



AI product cards: a framework for code-bound formal documentation cards in the public administration

Albana Celepija^{1,2} , Alessio Palmero Aprosio¹ , Bruno Lepri¹ and Raman Kazhamiakin¹

¹Fondazione Bruno Kessler, Digital Society Center, Trento, Italy

²Department of Information Engineering and Computer Science, University of Trento, Trento, Italy

Corresponding author: Albana Celepija; Email: albana.celepija@unitn.it

Received: 01 December 2023; **Revised:** 17 May 2024; **Accepted:** 01 June 2024

Keywords: adaptable AI; AI documentation; AI product cards; public administration; reusable AI solutions

Abstract

Currently, artificial intelligence (AI) is integrated across various segments of the public sector, in a scattered and fragmented manner, aiming to enhance the quality of people's lives. While AI adoption has proven to have a great impact, there are several aspects that hamper its utilization in public administration. Therefore, a large set of initiatives is designed to play a pivotal role in promoting the adoption of reliable AI, including documentation as a key driver. The AI community has been proactively recommending a variety of initiatives aimed at promoting the adoption of documentation practices. While currently proposed AI documentation artifacts play a crucial role in increasing the transparency and accountability of various facts about AI systems, we propose a code-bound declarative documentation framework that aims to support the responsible deployment of AI-based solutions. Our proposed framework aims to address the need to shift the focus from data and models being considered in isolation to the reuse of AI solutions as a whole. By introducing a formalized approach to describing adaptation and optimization techniques, we aim to enhance existing documentation alternatives. Furthermore, its utilization in the public administration aims to foster the rapid adoption of AI-based applications due to the open access to common use cases in the public sector. We further showcase our proposal with a public sector-specific use case, such as a legal text classification task, and demonstrate how the AI Product Card enables its reuse through the interactions of the formal documentation specifications with the modular code references.

Policy Significance Statement

The framework proposes an approach that allows the acceleration of adopting AI solutions in the public sector by offering a set of techniques that facilitate the adaptation and reuse of AI-based use cases. The benefits of adopting the framework are twofold. First, the policymakers are able to disseminate the use cases in the public sector that have been solved through the use of AI, so that they can be reused by additional public entities. Second, the documentation approach of different components that constitute the whole solution is bound with the specifications and the relevant code that enables its responsible reuse. Therefore, the documentation initiatives that support the regulatory objectives of building AI systems should be translated into actionable operations that enhance public policy processes.



1. Introduction

With the increasing ubiquity of digital technologies, companies and organizations constantly restructure their processes to generate and capture value. This is considered a transformative routine that affects fundamental changes by leveraging existing competencies or developing innovative ones (Liu et al., 2011; Alvarenga et al., 2020). Innovation encompasses various forms such as novel practices, ideas, service delivery methods, or technologies that generate value including Artificial Intelligence (AI). When it comes to public innovation, it specifically pertains to innovation within the context of public administration (van Noordt & Misuraca, 2022). The benefits of public innovation result in significant impacts such as increased citizen engagement, enhanced decision-making processes, and broader influence on businesses (Bertot et al., 2016).

Currently, in the context of public administration, there is increased evidence of the utilization of AI tasks (Mehr, 2017) using techniques, such as, Machine Learning (ML), Computer Vision (CV), Natural Language Processing (NLP) and Automatic Speech Recognition (ASR). Similar techniques that address the problems of digital transformation may combine diverse models or tasks that interact with each other as part of a system that solves a larger problem (Martie et al., 2023; Kamal Paul & Bhaumik, 2022).

Prior studies emphasize that the utilization of AI technologies within public administration has the potential to enhance policy-making processes, improve the delivery of public services, and optimize the internal management of public administrations (Madan & Ashok, 2023). Furthermore, they demonstrate the current utilization of AI technologies for various objectives, such as virtual agents, recommendation systems, cognitive robotics and autonomous systems, and AI-driven process automation systems (Mikalef et al., 2019). However, limited empirical evidence exists regarding their impact due to a large number of barriers to the implementation and evaluation of AI (Bérubé et al., 2021).

Previous research findings into aspects that hamper AI adoption elicit concern about the lack of a unified integrated approach to deploying AI solutions in the public sector. Given the significant lag compared to the private sector, it is crucial to ensure the effective and consistent progression of AI in public organizations. Factors thought to be influencing AI adoption in the public sector have been explored in several studies (Mikalef et al., 2022). For (Mikalef et al., 2019) the main hindering aspects to adopting AI technologies are of four kinds: (1) lack of expertise, (2) organizational implementation challenges, (3) inability to integrate systems and data, and (4) data availability and privacy concerns. On the other hand, (John et al., 2021) and (Zhang et al., 2023) argue that there are architecture decisions, such as technology, component and data decisions, that affect the success of AI-based systems in terms of deployment and maintenance. In relation to insufficient expertise, AI technologies are not fully embraced due to the lack of trained staff, AI engineers and practitioners, which requires dedicated specialization and deep expertise in building and maintaining AI tasks. In addition, organizational challenges in implementing AI solutions often stem from a combination of uncertainty about where to begin and resistance to changing existing work practices (Medaglia et al., 2023; Desouza et al., 2020).

Consequently, despite the groundbreaking potential of AI, its implementation in the public sector is struggling due to the fact that it currently considers AI applications in isolation and fragmentation (Wirtz et al., 2019). This is mainly because public authorities may not be aware of the full range of AI application opportunities and do not have a governance scheme to guide the AI adoption process in different domains. Moreover, the complexity of developing AI-based systems increases as a result of the absence of a recognizable pattern, particularly when transferring or adapting the same AI task across different domains or settings in public administration.

The challenge of transferring AI products across different domains arises mainly due to the varying statistical characteristics of production data compared to the data encountered during training. Traditional learning algorithms often exhibit limitations when reused in dissimilar conditions due to data distribution shifts, commonly known as domain shifts (Sun & Saenko, 2016; Ganin & Lempitsky, 2015; Cobb & Looveren, 2022).

Although the potential benefits of deploying AI-based systems are enticing, effectively assessing their successful deployment and utilization in different contexts remains an unresolved challenge (Wirtz et al., 2019).

Given the potential benefits of leveraging existing solutions and best practices, it is not necessary for every local government to independently develop an AI-based solution (Madan & Ashok, 2023). Instead, they can draw upon the experiences and successes of other jurisdictions that have faced similar challenges. By doing so, they can avoid reinventing the wheel and expedite the adoption of AI technologies in their own organizations. To achieve this objective, classifying, categorizing, and documenting AI solutions employed in public administration could facilitate the sharing of common scenarios among various administrative bodies. The sharing of software in the public sector has been implemented in different countries¹ with the intention of transforming the way public administrations work. For example, the E-Government Action Plan 2016–2020², Sharing and Reuse of IT solutions³, and the European Interoperability Framework⁴ initiatives promote the reuse of information and solutions in member countries in order to reduce costs, facilitate interoperability and develop services more quickly and transparently by using existing components.

