-recognition-falsely-matched-28>
- [4] Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, "Dissecting racial bias in an algorithm used to manage the health of populations," *Science*, vol. 366, no. 6464, pp. 447–453, Oct. 2019, doi: [10.1126/science.aax2342](https://doi.org/10.1126/science.aax2342).
- [5] N. Singer and C. Metz. (Dec. 2019). *Many Facial-Recognition Systems are Biased, Says U.S. Study*. The New York Times. Accessed: Nov. 22, 2021. [Online]. Available: <https://www.nytimes.com/2019/12/19/technology/facial-recognition-bias.html>
- [6] R. Mac. (Sep. 2021). *Facebook Apologizes After A.I. Puts 'Primates' Label on Video of Black Men*. The New York Times. Accessed: Sep. 9, 2021. [Online]. Available: <https://www.nytimes.com/2021/09/03/technology/facebook-ai-raceprimates.html>
- [7] J. Sarlin. (Apr. 2021). *A False Facial Recognition Match Sent This Innocent Black Man to Jail*. CNN. Accessed: Sep. 9, 2021. [Online]. Available: <https://edition.cnn.com/2021/04/29/tech/nijeer-parks-facial-recognition-police-arrest/index.html>
- [8] N. Kayser-Bril. (Apr. 2021). *Google Apologizes After its Vision AI Produced Racist Results*. Algorithm Watch. Accessed: Nov. 22, 2021. [Online]. Available: <https://algorithmwatch.org/en/google-vision-racism/>
- [9] J. Vincent. (Nov. 2020). *AI Camera Operator Repeatedly Confuses Bald Head for Soccer Ball During Live Stream*. The Verge. Accessed: Sep. 9, 2021. [Online]. Available: <https://www.theverge.com/tldr/2020/11/3/21547392/aicamera-operator-football-bald-head-soccer-mistakes>
- [10] J. Vincent. (Sep. 2021). *Automated Hiring Software is Mistakenly Rejecting Millions of Viable Job Candidates*. The Verge. Accessed: Sep. 9, 2021. [Online]. Available: <https://www.theverge.com/2021/9/6/22659225/automated-hiring-software-rejecting-viable-candidates-harvard-business-school>
- [11] F. Shah. (Oct. 2021). *Couple Fined After Bus Lane Camera Photographs Text on Woman's Top*. Independent. Accessed: Oct. 25, 2021. [Online]. Available: <https://www.independent.co.uk/news/uk/home-news/couple-fined-camera-writing-b1940356.html>
- [12] P. McCausland. (Nov. 2019). *Self-Driving Uber Car That Hit and Killed Woman did Not Recognize That Pedestrians Jaywalk*. NBC News. Accessed: Sep. 9, 2021. [Online]. Available: <https://www.nbcnews.com/tech/tech-news/self-driving-uber-car-hit-killed-woman-did-not-recognize-n1079281>
- [13] A. J. Hawkins. (Mar. 2019). *Everything you Need to Know About the Boeing 737 Max Airplane crashes*. The Verge. Accessed: Sep. 9, 2021. [Online]. Available: <https://www.theverge.com/2019/3/22/18275736/boeing-737-max-plane-crashes-grounded-problems-info-details-explained-reasons>
- [14] P. J. Scott and R. V. Yampolskiy, "Classification schemas for artificial intelligence failures," *Delphi-Interdiscipl. Rev. Emerg. Technol.*, vol. 2, no. 4, pp. 186–199, 2019, doi: [10.21552/DELPHI/2019/4/8](https://doi.org/10.21552/DELPHI/2019/4/8).
- [15] R. V. Yampolskiy, "Predicting future AI failures from historic examples," *Foresight*, vol. 21, no. 1, pp. 138–152, Mar. 2019, doi: [10.1108/FS-04-2018-0034](https://doi.org/10.1108/FS-04-2018-0034).
- [16] (2021). *AI Fails and How We Can Learn From Them*. MITRE Partnership Network. Accessed: Sep. 9, 2021. [Online]. Available: <https://sites.mitre.org/aifails/>
- [17] S. McGregor, "Preventing repeated real world AI failures by cataloging incidents: The AI incident database," 2020, *arXiv:2011.08512*.
- [18] *House of Lords Select Committee on Artificial Intelligence. AI in the UK: Ready, Willing and Able*, UK Parliament, Apr. 2017. [Online]. Available: <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/10002.htm>
- [19] (Apr. 2019). *The European Commission Independent High-Level Expert Group on Artificial Intelligence*. Ethics Guidelines for Trustworthy AI. [Online]. Available: https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=60419
- [20] Y. D. Clarke, "H.R.2231—116th congress (2019–2020)," USA Congr. Bill, Tech. Rep., 2019. [Online]. Available: <https://www.congress.gov/bills/116/congress-house-bill/2231>
- [21] (Apr. 2021). *European Commission. Europe Fit for the Digital Age: Commission Proposes New Rules and Actions for Excellence and Trust in Artificial Intelligence*. [Online]. Available: https://ec.europa.eu/commission/presscorner/detail/en/IP_21_1682
- [22] ACM U.S. Public Policy Council (USACM) and ACM Europe Council Policy Committee (EUACM). (May 2017). *Statement on Algorithmic Transparency and Accountability*. Accessed: Jan. 14, 2019. [Online]. Available: https://www.acm.org/binaries/content/assets/publicpolicy/2017_joint_statement_algorithms.pdf

