

AIUP: an ODRL Profile for Expressing AI Use Policies to Support the EU AI Act

Delaram Golpayegani^{1,*}, Beatriz Esteves², Harshvardhan J. Pandit³ and Dave Lewis¹

¹ADAPT Centre, Trinity College Dublin, Dublin, Ireland

²IDLab, Ghent University – imec, Ghent, Belgium

³ADAPT Centre, Dublin City University, Dublin, Ireland

Abstract

The upcoming EU AI Act requires providers of high-risk AI systems to define and communicate the system’s intended purpose – a key and complex concept upon which many of the Act’s obligations rely. To assist with expressing the intended purposes and uses, along with precluded uses as regulated by the AI Act, we extend the Open Digital Rights Language (ODRL) with a profile to express the AI Use Policy (AIUP). This open approach to declaring use policies enables explicit and transparent expression of the conditions under which an AI system can be used, benefiting AI application markets beyond the immediate needs of high-risk AI compliance in the EU. AIUP is available online at <https://w3id.org/aiup> under the CC-BY-4.0 license.

Keywords

AI Act, ODRL, AI use policy, AI risk management, regulatory enforcement, trustworthy AI

1. Introduction

Within the EU AI Act [1] there is a strong emphasis on intended purpose – a legal term-of-art described as the use of the system specified by the provider, which should include information regarding context and conditions of use (AI Act, Art. 3). Given its importance in assessment of risk level under the Act [2], and in turn in ensuring safe and trustworthy use of AI, intended purpose of an AI system should be communicated to its deployers in a transparent manner. In this paper, we aim to simplify the specification of this key concept by adopting a policy-based approach. As such, we propose to extend the W3C’s recommendation on Open Digital Rights Language (ODRL)¹ to fulfil the representation of intended purpose through an AI Use Policy (AIUP) profile. AIUP serves as a mechanism for expressing AI intended and precluded uses as well as conditions of use by modelling them as permissions, prohibitions, and duties within a policy.


SEMANTiCS’24: 20th International Conference on Semantic Systems, September 17–19, 2024, Amsterdam, Netherlands

*Corresponding author.

✉ sgolpays@tcd.ie (D. Golpayegani); beatriz.esteves@ugent.be (B. Esteves); me@harshp.com (H. J. Pandit); delewis@tcd.ie (D. Lewis)

🆔 0000-0002-1208-186X (D. Golpayegani); 0000-0003-0259-7560 (B. Esteves); 0000-0002-5068-3714 (H. J. Pandit); 0000-0002-3503-4644 (D. Lewis)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 ¹<https://www.w3.org/ns/odrl/2/>

2. Related Work

ODRL has been leveraged for legal compliance and policy enforcement, particularly in EU GDPR compliance tasks such as automated checking of consent permissions [3], expressing legal obligations [4], and modelling the obligations in terms of permissions and prohibitions regarding executing business processes [5]. In the context of data governance, ODRL was extended for expressing policies related to access control over data stored in Solid Pods [3], utilised for modelling policies associated with responsible use of genomics data [6], and used in expressing data spaces' usage and access control policies [7, 8].

3. AIUP

3.1. AIUP Requirements

AIUP is intended to be used by AI providers and deployers to communicate and negotiate the conditions under which an AI system can/cannot be used. The competency questions, which shape the requirements of the policy profile, are extracted from the AI Act and listed in the following:

- CQ1. What is the intended use(s) of the AI system? (Art. 13 and Annex IV(1a))
- CQ2. What is the precluded use(s)² of the AI system? (Recital 72)
- CQ3. To use the system as intended, what human oversight measure(s) should be implemented by the deployer? (Art. 14 (3)(b))
- CQ4. What is the reporting obligation(s) of the deployer? (Art. 26(5))

To express intended and precluded uses, we utilise the 5 concepts identified in our previous work [9] that are domain, purpose, AI capability, AI deployer, and AI subject. To further capture the context of use, we also include locality of use.

3.2. AIUP Overview

An overview of the AIUP's profile is illustrated in Figure 1. Expressing intended and precluded uses of an AI system or component within a policy are enabled by employing `odr1:permission` and `odr1:prohibition` rules respectively. For expressing the conditions of use, i.e., obligations that should be fulfilled by a party in order to use a system or component, the `odr1:duty` property should be employed. The vocabulary used in AIUP is defined in alignment with the AI Risk Ontology (AIRO) [10] and the Data Privacy Vocabulary (DPV) [11]. The development follows the ODRL V2.2 Profile Best Practices³, which requires the terms to be defined in the policy namespace (in this case `aiup`) with `skos:exactMatch` to link the proposed terms to existing vocabularies.