Several studies have highlighted the importance of documentation in achieving transparency and accountability (Königstorfer & Thalmann, 2022) in AI systems. The AI community has put a substantial amount of effort towards standardised approaches for documenting models (Mitchell et al., 2019), data (Gebu et al., 2021; Pushkarna et al., 2022; Holland et al., 2018), use cases (Hupont et al., 2023), systems⁵ as well as their related potential risks (Derczynski et al., 2023).

However, existing studies have failed to address the need for using the documentation of AI systems as a means to the reuse and rapid deployment of AI products in different contexts. There is still little research on documentation that considers the AI product as a whole (Micheli et al., 2023) in order to ensure its responsible deployment and offer a basis for evaluating adherence to regulatory standards. Despite efforts to create documentation protocols that increase transparency in the development of various AI system components, there is currently a gap in documenting how to reuse the entire product in compliance with regulatory requirements. The main aim of this manuscript is therefore to address the following research question: How can documentation be utilized as a means to facilitate the rapid deployment of AI-driven systems in contexts that allow their effective and responsible reuse?

In this paper, we aim to provide a framework for reusable AI solutions suitable for the public administration sector. For this purpose, our approach builds upon preliminary studies on AI documentation initiatives in order to provide a comprehensive view of a complete AI product that is able to be adapted and reused in different contexts.

Furthermore, our proposal caters to the needs of policymakers and AI practitioners who would be able to systematize and classify AI applications for public sector-specific use cases. Overall, the main contributions of our work are as follows:

- Given that typical activities in the context of public administration exhibit substantial uniformity or similarity we propose standardised documentation protocols that facilitate the reuse and adaptation of AI solutions.
- An AI system is made up of different components, including models and data, that are utilized to solve a problem. Therefore, merely disclosing the documentation for each component's individual functioning is insufficient for leveraging them within a system that integrates them. We propose to formalize a documentation procedure that offers a holistic perspective on the elements constituting the final AI product, as opposed to focusing solely on individual and fragmented pieces of documentation.

¹ <https://digital-strategy.ec.europa.eu/en/library/egovernment-benchmark-2023>

² <https://digital-strategy.ec.europa.eu/en/policies/egovernment-action-plan>

³ https://joinup.ec.europa.eu/sites/default/files/custom-page/attachment/2017-10/sharing_and_reuse_of_it_solutions_framework_final.pdf

⁴ <https://joinup.ec.europa.eu/collection/nifo-national-interoperability-framework-observatory/2-underlying-principles-european-public-services>

⁵ <https://ai.meta.com/research/publications/system-level-transparency-of-machine-learning/>

- We aim to align the documentation with functional code that enables the automation of AI product design, reuse, and deployment in order to use the documentation as a tool to build AI-based systems.

The remainder of this paper is structured in the following manner. We offer background information that our work relies on in [Section 2](#). In [Section 3](#), we detail the components that make up the AI Product Card and provide a sample to demonstrate how it should be completed. Lastly, in [Section 4](#), we draw conclusions and identify some limitations of our proposed framework.

2. Related work

In this section, we discuss the state-of-the-art AI systems documentation approaches from the perspective of the need for reuse and adaptation of typical public sector AI use cases.

A number of methodologies for documentation practices have emerged in recent times. Their primary objective is related to the enhancement of transparency, quality or discoverability in the design, deployment and evaluation of trustworthy AI solutions. The initial wave of documentation initiatives is dedicated to more horizontal and general-purpose disclosure of the information regarding the AI models and the characteristics of the data used during the training and validation phases. In light of the growing demand for transparency and accountability in the creation of AI-based applications, the concept of a Model Card has been proposed as a response (Mitchell et al., 2019). Model Card takes a descriptive approach to provide detailed information regarding the way the model was built, its architectural type, the intended purpose, and the performance metrics that influence its behavior. Additionally, it may provide general descriptions of the datasets employed for both training and testing the model. In this regard, to offer more comprehensive information about the data employed in constructing AI systems, various specialized data documentation initiatives have been suggested such as DataSheets for datasets (Gebru et al., 2021), Data Cards (Pushkarna et al., 2022), and the dataset nutrition labels (Holland et al., 2018).

While models and data cards adopt a broad focus, the recent trends in documenting AI systems tend to narrow down the level of the details provided for its elements. Customization and adaptation of documentation approaches are becoming more prevalent. They tend to tailor and adjust the documentation approaches based on specific domain requirements, data categories or risk-oriented regulations. To address the needs of specific domains, documentation protocols for AI systems have been proposed in healthcare (Rostamzadeh et al., 2022) and educational (Chaudhry et al., 2022) sectors. On the other hand, to accommodate diverse data categories, initiatives to document particular images (Prabhu & Birhane, 2020) and text (McMillan-Major et al., 2021) datasets have been implemented. Lastly, Risk Cards (Derczynski et al., 2023) have been proposed to create a common knowledge base for understanding and mitigating risks associated with language models.

However, although these initiatives contribute to the further adoption of documentation practices over a wide range, they can not provide a global perspective on how each component contributes to the overall system. While model and data cards place emphasis on the single component description, the need for a different perspective that focuses on the final system or service provided has been addressed by recent studies (Königstorfer & Thalmann, 2022). To offer a thorough overview of an AI service, FactSheets (Arnold et al., 2019) have been suggested as a standardized method for service creators to convey how they have tackled various concerns in a deployed version (Hind et al., 2019). Alternatively, the proposition to document the entire pipeline (Tagliabue et al., 2021) of the process of designing, developing and deploying the AI solution has gained interest, especially from an industrial point of view. Although more recent studies state the need for a fully dynamic documentation protocol for tracking AI system's behavior over time (Mehta et al., 2023), they consider models and data cards as narrowly focused and unable to provide actionable insights over a composed AI system. From this perspective, work still needs to be done to consider documentation as an active tool for constructing AI-based solutions.