- [23] "The IEEE global initiative on ethics of autonomous and intelligent systems. Ethically aligned design: A vision for prioritizing human well-being with autonomous and intelligent systems," Tech. Rep., 2019. [Online]. Available: <https://ethicsinaction.ieee.org/wp-content/uploads/eald1e.pdf>
- [24] N. Diakopoulos, "Algorithmic accountability reporting: On the investigation of black boxes," Tow Center for Digital Journalism, Columbia Univ., New York, NY, USA, Tech. Rep., 2014, doi: [10.7916/D8ZK5TW2](https://doi.org/10.7916/D8ZK5TW2).
- [25] D. Reisman, J. Schultz, K. Crawford, and M. Whittaker, "Algorithmic impact assessments: A practical framework for public agency accountability," AI Now Inst., New York, NY, USA, Tech. Rep., Apr. 2018. [Online]. Available: <https://ainowinstitute.org/aiareport2018.pdf>
- [26] IBM Corp. (2019). *Everyday Ethics for Artificial Intelligence*. [Online]. Available: <https://www.ibm.com/watson/assets/duo/pdf/everydayethics.pdf>
- [27] M. Wieringa, "What to account for when accounting for algorithms: A systematic literature review on algorithmic accountability," in *Proc. Conf. Fairness, Accountability, Transparency*, New York, NY, USA, Jan. 2020, pp. 1–18.
- [28] S. Cranefield, N. Oren, and W. W. Vasconcelos, "Accountability for practical reasoning agents," in *Proc. Int. Conf. Agreement Technol.* Cham, Switzerland: Springer, 2018, pp. 33–48.
- [29] J. J. Bryson, M. E. Diamantis, and T. D. Grant, "Of, for, and by the people: The legal lacuna of synthetic persons," *Artif. Intell. Law*, vol. 25, no. 3, pp. 273–291, Sep. 2017, doi: [10.1007/S10506-017-9214-9](https://doi.org/10.1007/S10506-017-9214-9).
- [30] I. D. Raji, A. Smart, R. N. White, M. Mitchell, T. Gebru, B. Hutchinson, J. Smith-Loud, D. Theron, and P. Barnes, "Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing," in *Proc. Conf. Fairness, Accountability, Transparency*, New York, NY, USA, Jan. 2020, pp. 33–44, doi: [10.1145/3351095.3372873](https://doi.org/10.1145/3351095.3372873).
- [31] N. Diakopoulos, S. Friedler, M. Arenas, S. Barocas, M. L. Hay, B. Howe, H. V. Jagadish, K. Unsworth, A. Sahuguet, S. Venkatasubramanian, C. Wilson, C. Yu, and B. Zevenbergen, *Principles for Accountable Algorithms and a Social Impact Statement for Algorithms*. FAT/ML. Accessed: Jan. 10, 2019. [Online]. Available: <http://www.fatml.org/resources/principles-for-accountable-algorithms>
- [32] R. Williams, R. Cloete, J. Cobbe, C. Cottrill, P. Edwards, M. Markovic, I. Naja, F. Ryan, J. Singh, and W. Pang, "From transparency to accountability of intelligent systems: Moving beyond aspirations," *Data Policy*, vol. 4, no. e7, pp. 1–23, 2022, doi: [10.1017/DAP.2021.37](https://doi.org/10.1017/DAP.2021.37).
- [33] J. A. Kroll, S. Barocas, E. W. Felten, J. R. Reidenberg, D. G. Robinson, and H. Yu, "Accountable algorithms," *Univ. Pennsylvania Law Rev.*, vol. 165, pp. 633–705, Jan. 2017.
- [34] (2020). *Information Commissioner's Office and The Alan Turing Institute. Explaining Decisions Made With AI*. Accessed: Dec. 12, 2021. [Online]. Available: <https://ico.org.uk/for-organisations/guide-to-data-protection/key-dp-themes/explaining-decisions-made-with-artificial-intelligence>
- [35] J. Cobbe and J. Singh, "Reviewable automated decision-making," *Comput. Law Secur. Rev.*, vol. 39, Nov. 2020, Art. no. 105475, doi: [10.1016/j.clsr.2020.105475](https://doi.org/10.1016/j.clsr.2020.105475).
- [36] The European Parliament and the Council of the European Union. (Apr. 2016). *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons With Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation)*, OJ 2016 L 119/1. Accessed: Oct. 1, 2020. [Online]. Available: <https://eurlex.europa.eu/eli/reg/2016/679/oj>
- [37] J. Singh, J. Cobbe, and C. Norval, "Decision provenance: Harnessing data flow for accountable systems," *IEEE Access*, vol. 7, pp. 6562–6574, 2019, doi: [10.1109/ACCESS.2018.2887201](https://doi.org/10.1109/ACCESS.2018.2887201).
- [38] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. Daumeé, and K. Crawford, "Datasheets for datasets," in *Proc. 5th Workshop Fairness, Accountability, Transparency Mach. Learn.*, 2018, pp. 86–92.
- [39] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, and T. Gebru, "Model cards for model reporting," in *Proc. Conf. Fairness, Accountability, Transparency*, New York, NY, USA, Jan. 2019, pp. 220–229, doi: [10.1145/3287560.3287596](https://doi.org/10.1145/3287560.3287596).