AIUP introduces 3 types of `aiup:UsePolicy`, that are `aiup:UseOffer`, `aiup:UseRequest`, and `aiup:UseAgreement`. These enable expressing offers, requests, and agreements from/between AI providers and deployers. To address the

²Refers to the uses of an AI system that are prohibited by the provider.

³<https://w3c.github.io/odrl/profile-bp/>

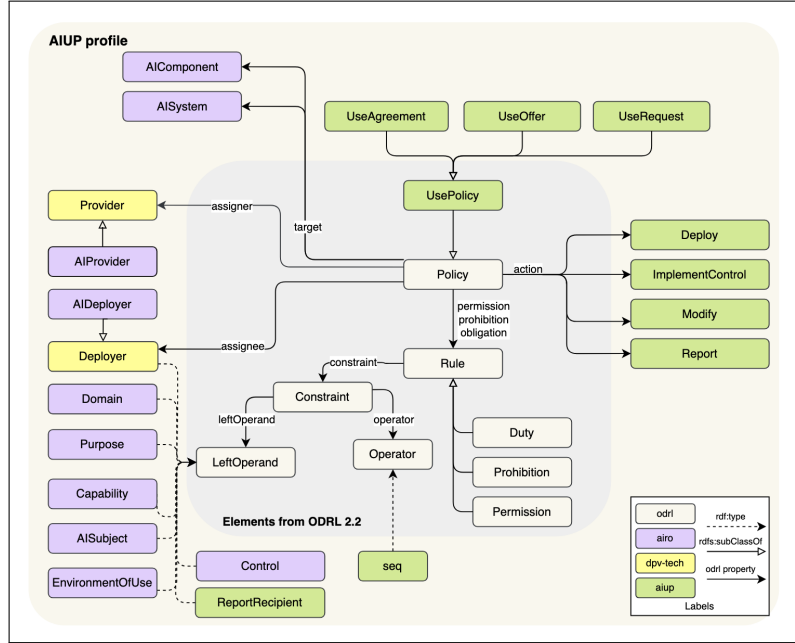


Figure 1: AIUP core classes and properties.

ambiguities around the function of `odr1:isA` in the inclusion of “sub-class of” relations, we introduce semantic equality (`aiup:seq`) that indicates presence of either “instance of” or “sub-class of” relations. AIUP allows describing use policies for AI components, such as general-purpose AI models, by specifying general concepts of `aiup:AIComponent`, `aiup:Provider`, and `aiup:Deployer`. However, it leaves out the inclusion of more specific elements required for expressing component use policies for future work. AIUP is made available online at <https://w3id.org/aiup> under the CC-BY-4.0 license.

3.3. AIUP Example

As an example scenario, we consider a policy for an online student proctoring system called Proctify, previously described in [12]. The conditions of deploying Proctify, as an `aiup:UseOffer` policy, are presented in Listing 1. For brevity, we only include 3 constraints for describing the intended domain, purpose, and AI subjects. The offer indicates that the deployer should provide training to end-users of the system as a control measure to address the risk of over-reliance on the system’s output.

4. Conclusion

In this paper, we proposed AIUP as a novel technical solution for declaring AI use policies in an open, machine-readable, and interoperable format based on the evolving requirements of the AI value chain, particularly the obligations of the EU AI Act. The AIUP profile supports modelling and comparison of use policies related to AI systems

```

1 @prefix odr1: <https://www.w3.org/ns/odrl/2/> .
2 @prefix aiup: <https://w3id.org/aiup#> .
3 @prefix vair: <http://w3id.org/vair#> .
4 @prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
5 @prefix dct: <http://purl.org/dc/terms/> .
6 @prefix ex: <http://example.org/> .
7
8 ex:proctify-offer-01 a aiup:UseOffer ;
9   odr1:uid ex:proctify-offer-01 ;
10  odr1:profile aiup: ;
11  rdfs:comment "Offer for using Proctify"@en ;
12  odr1:permission [
13    a odr1:Permission ;
14    odr1:assigner ex:aiedux ;
15    odr1:target ex:proctify ;
16    odr1:action aiup:Deploy ;
17    odr1:constraint [
18      odr1:and [
19        odr1:leftOperand aiup:Domain ;
20        odr1:operator aiup:seq ;
21        odr1:rightOperand vair:Education ] ,
22      [
23        odr1:leftOperand aiup:Purpose ;
24        odr1:operator aiup:seq ;
25        odr1:rightOperand vair:DetectCheating ] ,
26      [
27        odr1:leftOperand aiup:AISubject ;
28        odr1:operator aiup:seq ;
29        odr1:rightOperand vair:Student ] ] ] ;
30  odr1:duty [
31    dct:title "User training to address over-reliance" ;
32    odr1:action aiup:ImplementControl ;
33    odr1:constraint [
34      odr1:leftOperand aiup:Control ;
35      odr1:operator aiup:seq ;
36      odr1:rightOperand vair:Training ] ] ] .