On the other hand, recent advancements in AI systems documentation call for the involvement of a larger set of stakeholders in a participatory manner in order to respond to the various needs they may have toward complex AI systems. Recent initiatives include Use Case Cards (Hupont et al., 2023) and System

Table 1. Categorisation of state-of-the-art AI/data/system documentation approaches. The black circle indicates the compliance of the documentation approach with the relevant perspective. The white circle indicates a lack of adherence to the relevant perspective of the documentation approach

Card name	Fragmented perspective	System perspective	Adaptation perspective
Model card (Mitchell et al., 2019)	●	○	Transparency, fairness
DataSheets for datasets, data cards, The dataset nutrition label (Gebru et al., 2021; Pushkarna et al., 2022; Holland et al., 2018)	●	○	Ethics, reproducibility
DAG card (Tagliabue et al., 2021)	○	●	Reproducibility, quality
AI FactSheets (Arnold et al., 2019)	○	●	Trust, accountability
Method card (Adkins et al., 2022)	●	○	Reproducibility
System cards ⁸	○	●	Transparency, explainability
Risk cards (Derczynski et al., 2023)	●	○	Ethics
Reward reports (Mehta et al., 2023)	○	●	AI audits, accountability
Use case cards (Hupont et al., 2023)	●	○	Transparency, risk assessment
AI Product Card	●	●	Adaptation, reuse

Cards⁶ which aims to provide an overview of the ways different components of the system interact with each other or how the system interfaces with the end users. Regarding the impact this interaction may have on citizens, a series of tools and checklists have been proposed for assessing and mitigating biases and ethical issues in AI systems in the public sector⁷. These proposals have a positive impact in terms of the increased transparency and explainability they provide to a broad audience. However, they lack the ability to provide the designers, engineers or policymakers with the right approach and directives that enable the reuse of the AI solution.

Current documentation methods generally adopt either a descriptive approach (Mitchell et al., 2019), which seeks to describe an existing system or a prescriptive approach (Adkins et al., 2022), which focuses on creating or replicating the system under identical conditions. Nevertheless, these approaches often neglect the need for adapting and reusing solutions to different target data and task characteristics.

Table 1 summarizes the works discussed so far. For each approach, we provide the following information:

- **Card Name:** the name of the documentation type proposed by the paper's authors.
- **Fragmented perspective:** indicates that the documentation methodology focuses on a single component of the AI product, such as the model or the dataset.
- **System perspective:** indicates that the documentation approach is concerned with aspects that consider the AI system as a whole.
- **Adaptation perspective:** specifies whether the documentation approach is designed to meet the system's adaptation requirements or specifies its intention otherwise.

⁶ <https://ai.meta.com/research/publications/system-level-transparency-of-machine-learning/>

⁷ Algorithmic Transparency Recording Standard, Ethics and Algorithm Toolkit, Aequitas, Algorithmic Impact Assessment

⁸ <https://ai.meta.com/research/publications/system-level-transparency-of-machine-learning/>

In sum, the existing body of knowledge related to the AI documentation approaches lacks to provide guidance regarding the adaptation techniques in place for reusing an AI-based system, highly advantageous characteristics for use cases specific to the public sector and beyond. On the other hand, there is a need to integrate and formalize the proposed methodologies into a unified machine-readable representation of the final reproducible solution in similar use cases. Only a few documentation initiatives incorporate partial automation, focusing solely on automating the documentation generation process⁹ (Tagliabue et al., 2021), while overlooking the automation of deploying the components that the documentation aims to cover. In accordance with insights from recent research (Micheli et al., 2023), the majority of documentation methodologies center mostly on data and AI models, neglecting crucial aspects associated with the overall AI system.

This paper proposes a new documentation framework, called AI Product Card, dedicated to mitigating the shortcomings of current documentation practices in relation to the reuse and accelerated adoption of AI-based solutions. Furthermore, it aims to provide a new method of documentation focusing on both aspects of a system's proper functioning—the descriptive side, as well as the machine-readable format that aligns the descriptions with the corresponding executable scripts or artefacts.

3. AI Product Cards

In this manuscript, we introduce AI Product Cards as a descriptive and declarative documentation approach tightly coupled with the necessary code that drives the reuse and adaptation of AI products across diverse contexts. The novelty of our proposed framework lies in its emphasis on providing a comprehensive and holistic approach to the development of AI products, taking into account the entire system as a whole. The framework addresses the fragmented nature of available documentation approaches (Micheli et al., 2023) and advances them in order to automate the deployment and verification of the AI product's transferability.

An AI product undergoes various phases throughout its lifecycle (Li et al., 2023), including data validation and data processing, model development, evaluation, and monitoring. According to the challenges mentioned in the Introduction Section, from a public sector perspective, this process may require time and effort that can be reduced by adopting a shared framework that discloses AI product documentation and implementation references for each of these phases in an integrated manner. To achieve this objective, we introduce a structure for the AI Product Card as depicted in Figure 1.

First of all, the conceptualization of the final AI product starts by framing the problem explicitly, focusing on the possibility of reusing existing solutions in place. The card's goal is to assist policymakers in exploring past instances of similar issues encountered by other public entities or proposing a new problem definition, enhancing its potential for future reuse. The proposed framework is designed to function as a repository of potential use cases in public administration that are addressed by existing AI products produced by other public sector organizations. It is important to note that the AI Product Card is a comprehensive system including various elements that make up the AI product. The Domain Problem Card, which is responsible for determining the context in which the AI product will be used, serves as a single instance in the future catalog of problem definitions within the public sector. This means that once the framework is integrated into a platform, policymakers, acting as AI product consumers, will have the ability to search for and identify relevant use cases that they can repurpose for their specific scenario. This will enable them to take advantage of already implemented solutions for their problem. The attributes within the Domain Problem Card, including name, description, and tag, make the AI product easily discoverable, accessible, and reusable. Consider, for example, a scenario in which a specific municipal administration is implementing a solution that utilizes computer vision to determine whether drivers are holding their mobile phones. It would be beneficial if this problem were defined in a unique manner including relevant tags, its input data, the category of the AI task it employs, a detailed description, and the

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




Figure 1. Schema of the AI Product Card structure.

context in which it has been used. In future scenarios, another governmental entity would find and reuse the solution to the same problem that has been faced in similar circumstances.

Second, one of the primary concerns that must be addressed by the AI Product Card is the validation of the target data to which the model will be applied, in terms of quantity and characteristics they may have. Based on the outcome of the data validation, the stakeholders are able to decide whether to reproduce the same solution or to adapt their model according to the proprietary data. As a result, they are able to validate the transferability level of the found solution.

Third, due to the continually evolving nature of real-world data and the diverse characteristics of various governmental entities, it is necessary for the AI product to be adapted to suit the aspects of production data. This is achieved through the dedicated section for the adaptation procedure which may encompass fine-tuning techniques that have to be disclosed through the completion of the Adaptation specification section. This entails providing the necessary fine-tuning script that specializes the model according to the production data. On the other hand, the specifications should cover also aspects that drive the optimization phase of the solution so as to improve the time and resources employed during the training and inference time. These stages of the AI Product should be taken into account when populating the documentation and specification sections of the framework. The policymakers need the tools that facilitate the implementation of adaptation strategies, thereby guiding the model towards an optimal state or specializing it in order to conform to the characteristics of the available data.