- [40] M. Arnold, R. K. E. Bellamy, M. Hind, S. Houde, S. Mehta, A. Mojsilovic, R. Nair, K. N. Ramamurthy, A. Olteanu, D. Piorkowski, D. Reimer, J. Richards, J. Tsay, and K. R. Varshney, "FactSheets: Increasing trust in AI services through supplier's declarations of conformity," *IBM J. Res. Develop.*, vol. 63, no. 4/5, pp. 6:1–6:13, Jul. 2019, doi: [10.1147/JRD.2019.2942288](https://doi.org/10.1147/JRD.2019.2942288).
- [41] T. Menzies, "The five laws of SE for AI," *IEEE Softw.*, vol. 37, no. 1, pp. 81–85, Jan. 2020, doi: [10.1109/MS.2020.2967196](https://doi.org/10.1109/MS.2020.2967196).
- [42] S. Amershi, A. Begel, C. Bird, R. DeLine, H. Gall, E. Kamar, N. Nagappan, B. Nushi, and T. Zimmermann, "Software engineering for machine learning: A case study," in *Proc. IEEE/ACM 41st Int. Conf. Softw. Eng., Softw. Eng. Pract. (ICSE-SEIP)*, May 2019, pp. 291–300, doi: [10.1109/ICSE-SEIP.2019.00042](https://doi.org/10.1109/ICSE-SEIP.2019.00042).
- [43] T. Lebo, S. Sahoo, and D. McGuinness. (Apr. 2013). *PROV-O: The PROV Ontology*. W3C recommendation. [Online]. Available: <http://www.w3.org/TR/2013/REC-prov-o-20130430/>
- [44] M. Markovic, D. Garijo, P. Edwards, and W. Vasconcelos, "Semantic modelling of plans and execution traces for enhancing transparency of IoT systems," in *Proc. 6th Int. Conf. Internet Things, Syst., Manage. Secur. (IOTSMS)*, Oct. 2019, pp. 110–115, doi: [10.1109/IOTSMS48152.2019.8939260](https://doi.org/10.1109/IOTSMS48152.2019.8939260).
- [45] I. Naja, M. Markovic, P. Edwards, and C. Cottrill, "A semantic framework to support AI system accountability and audit," in *The Semantic Web*, R. Verborgh, K. Hose, P.-A. Champin, M. Maleshkova, O. Corcho, P. Ristoski, and M. Alam, Eds. Cham, Switzerland: Springer, 2021, pp. 160–176, doi: [10.1007/978-3-030-77385-4_10](https://doi.org/10.1007/978-3-030-77385-4_10).
- [46] M. Markovic, I. Naja, P. Edwards, and W. Pang, "The accountability fabric: A suite of semantic tools for managing AI system accountability and audit," in *Proc. CEUR Workshop*, Oct. 2021, pp. 1–5.
- [47] M. Uschold and M. Gruninger, "Ontologies: Principles, methods and applications," *Knowl. Eng. Rev.*, vol. 11, no. 2, pp. 93–136, Jun. 1996.
- [48] Department of Health & Social Care (UK). (Jul. 2019). *Code of Conduct for Data-Driven Health and Care Technology*. Accessed: Sep. 10, 2019. [Online]. Available: <https://www.gov.uk/government/publications/code-of-conduct-for-data-driven-health-and-care-technology/initial-code-of-conduct-for-data-driven-health-and-care-technology>
- [49] Digital Catapult Machine Intelligence Garage Ethics Committee. (Apr. 2020). *Ethics Framework*. Accessed: Jan. 14, 2022. [Online]. Available: https://migarage.digicatapult.org.uk/wpcontent/uploads/2020/11/MIG_Ethics-Report_2020_v5.pdf
- [50] V. Curcin, E. Fairweather, R. Danger, and D. Corrigan, "Templates as a method for implementing data provenance in decision support systems," *J. Biomed. Inform.*, vol. 65, pp. 1–21, Jan. 2017, doi: [10.1016/j.jbi.2016.10.022](https://doi.org/10.1016/j.jbi.2016.10.022).
- [51] T. D. Huynh, S. Stalla-Bourdillon, and L. Moreau, "Provenance-based explanations for automated decisions: Final IAA project report," King's College London, Tech. Rep., 2019. [Online]. Available: [https://kclpure.kcl.ac.uk/portal/en/publications/provenancebased-explanations-for-automated-decisions\(5b1426ce-d253-49fa-8390-4bb3abe65f54\).html](https://kclpure.kcl.ac.uk/portal/en/publications/provenancebased-explanations-for-automated-decisions(5b1426ce-d253-49fa-8390-4bb3abe65f54).html)
- [52] G. C. Publio, D. Esteves, A. Ławrynowicz, P. Panov, L. Soldatova, T. Soru, J. Vanschoren, and H. Zafar, "ML-schema: Exposing the semantics of machine learning with schemas and ontologies," in *Proc. 2nd Reproducibility Mach. Learn. Workshop ICML*, Stockholm, Sweden, Jul. 2018, pp. 1–5.
- [53] Y. Gil, V. Ratnakar, and D. Garijo, "OntoSoft: Capturing scientific software metadata," in *Proc. 8th Int. Conf. Knowl. Capture*, New York, NY, USA, Oct. 2015, pp. 1–4, doi: [10.1145/2815833.2816955](https://doi.org/10.1145/2815833.2816955).
- [54] M. Hind, S. Houde, J. Martino, A. Mojsilovic, D. Piorkowski, J. Richards, and K. R. Varshney, "Experiences with improving the transparency of AI models and services," in *Proc. Extended Abstr. CHI Conf. Hum. Factors Comput. Syst.*, New York, NY, USA, Apr. 2020, pp. 1–8, doi: [10.1145/3334480.3383051](https://doi.org/10.1145/3334480.3383051).



