

University of Trento

Doctoral Programme in Information Engineering and Computer Science

Human-Centred AI Compliance in Common European Data Spaces

Supervisor
Dott. Maurizio Napolitano
Fondazione Bruno Kessler

PhD Candidate
Dott. Pietro Giovanni Bizzaro
239766

Final examination date: 13 October 2026

Abstract

Common European Data Spaces (CEDS) are intended to support governed data sharing and reuse across strategic sectors. Public administrations supply data to these spaces and seek to reuse it in AI projects. Access alone, however, does not establish whether a downstream use is authorised, ethically justified, or supported by adequate evidence. This thesis investigates how EU data and AI regulation can be translated into organisational and metadata-based tools that support public administrations without displacing competent human and institutional judgment.

The research combines comparative legal and policy analysis; formative development of the AI Ethics Canvas through action research, value-sensitive design, twelve interviews, and a workshop with eighteen participants; and a conceptual examination of selected legal sources and technical standards. It finds that regulatory ambition and formal rights yield meaningful agency only when backed by material, technical, and administrative capacity. CEDS can strengthen governed access and exchange, but address only part of those dependencies. Within public administrations, the Canvas provides a shared setting in which technical and non-technical participants can formulate and revisit ethical, legal, social, and governance questions, identify evidence needs, and expose unresolved responsibilities. Data stewardship maintains continuity among governed data, metadata, evidence, decisions, versions, and review states, while competent actors retain authority.

The informational contribution is a five-stage, source-linked framework that moves from a legal source to a reviewed legal concept or requirement, a metadata need, an exact standard element, and a bounded representational outcome. Two bounded micro-applications show how exact provenance elements can support a collection-methodology relation and why descriptive sufficiency remains a question for competent judgment.

Together, these contributions define human-centred compliance as the maintained connection among institutional capacity, structured deliberation, data stewardship, technical traceability, and reassessment. They support reviewable decisions without equating metadata, workshop completion, or automated checks with compliance.

Keywords

Common European Data Spaces; human-centred AI; digital sovereignty; public-sector artificial intelligence; public administration; metadata standards; regulatory compliance; AI ethics.

Contents

Abstract i

Keywords iii

List of Tables ix

List of Figures xi

1 Introduction 1
- 1.1 A dataset's journey . . . 2
- 1.2 The Problem . . . 3
- 1.3 Human-centred compliance . . . 4
- 1.4 Three complementary capacities . . . 7
  - 1.4.1 Strategic and infrastructural capacity . . . 8
  - 1.4.2 Organisational and deliberative capacity . . . 8
  - 1.4.3 Informational and technical capacity . . . 9
- 1.5 Scope, method, and evidence boundaries . . . 10
- 1.6 Research questions . . . 12
- 1.7 Contribution of the Thesis . . . 12
- 1.8 Structure of the Thesis . . . 13

2 State of the Art 15
- 2.1 Review Method . . . 15
- 2.2 Data Spaces and Downstream Authority . . . 17
  - 2.2.1 Sectoral implementation initiatives . . . 18
- 2.3 Access and Downstream Use . . . 20
- 2.4 Conditional Applicability . . . 22
- 2.5 EHDS as a Sectoral Reference . . . 26
- 2.6 Documentation and Traceability . . . 27
- 2.7 Formalisation and Legal Judgment . . . 31
- 2.8 Limits of Technical Enforcement . . . 33
- 2.9 Authority in Human Oversight . . . 35
- 2.10 Bounded Evaluation . . . 37
- 2.11 The Integration Gap . . . 38

3 Digital Sovereignty, Regulatory Capacity, and Common European Data Spaces **41**
3.1 Sovereignty and Capacity . . . 41
3.1.1 From territorial authority to individual agency . . . 43
3.1.2 Data, platforms, and infrastructure . . . 44
3.2 Question, Comparison, and Method . . . 45
3.2.1 Rationale for the EU–China Comparison . . . 46
3.2.2 Boundaries of comparability . . . 47
3.3 Comparative Legal Analysis . . . 47
3.3.1 Regulatory architectures and the allocation of authority . . . 47
3.3.2 Platforms, algorithms, and the position of the user . . . 48
3.3.3 Personal data, data flows, and extraterritorial influence . . . 49
3.4 Comparative Findings . . . 51
3.4.1 Results and Evidentiary Boundaries . . . 52
3.4.2 Data, cloud, and compute in the AI competition . . . 53
3.4.3 Common European Data Spaces as a partial response . . . 54
3.4.4 Limits of Data Spaces . . . 55
3.5 Three Capacities in Data Governance . . . 56
3.5.1 Public Administrations in Data Governance . . . 57
3.6 Organisational Capacity and Deliberation . . . 58
3.7 Conclusion . . . 59