The sections that compose the AI Product Card are the metadata representation of the documentation of each component of the solution, putting the main focus on the transformation of the documentation into a tool that drives the design and evaluation of the AI system. Given the diverse range of stakeholders utilizing AI systems, their transparency requirements exhibit various perspectives, necessitating a thorough and cohesive alignment in documenting each component of the system. According to the literature (Mehta et al., 2023), for the framework to achieve success, it must strike a balance between being comprehensive enough to encompass the majority of use cases for public sector-specific AI solutions and being adaptable and extensible to accommodate specific customizations.

To date, there is no unified AI documentation approach focusing on metadata specification that allows for seamless integration and deployment of the components of an AI-based system. To demonstrate our approach, we present metadata definitions of the Cards and show how these formal definition paves the way towards the implementation and deployment of the AI Solution.

3.1. *An AI Product Card template*

In accordance with the suggested descriptive and prescriptive documentation methodologies, we identify certain components as indispensable for the formation of a final AI product that can subsequently be reused in a comparable context. In this section, we describe the components of the AI Product Card, a novel documentation approach that presents a comprehensive overview of an AI system by incorporating documentation techniques for each of its constituents and expressing it as a set of metadata specifications. The paper proposes a way to formalize the stages of the lifecycle of an AI product to promote its reuse across various contexts. By establishing standardized procedures and necessary components, the approach aims to accelerate adoption and enhance the productivity of engineering teams. The systematic methodology not only offers a descriptive framework but also includes implementation and artefact references, enabling the potential deployment of the AI product. On the whole, a standard AI Product Card typically comprises a range of components that facilitate the prototyping, reuse, and deployment of AI solutions, as per the specific use case.

This work builds upon earlier work that emphasized the importance of dynamic documentation (Mehta et al., 2023) by incorporating metadata specifications for each component. This approach allows the documentation cards to become more dynamic and be ready for deployment by associating declarative specifications with each documentation section and aligning them with their corresponding implementation. Each part of the code is able to solve a modular and well-defined piece of the system that is accompanied by dedicated documentation that describes the same functionalities that the code implements. The necessary specifications that may be included in each of these components are organised in the following eight sections.

3.1.1. *Domain problem specification*

The AI Product Card approach offers the advantage of examining the AI system from various viewpoints, including those of policymakers, AI practitioners, and domain experts with specialized knowledge in AI. The **Domain Problem specification** is the portion of the AI Product Card that serves to outline the domain problem that the Card aims to address. It provides a comprehensive explanation of the use case from the perspective of the policymaker who will utilize the solution offered by the Card.

The catalogue containing the list of the **Domain Problems** that have been formulated and used by diverse public institutions may be considered a potential contribution to the implementation of the FAIR principles (Huerta et al., 2023) within the public sector. Providing open access to commonly used governmental use cases, the catalog with the **Domain Problems** aims to facilitate their reuse, thereby reducing the time and effort required for their adoption and realization of benefits. While existing efforts to provide open access to datasets and AI models¹⁰ are important (Feurer et al., 2021; Chard et al., 2018),

¹⁰ <https://huggingface.co/models>

further work is needed to contextualise these models and data to the public administration domain in order to transform them into real AI products.

This section describes the problem within the context of public administration and aims to make it easily discoverable for future reuse by other public authorities. The catalog of **Domain Problems** contained within our proposed framework will enable policymakers to quickly identify and access previously solved problems within their own domain. Additionally, the framework offers indicators of the applicability level of the AI solution that addresses the **Domain Problem**, based on mechanisms for detecting, measuring, and reporting deviations in target domain data distribution. Some of the attributes of the formal definition of this specification may include the unique name of the problem, its description, available input data specifications, output data specifications, and regulatory requirements needed. This section of the AI Product Card contains the following attributes, among which ‘regulatory requirements’ is optional:

- **id-1**: Unique identifier of the domain problem;
- **name-1**: Name of the problem;
- **description-1**: Provide some description of the problem;
- **tags-1**: Label the problem so that it can be easily filtered and accessible;
- **input-data-specifications**: Define input data specifications that the domain problem should expect to work with;
- **output-data-specifications**: Define the produced output data specification;
- **regulatory-requirements**: Validation requirements to use as a certification for the solution.

3.1.2. AI solution specification

An effective AI documentation system must establish a structure that allows for a balance between considering the big picture and focusing on the details, aligning general objectives with specific design choices (Mehta et al., 2023). Therefore, the **AI solution** serves as a comprehensive wrapper for a series of cooperative **AI tasks** that work together to tackle a complex issue outlined in the **Domain Problem** Specification. It comprises a collection of AI tasks designed to address a larger problem. This section of the Card can be likened to the **Domain Problem**, as it has the ability to address its requirements and solve the specified challenge. On the other hand, this section of the Card achieves a balance between the broader perspective and the finer details of implementing a final AI product that effectively addresses the problem at hand. This section of the AI Product Card contains the following mandatory attributes:

- **id-2**: Unique identifier of the AI solution;
- **name-2**: Name of the solution;
- **description-2**: Provide some description of the solution;
- **tags-2**: Label the solution so that it be easily filtered and accessible;
- **domain-problem-id**: Domain problem the solution implements;
- **implementation-reference**: Definition of the entry point for the solution as the implementation of the domain problem functionality;
- **input-data-specifications**: Define input data specifications that the solution should expect to work with;
- **output-data-specifications**: Define the produced output data specification.

3.1.3. AI task specification

The **AI task specifications** encompass the formal documentation of the various components that have been proposed in the literature such as data, model, validation, optimization, and adaptation techniques. An **AI task** constitutes the discrete elements of an **AI solution**. It serves as the primary building block for a complex system that can be integrated into different **AI solutions** to address specific Domain Problems.