```

Listing 1: An example of aiup:UseOffer describing Proctify’s use policy.

and their components. It further assists AI auditors and authorities in investigation of non-compliance and ascertaining liable parties when investigating claims concerning AI.

Acknowledgments

This project has received funding from the EUs Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No 813497 (PROTECT ITN) and from Science Foundation Ireland under Grant#13/RC/2106_P2 at the ADAPT SFI Research Centre. Beatriz Esteves is funded by SolidLab Vlaanderen (Flemish Government, EWI and RRF project VV023/10). Harshvardhan Pandit has received

funding under the SFI EMPOWER program.

References

- [1] Regulation (EU) 2024/1689 of the european parliament and of the council of 13 june 2024 laying down harmonised rules on artificial intelligence and amending regulations (EC) no 300/2008, (EU) no 167/2013, (EU) no 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (artificial intelligence act), 2024. URL: <http://data.europa.eu/eli/reg/2024/1689/oj>.
- [2] I. Hupont, D. Fernández-Llorca, S. Baldassarri, E. Gómez, Use case cards: A use case reporting framework inspired by the european AI act, *Ethics and Information Technology* 26 (2024).
- [3] B. Esteves, H. J. Pandit, V. Rodríguez-Doncel, ODRL profile for expressing consent through granular access control policies in solid, in: 2021 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW), IEEE, 2021, pp. 298–306.
- [4] S. Agarwal, S. Steyskal, F. Antunovic, S. Kirrane, Legislative compliance assessment: Framework, model and GDPR instantiation, in: *Privacy Technologies and Policy*, Springer International Publishing, Cham, 2018, pp. 131–149.
- [5] M. De Vos, S. Kirrane, J. Padget, K. Satoh, ODRL policy modelling and compliance checking, in: P. Fodor, M. Montali, D. Calvanese, D. Roman (Eds.), *Rules and Reasoning*, SpringerInternational Publishing, 2019, pp. 36–51.
- [6] H. J. Pandit, B. Esteves, Enhancing data use ontology (DUO) for health-data sharing by extending it with ODRL and DPV, *Semantic Web Journal* (2024).
- [7] T. Dam, A. Krimbacher, S. Neumaier, Policy patterns for usage control in data spaces, *arXiv preprint arXiv:2309.11289* (2023).
- [8] I. Akaichi, W. Slabbinck, J. A. Rojas, C. Van Gheluwe, G. Bozzi, P. Colpaert, R. Verborgh, S. Kirrane, Interoperable and continuous usage control enforcement in dataspace, in: *The Second International Workshop on Semantics in Dataspace*, co-located with the Extended Semantic Web Conference, 2024.
- [9] D. Golpayegani, H. J. Pandit, D. Lewis, To be high-risk, or not to be—semantic specifications and implications of the AI act’s high-risk AI applications and harmonised standards, in: *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 2023, pp. 905–915.
- [10] D. Golpayegani, H. J. Pandit, D. Lewis, AIRO: An ontology for representing AI risks based on the proposed EU AI Act and ISO risk management standards, in: *Towards a Knowledge-Aware AI*, volume 55, IOS Press, 2022, pp. 51–65.
- [11] H. J. Pandit, B. Esteves, G. P. Krog, P. Ryan, D. Golpayegani, J. Flake, Data privacy vocabulary (DPV)—version 2, *arXiv preprint arXiv:2404.13426* (2024).
- [12] D. Golpayegani, I. Hupont, C. Panigutti, H. J. Pandit, S. Schade, D. O’Sullivan, D. Lewis, AI cards: Towards an applied framework for machine-readable AI and risk documentation inspired by the EU AI act, in: *Privacy Technologies and Policy*, Springer Nature Switzerland, 2024, pp. 48–72.