IMAN NAJA is a Research Fellow in computing science with the University of Aberdeen. Her research interests include accountability, audit, trustworthiness, and transparency of AI systems, and her work has been in applying Semantic Web technologies and provenance models to address those topics.



WEI PANG is an Associate Professor in computer science with Heriot-Watt University. His research interests include machine learning, bio-inspired computing, and accountable AI. He has authored over 130 papers in the above fields.



MILAN MARKOVIC is a Research Fellow in computing science with the University of Aberdeen. His research interests include enhancing transparency and accountability in complex socio-technical systems through intelligent processes based on provenance data models and Semantic Web technologies.



CAITLIN COTTRILL was born in December 1977. She was a Senior Lecturer and the Director of the Centre for Transport Research with the School of Engineering, University of Aberdeen. Her research explored the use of technology in urban spaces, with particular focus on transport and data privacy, and her work had a strong focus in the support of public and active transport modes, especially for vulnerable and under-served users.



PETER EDWARDS is a Professor of computing science with the University of Aberdeen. His research interests include provenance, trust, and information quality, with a particular emphasis on the role of Semantic Web technologies in helping realize transparent and accountable systems.



REBECCA WILLIAMS received the undergraduate degree from Worcester College, Oxford, and the Ph.D. degree from the University of Birmingham. She is a Professor of public law and criminal law, in association with the Pembroke College. She was previously a fellow of the Robinson College, Cambridge. Her principal teaching interests include criminal law and public law. Her research interests include law and computer science, criminal law, public law, and the interrelationship between public law and unjust enrichment.

...