The goal of the AI Product Card framework is to facilitate a wide range of AI tasks, with a focus on ease of reuse and adaptability. This section contains the metadata that matches the descriptive documentation parts of the task with the corresponding executable scripts or commands. In other words, it provides

different scripts that implement the logic behind running inference on a particular model or executing a training phase. This section pertains to the practical details and resources required for the training and inference phases of the *AI task*. Additionally, the validation-reference attribute offers a means to assess the quality and compliance of the available data, ensuring that it aligns with the characteristics and distribution of the data used to train the model. This is an essential step in ensuring trustworthy AI development, as it helps to prevent the incorrect use of the task on data that exhibit different characteristics or distributions. Furthermore, the *AI task* should reference the adaptation techniques that may be employed to specialize the model for its specific use case. This section of the AI Product Card contains the following attributes, among which ‘adaptation-reference’ and ‘optimization-reference’ are optional:

- **id-3:** Unique identifier of the AI task;
- **name-3:** Name of the task;
- **description-3:** Provide some description of the task;
- **task:** AI domain task categories, such as image classification, object tracking, and text classification. A value from a predefined, yet extendable, domain;
- **inputs:** Define input data specifications that the task should expect to work with;
- **outputs:** Optionally define the output data specification produced as a result of the task;
- **model-reference:** Reference to the ML model specification; it contains the reference ID of the Model Specification that makes up the AI task;
- **validation-reference:** Validation specification to be applied in order to check whether the task may be engaged in the context; it contains the reference ID of the Validation Specification that makes up the AI task;
- **adaptation-reference:** Adaptation specification that defines when the task may be adapted to the given context and how; it contains the reference ID of the Adaptation Specification that makes up the AI task;
- **optimization-reference:** Optimization specification that defines how the solution may be optimized to a context; it contains the reference ID of the Optimization Specification that makes up the AI task;
- **resources:** Specification of the resources required or recommended for the task.

3.1.4. Model specification

The appropriate specifications inside this section include the formalised version of the proposed Model Card (Mitchell et al., 2019). Some related work has already been done to provide a formalised and machine-readable format for the representation of the Model Card¹¹. Additional studies aim to build the ontology corresponding to the details within the sections of the Model Card (Amith et al., 2022; Naja et al., 2022). Considering that the Model Card discloses information about different aspects of model development and evaluation, we propose the following attributes as part of the declarative specifications for replicating it in an automated manner. This representation of the model facilitates seamless integration into the AI task definition:

- **id-4:** Unique identifier of the model;
- **name-4:** Name of the model;
- **description-4:** Description of the model;
- **library-name:** Name of the library underlying the model, such as keras, spacy, pytorch;
- **training-data:** Reference to the data used during the training phase;
- **training-script:** The path to the training script;
- **validation-data:** Reference to the data used during the validation phase;
- **metrics:** Metrics of the model specific to the type of the library

¹¹ <https://github.com/huggingface/hub-docs/blob/main/modelcard.md?plain=1>

3.1.5. Data specification

The Data Specification section contains the set of parameters that represent the sections of a typical Data Card (Pushkarna et al., 2022) documentation approach. Data Cards play a crucial role in understanding the structure and content of the dataset used for training and validating the model. They are used to assess the discrepancies between the characteristics of the training data and the characteristics of the production data, which may be utilized by public administrations. The specifications for a dataset should not only provide its structural schema but also include any necessary constraints that the dataset must meet. The metadata for the data specifications encompasses the attributes that facilitate its preparation for the training phase of the AI product. As per the data card template, we have incorporated the specifications for the data format, schema, one instance of the dataset, and a set of constraints that the dataset must meet. The list of these attributes is essential to achieve transparency and promote the reuse of the solution. This section of the AI Product Card contains the following mandatory attributes:

- **id-5:** Unique identifier of the dataset
- **name-5:** Name of the dataset
- **tags-5:** label the domain dataset so that it can be easily filtered and accessible
- **dataset-format:** Specify the category of data to be used, such as text, audio, video or tabular data
- **schema:** Depending on the type of the dataset, reference to its schema
- **reference-data-implementation-example:** Reference the dataset implementation example
- **constraints:** Define the criteria that the dataset needs to meet

3.1.6. Validation specification

The validation specification enables the provider to develop a certifiable AI solution that encompasses all crucial aspects for assessing the suitability of existing solutions from the catalog. It is essential to guarantee that the solution can be effortlessly adapted to a new context. The validation outcome provides a percentage of confidence for the repurposing of the solution, enabling policymakers to evaluate its applicability and efficacy. Data validation is often considered an iterative process that addresses the quality and distribution of real-world data that are continuously evolving, leading to a phenomenon known as domain shift. Research on how to deal with domain shift has been extensively conducted in the literature (Rabanser et al., 2019). However, systems incorporating AI solutions in the public sector nowadays require a mechanism to identify, measure, and notify deviations in terms of changes in data distribution and other profiling features, including the quantity of data accessible. In our work, the framework will provide AI practitioners with the ability to design AI Product Cards with a domain shift mindset, aimed at breaking the assumption that source and target data are independent and identically distributed (i.i.d).

To this end, the framework will integrate and establish mechanisms to identify the level of data drifting and offer essential assistance to alleviate its impact for the purpose of developing models that are robust to domain change or domain shift. As a potential open-source library for detecting data drifting in specific categories of data types, it evidently¹² may be well-suited for conducting test suites on relevant data sets. This section of the AI Product Card contains the following mandatory attributes:

- **id-6:** Unique identifier of the validation
- **library:** The name of the library the framework uses for running the validation procedures
- **description-6:** The description of the validation type
- **implementation:** Definition of the implementing function for the validation ready for execution
- **input data:** The data to be considered for validation
- **reference data:** Some referring data used for testing or benchmarks

¹² <https://docs.evidentlyai.com/>

3.1.7. *Adaptation specification*

This section of the AI Product Card aims to address the issues related to the efficient adaptation of the AI solution into a different context. When referring to context, we are describing different target domain data characteristics. Taking into account the intended reuse of the same AI system, this section will outline the necessary steps and specifications for the reference implementations of the specialized solution. Considering the use case of object classification and detection from traffic cameras in the city of Tallinn (van Noordt & Misuraca, 2022), it is highly beneficial to reproduce the same AI solution in a different context, whether it is in another urban area or even another city (Wei et al., 2016). On the other hand, when addressing Natural Language Processing (NLP) tasks, domain adaptation techniques are commonly used to minimize the disparity between different domains. For example, the task of classifying legal documents with multilingual Eurovoc thesaurus descriptors (Avram et al., 2021) is commonly employed by various EU institutions, each of which employs distinct concepts that are categorized into distinct areas or subareas. Although the same document classification task may be applicable in diverse domains, it may result in incorrect classification due to domain shifts caused by disparate distribution of target labels or classes. It would be highly advantageous in terms of time-saving and optimizing resource usage to adapt or fine-tune an existing AI solution for document classification to a different government body, rather than creating a new system from scratch that would perform the same task. This section of the AI Product Card contains the following attributes, among which the ‘base-opensource-model’ is optional:

- **id-7:** Unique identifier of the adaptation procedure
- **description-7:** The name of the library the framework uses for running the validation procedures
- **implementation:** Definition of the implementing function for the adaptation ready for execution
- **compliance-check:** Validation specification to be applied in order to check whether the task may be engaged in the context with this adaptation spec
- **prepare-local-dataset:** In particular cases, each model may employ distinct tokenizers, leading to minor variations in the preprocessed datasets.
- **base-opensource-model:** The baseline pre-trained open-source model the AI task is going to use (optional)

3.1.8. *Optimization specification*

This section presents the metadata format for the optimization techniques that may be utilized for AI-based solutions. It includes the machine-readable version of the prescriptive approach outlined in the Method Card (Adkins et al., 2022), which provides AI engineers with the necessary guidance to optimize the system. Smaller public sector institutions might encounter challenges when fine-tuning a pre-trained model since it will maintain all the parameters of the original model resulting in high resource utilisation. Many recent studies have put forth an array of optimization techniques that could potentially influence the efficacy of the AI product in terms of time and resources allotted to it. Parameter-efficient fine tuning




Figure 2. Example of domain problem specification for legislative text classification.

AI Solution Specification——

- **id-2:** 1-legislative-text-classification-solution
- **name-2:** Text classification of legal documents
- **description-2:** Text classification of legal documents
- **version:** 0.1
- **domain-problem-id:** 1-legislative-text-classification
- **tags-2:** text classification
- **implementation-reference:** solution-implementation.py
- **input-data-specifications** input-data.json
- **output-data-specifications** output-data.json

AI Task Specification——

- **id-3:** 1-text-classification
- **name-3:** Text Classification for short text
- **description-3:** Text Classification for short text
- **task:** text-classification
- **model-reference:** 1-text-classification-model // contains the value of the field labeled 'id-4'
- **validation-reference:** 1-text-classification-validation // contains the value of the field labeled 'id-6'
- **adaptation-reference:** 1-text-classification-adaptation // contains the value of the field labeled 'id-7'
- **optimization-reference:** 1-text-classification-optimization // contains the value of the field labeled 'id-8'

Model Specification——

- **id-4:** 1-text-classification-model
- **name-4:** Text Classification model
- **description-4:** Text Classification model
- **library-name:** transformers
- **training-data:** training-data.json
- **training-script:** training-model.py
- **validation-data:** validation-data.json
- **metrics:** accuracy, recall-micro

Data Specification——

- **id-5:** text-classification-eurovoc
- **name-5:** Automatic classification of legal documents dataset
- **version:** 0.1
- **tags-5:** text classification data
- **dataset-format:** text
- **reference-data-implementation-example:** 0.1
- **constraints:** constraints.json
- **schema-definition:** schema.json

Validation Specification——

- **id-6:** 1-text-classification-validation
- **description-6:** detect out of vocabulary
- **library:** evidently
- **implementation:** testsuite.py
- **input-data:** input-data.json
- **reference-data:** reference-data.json

Adaptation Specification——

- **id-7:** 1-text-classification-adaptation
- **description-7:** Adapt the task to classify legal documents of the municipal administration of another city
- **implementation:** adaptation-script.py
- **compliance-check:**
- **base-opensource-model:**
- **prepare-local-dataset:**

Optimization Specification——

- **id-8:** 1-text-classification-optimization
- **description-8:** apply LoRA optimization technique
- **implementation:** finetune/lora.py
- **metric:** memory-saving
- **range:** [16 - 20 GB]

Figure 3. Example of AI product card for legislative text classification.

(Houlsby et al., 2019) and Low-rank adaptation (LoRA) (Hu et al., 2021) are widely recognized and highly effective techniques for adapting custom LLMs in an efficient manner on downstream tasks. Accordingly, the optimization section will contain the most effective configurations for implementing them in various applications. These configurations may encompass the base model, the batch size, the matrix rank, the alpha parameter, the target dataset, and the adaptation scripts that need to be applied to it. This section of the AI Product Card contains the following attributes, among which 'range' is optional:

- **id-8:** Unique identifier of the optimization technique
- **description-8:** Description of the optimization technique

- **implementation:** Definition of the implementing function for the optimization ready for execution
- **metric:** The metric specification
- **range:** Possible value ranges

3.2. An example of an AI Product Card

The AI Product Card provided in this section exemplifies how a team of different stakeholders in the public sector, including policymakers, AI engineers and functional specialists can populate the attributes of the specifications provided in [Section 3.1](#). The card describes the task of text classification for legal documents (Rovera et al., n.d.). From an industrial perspective, many practical tools have been developed for orchestrating Machine Learning (ML) applications lifecycle. Our objective is to establish a connection between the orchestration of these applications and the specifications outlined in the AI Product Card documentation framework. The different stages of the code-bound documentation can be interpreted as the various steps of the ML pipeline, which may include services, datasets, functions or other executable code or artifacts. In this context, MLRun¹³ has been deemed suitable for experimenting with the proposed framework as its organizational structure of a Machine Learning (ML) project conforms to the modular documentation concepts presented in the AI Product Card. The [Figure 2](#) and [Figure 3](#) represent the suggested sections of the AI Product Card for the problem of legal document classification. We provide an initial version of the AI Product Card specifications, along with accompanying reference implementations on an openly accessible github repository.¹⁴

4. Conclusion

In contrast to other commonly employed methodologies for AI documentation that aim to provide descriptive or prescriptive documentation directives to the AI system design, we have proposed a machine-readable declarative framework that allows the adaptation of the same solution in different contexts. Indeed, the proposed AI Product Cards framework aims to improve the reuse and adaptation of AI products in public administration by leveraging the machine-readable specifications obtained by the descriptive documentation cards. Specifically, the framework encompasses necessary descriptive and declarative specifications that orchestrate the execution of the accompanying code references and other artifacts that allow for the seamless reuse of the entire AI product. As an example, we have discussed in detail an AI Product Card designed for the legal document classification task.

The framework allows to accommodate diverse AI tasks in a generalizable template. The initial specification template is capable of capturing the needs of the lifecycle of an AI solution. However, one limitation of our proposed framework is the difficulty of formalizing each phase of AI-based solutions due to their diverse needs that vary on a per-task basis. To address this, it may be necessary to extend the metadata specifications over time using a repeatable approach, in order to accommodate the specific needs of training procedures or optimization parameters. Another limiting aspect of the framework, in certain situations, is the absence of a well-defined pattern which can make its representation difficult in a concise and explicit manner. For instance, evaluating the output generated from Large Language Models (LLM) may require manual and empirical processes that involve human feedback.

A possible future extension of the proposed framework could be the exploitation of the recent advancements in the research direction of formalising model card report information through semantic representation (Amith et al., 2022; Naja et al., 2022) as well as in the application to scenarios outside the public administration.

Data availability statement. We provide an initial version of the AI Product Card specifications, along with accompanying reference implementations on an openly accessible GitHub repository¹⁵.

¹³ <https://www.mlrun.org/>

¹⁴ https://github.com/smartcommunitylab/ai_product_card_templates

¹⁵ https://github.com/smartcommunitylab/ai_product_card_templates

Author contribution. A.C. wrote the paper. A.C. and R.K. conceptualised and designed the proposed approach. B.L. and R.K. supervised the research direction and the writing process. A.C. and A.P.A. contextualised the example of the proposed AI Product Card. All authors read and approved the final manuscript.

Provenance. This article is part of the Data for Policy 2024 Proceedings and was accepted in Data & Policy on the strength of the Conference’s review process.

Funding statement. The authors acknowledge the support of the “AIxPA—Artificial Intelligence in the Public Administration system” area of Flagship Project—Progetto Bandiera PNC-A.1.3 “Digitalizzazione della pubblica amministrazione della Provincia Autonoma di Trento.” Bruno Lepri also acknowledges the support of the PNRR project FAIR—Future AI Research (PE00000013), under the NRRP MUR program funded by the NextGenerationEU.

Competing interest. The authors declare no competing interests exist.

References

- Adkins D, Alsallakh B, Cheema A, Kokhlikyan N, McReynolds E, Mishra P, Procope C, Sawruk J, Wang E and Zvyagina P (2022) Method cards for prescriptive machine-learning transparency. In *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI*, pp. 90–100. <https://doi.org/10.1145/3522664.3528600>
- Alvarenga A, Matos F, Godina R and C O Matias J (2020) Digital transformation and knowledge management in the public sector. *Sustainability* 12(14), 5824. <https://doi.org/10.3390/su12145824>
- Amith MT, Cui L, Zhi D, Roberts K, Jiang X, Li F, Yu E and Tao C (2022) Toward a standard formal semantic representation of the model card report. *BMC Bioinformatics* 23(S6), 281. <https://doi.org/10.1186/s12859-022-04797-6>
- Arnold M, Bellamy RKE, Hind M, Houde S, Mehta S, Mojsilovic A., Nair R, Ramamurthy KN, Reimer D, Olteanu A, Piorkowski D, Tsay J and Varshney KR (2019, February) FactSheets: Increasing trust in AI services through supplier’s declarations of conformity. <http://arxiv.org/abs/1808.07261>. Comment: 31 pages.
- Avram A-M, Pais V, and Tufis D (2021, August) PyEuroVoc: A tool for multilingual legal document classification with EuroVoc descriptors [arXiv:2108.01139 [cs]]. <http://arxiv.org/abs/2108.01139>. Comment: RANLP2021.
- Bertot J, Estevez E and Janowski T (2016) Universal and contextualized public services: Digital public service innovation framework. *Government Information Quarterly* 33(2), 211–222. <https://doi.org/10.1016/j.giq.2016.05.004>
- Bérubé M, Giannelia T and Vial G (2021) Barriers to the implementation of AI in organizations: Findings from a Delphi study. <https://doi.org/10.24251/HICSS.2021.805>
- Chard R, Li Z, Chard K, Ward L, Babuji Y, Woodard A, Tuecke S, Blaiszik B, Franklin MJ and Foster I (2018, November) DLHub: Model and data serving for science [arXiv:1811.11213 [cs, stat]]. <http://arxiv.org/abs/1811.11213>. Comment: 10 pages, 8 figures, conference paper.
- Chaudhry MA, Cukurova M and Luckin R (2022) A transparency index framework for AI in education.
- Cobb O and Looveren AV (2022) Context-aware drift detection, p. 4
- Derczynski L, Kirk HR, Balachandran V, Kumar S, Tsvetkov Y, Leiser MR and Mohammad S (2023, March) Assessing language model deployment with risk cards [arXiv:2303.18190 [cs]]. <http://arxiv.org/abs/2303.18190>
- Desouza KC, Dawson GS and Chenok D (2020) Designing, developing, and deploying artificial intelligence systems: Lessons from and for the public sector. *Business Horizons* 63(2), 205–213. <https://doi.org/10.1016/j.bushor.2019.11.004>
- Feurer M, van Rijn JN, Kadra A, Gijbbers P, Mallik N, Ravi S, Müller A, Vanschoren J and Hutter F (2021, June) OpenML-Python: An extensible Python API for OpenML [arXiv:1911.02490 [cs, stat]]. <http://arxiv.org/abs/1911.02490>
- Ganin Y. and Lempitsky V (2015) Unsupervised domain adaptation by backpropagation.
- Gebru T, Morgenstern J, Vecchione B, Vaughan JW, Wallach H, Iii HD and Crawford K (2021) Datasheets for datasets. *Communications of the ACM* 64(12), 86–92. <https://doi.org/10.1145/3458723>
- Hind M, Houde S, Martino J, Mojsilovic A, Piorkowski D, Richards J and Varshney KR (2019, November) Experiences with improving the transparency of AI models and services [arXiv:1911.08293 [cs]]. <http://arxiv.org/abs/1911.08293>
- Holland S, Hosny A, Newman S, Joseph J and Chmielinski K (2018) The dataset nutrition label: A framework to drive higher data quality standards.
- Houlsby N, Giurgiu A, Jastrzebski S, Morrone B, de Laroussilhe Q, Gesmundo A, Attariyan M and Gelly S (2019, June) Parameter-efficient transfer learning for NLP [arXiv:1902.00751 [cs, stat]]. <http://arxiv.org/abs/1902.00751>
- Hu EJ, Shen Y, Wallis P, Allen-Zhu Z, Li Y, Wang S, Wang L and Chen W (2021, October) LoRA: Low-rank adaptation of large language models [arXiv:2106.09685 [cs]]. <http://arxiv.org/abs/2106.09685>. Comment: Draft V2 includes better baselines, experiments on GLUE, and more on adapter latency.
- Huerta EA, Blaiszik B, Brinson LC, Bouchard KE, Diaz D, Doglioni C, Duarte JM, Emani M, Foster I, Fox G, Harris P, Heinrich L, Jha S, Katz DS, Kindratenko V, Kirkpatrick CR, Lassila-Perini K, Madduri RK, Neubauer MS, ... Zhu R (2023) FAIR for AI: An interdisciplinary and international community building perspective. *Scientific Data* 10(1), 487. <https://doi.org/10.1038/s41597-023-02298-6>

- Hupont I, Fernández-Llorca D, Baldassarri S and Gómez E** (2023, June) Use case cards: A use case reporting framework inspired by the European AI Act [arXiv:2306.13701 [cs] version: 1]. <http://arxiv.org/abs/2306.13701>. Comment: 33 pages, 9 figures.
- John MM, Holmström Olsson H and Bosch J** (2021) Architecting AI deployment: A systematic review of state-of-the-art and state-of-practice literature [Series Title: Lecture Notes in Business Information Processing]. In E. Klotins & K. Wnuk (eds.), *Software Business*, Vol. 407. Springer International Publishing, pp. 14–29. https://doi.org/10.1007/978-3-030-67292-8_2
- Kamal Paul S and Bhaumik P** (2022) Disaster Management through integrative AI. In *Proceedings of the 23rd International Conference on Distributed Computing and Networking*, pp. 290–293. <https://doi.org/10.1145/3491003.3493235>
- Königstorfer F and Thalmann S** (2022) AI documentation: A path to accountability. *Journal of Responsible Technology* 11, 100043. <https://doi.org/10.1016/j.jrt.2022.100043>
- Li B, Qi P, Liu B., Di S, Liu J, Pei J, Yi J and Zhou B** (2023) Trustworthy AI: From principles to practices. *ACM Computing Surveys* 55(9), 1–46. <https://doi.org/10.1145/3555803>
- Liu D-Y, Chen S-W and Chou T-C** (2011) Resource fit in digital transformation: Lessons learned from the CBC Bank global e-banking project. *Management Decision* 49(10), 1728–1742. <https://doi.org/10.1108/00251741111183852>.
- Madan R and Ashok M** (2023) AI adoption and diffusion in public administration: A systematic literature review and future research agenda. *Government Information Quarterly* 40(1), 101774. <https://doi.org/10.1016/j.giq.2022.101774>
- Martie L, Rosenberg J, Demers V, Zhang G, Bhardwaj O, Henning J, Prasad A, Stallone M, Lee JY, Yip L, Adesina D, Paikari E, Resendiz O, Shaw S and Cox D** (2023, February) Rapid development of compositional AI. <http://arxiv.org/abs/2302.05941>. Comment: Accepted to ICSE 2023, NIER track.
- McMillan-Major A, Osei S, Rodriguez JD, Ammanamanchi PS, Gehrman S and Jernite Y** (2021) Reusable templates and guides for documenting datasets and models for natural language processing and generation: A case study of the HuggingFace and GEM data and model cards. In *Proceedings of the 1st Workshop on Natural Language Generation, Evaluation, and Metrics (GEM 2021)* <https://doi.org/10.18653/v1/2021.gem-1.11>
- Medaglia R, Gil-Garcia J. R. and Pardo TA** (2023) Artificial intelligence in government: Taking stock and moving forward. *Social Science Computer Review* 41(1), 123–140. <https://doi.org/10.1177/08944393211034087>
- Mehr H** (2017) Artificial intelligence for citizen services and government.
- Mehta S, Rogers A and Gilbert TK** (2023, March) Dynamic documentation for AI systems [arXiv:2303.10854 [cs]]. <http://arxiv.org/abs/2303.10854>
- Micheli M, Hupont I, Delipetrev B and Soler-Garrido J** (2023) The landscape of data and AI documentation approaches in the European policy context. *Ethics and Information Technology* 25(4), 56. <https://doi.org/10.1007/s10676-023-09725-7>
- Mikalef P, Fjortoft SO and Torvatn HY** (2019) Artificial intelligence in the public sector: A study of challenges and opportunities for Norwegian municipalities. In IO Pappas, P Mikalef, YK Dwivedi, L Jaccheri, J Krogstie and M Mäntymäki (eds.), *Digital Transformation for a Sustainable Society in the 21st Century*, Vol. 11701. Springer International Publishing, pp. 267–277. https://doi.org/10.1007/978-3-030-29374-1_22. Series Title: Lecture Notes in Computer Science.
- Mikalef P, Lemmer K, Schaefer C, Ylinen M, Fjortoft SO, Torvatn HY, Gupta M and Niehaves B** (2022) Enabling AI capabilities in government agencies: A study of determinants for European municipalities. *Government Information Quarterly* 39 (4), 101596. <https://doi.org/10.1016/j.giq.2021.101596>
- Mitchell M, Wu S, Zaldivar A, Barnes P, Vasserman L, Hutchinson B, Spitzer E, Raji ID and Gebru T** (2019) Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. <https://doi.org/10.1145/3287560.3287596>
- Naja I, Markovic M, Edwards P, Pang W, Cottrill C and Williams R** (2022) Using knowledge graphs to unlock practical collection, integration, and audit of aiaccountability information. *IEEE Access* 10, 74383–74411. <https://doi.org/10.1109/ACCESS.2022.3188967>
- Prabhu VU and Birhane A** (2020) Large image datasets: A pyrrhic win for computer vision?
- Pushkarna M, Zaldivar A and Kjartansson O** (2022) Data cards: Purposeful and transparent dataset documentation for responsible ai.
- Rabanser S, Günemann S and Lipton ZC** (2019, October) Failing loudly: An empirical study of methods for detecting dataset shift [arXiv:1810.11953 [cs, stat]]. <http://arxiv.org/abs/1810.11953>. Comment: Advances in Neural Information Processing Systems (NeurIPS) 2019.
- Rostamzadeh N, Mincu D, Roy S, Smart A, Wilcox L, Pushkarna M, Schrouff J, Amironei R, Moorosi N and Heller K** (2022) Healthsheet: Development of a transparency artifact for health datasets.
- Rovera M, Aproso AP, Greco F, Lucchese M, Tonelli S and Antetomaso A** (n.d.) Italian legislative text classification for Gazzetta Ufficiale. *Proceedings of NLLP Workshop 2023*.
- Sun B and Saenko K** (2016) Deep coral: Correlation alignment for deep domain adaptation.
- Tagliabue J., Tuulos V, Greco C and Dave V** (2021, November) DAG card is the new model card [arXiv:2110.13601 [cs]]. <http://arxiv.org/abs/2110.13601> Comment: Accepted at DCAI @ Neurips 2021, camera-ready version. arXiv admin note: substantial text overlap with arXiv:2106.03907.
- Van Noordt C and Misuraca G** (2022) Artificial intelligence for the public sector: Results of landscaping the use of AI in government across the European Union. *Government Information Quarterly* 39(3), 101714. <https://doi.org/10.1016/j.giq.2022.101714>

- Wei Y, Zheng Y and Yang Q** (2016) Transfer knowledge between cities. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1905–1914. <https://doi.org/10.1145/2939672.2939830>
- Wirtz BW, Weyerer JC and Geyer C** (2019) Artificial intelligence and the public sector—Applications and challenges. *International Journal of Public Administration* 42(7), 596–615. <https://doi.org/10.1080/01900692.2018.1498103>
- Zhang B, Liu T, Liang P, Wang C, Shahin M and Yu J** (2023) Architecture decisions in AI-based systems development: An empirical study [ISSN: 2640–7574]. In *2023 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pp. 616–626. <https://doi.org/10.1109/SANER56733.2023.00063>