4 The AI Ethics Canvas: Designing and Reviewing Public-Sector AI **61**
4.1 From regulatory principles to an organisational practice . . . 61
4.2 A Canvas for Public-Sector AI . . . 63
4.2.1 AI Adoption as a Socio-Technical Decision . . . 63
4.2.2 Canvases and Structured Deliberation . . . 64
4.2.3 Value-sensitive design and action research . . . 64
4.3 Research and design process . . . 65
4.3.1 Institutional Setting . . . 65
4.3.2 Documentary and legal foundations . . . 66
4.3.3 Iterative drafting and the 2024 artifact . . . 67
4.3.4 Ethics, confidentiality, and analytical boundaries . . . 67
4.4 Canvas Structure . . . 68
4.4.1 Three complementary roles . . . 70
4.4.2 Data . . . 70
4.4.3 Algorithms . . . 72
4.4.4 Methods of analysis . . . 72
4.4.5 Socio-cultural elements . . . 75
4.4.6 Functional requirements . . . 75
4.5 Formative Evidence . . . 77
4.5.1 Contextual needs evidence: survey and focus-group activities . . . 77
4.5.2 Direct formative evaluation: twelve interviews . . . 78
4.5.3 Direct formative application: the 2025 workshop . . . 79
4.5.4 Findings and Evidentiary Boundaries . . . 81
4.6 Iterative Project Review . . . 81
4.6.1 Initial Framing . . . 82
4.6.2 Iterative review during design and operation . . . 82
4.6.3 Final Review . . . 83

4.7 Information and Evidence Touchpoints . . . 83
4.8 Relation to Common European Data Spaces . . . 84
4.9 Limits, risks, and conditions for responsible use . . . 84
4.10 Conclusion . . . 86

5 From EU Legal Concepts to Metadata Standards 87
5.1 The Informational Problem . . . 87
5.2 The Five-Stage Analytical Chain . . . 89
5.2.1 Stage 1: legal source . . . 89
5.2.2 Stage 2: source-linked legal concept or requirement . . . 90
5.2.3 Stage 3: metadata need . . . 91
5.2.4 Stage 4: exact standard element . . . 94
5.2.5 Stage 5: representational outcome . . . 94
5.2.6 Transversal Safeguards . . . 95
5.3 Operationalising the Analytical Chain . . . 95
5.3.1 A governed, source-linked corpus . . . 97
5.3.2 From a bounded exploration to operational use . . . 99
5.3.3 An assessment-supporting knowledge graph . . . 100
5.3.4 Standards as modular building blocks . . . 102
5.4 Applying the Conceptual Framework to Data Act Article 33(1)(a) . . . 103
5.4.1 Collection methodology . . . 104
5.4.2 Sufficiency of the description . . . 105
5.5 Sectoral Transfer: The EHDS . . . 106
5.6 A Prospective Minimal Profile . . . 107
5.7 Validity and Limitations . . . 107
5.8 Conclusion . . . 108

6 Synthesis and Conclusions 109
6.1 The Dataset Journey Revisited . . . 109
6.2 Operationalising EU Data and AI Regulation . . . 111
6.3 Digital sovereignty, material capacity, and CEDS . . . 112
6.4 The Canvas in public-sector AI design and review . . . 113
6.5 The Conceptual Framework . . . 115
6.6 Evidence and Proposals . . . 116
6.7 Original Contribution . . . 118
6.8 Implications for policy, administration, and standardisation . . . 119
6.9 Limitations . . . 120
6.10 Possible Directions for Empirical Evaluation . . . 121
6.11 Human-Centred Compliance in Practice . . . 122

Bibliography 123

A Literature review audit protocol 143
A.1 Purpose and scope . . . 143
A.2 Databases and access . . . 143
A.3 Complete search strings . . . 143
A.3.1 Scopus . . . 143
A.3.2 Supplementary OpenAlex search . . . 144
A.4 Contrary-evidence searches . . . 144
A.5 Eligibility criteria . . . 145

A.6 Screening procedure 146
A.7 Deduplication procedure 146

B Legal and Technical Source Inventory for the Conceptual Framework 147
B.1 Purpose and scope 147
B.2 Legal sources 147
B.3 Technical specifications 149
B.4 How the inventory should be read 150

Publications 151

Data Availability 153

List of Tables

2.1 Database retrieval yields and captured screening samples, 17 July 2026 16
2.2 Contributions of the reviewed approaches and remaining integration tasks . . 38

5.1 Metadata-need categories used by the framework 93
5.2 Representational outcomes defined by the framework 95
5.3 Transversal safeguards for the five-stage chain 96

6.1 A bounded illustration of the health-data journey 110
6.2 Evidentiary bases, conceptual proposals, and retained boundaries 117

A.1 Retrieval yields for the final literature-review queries 145

B.1 Selected legal instruments and their function in the conceptual framework . . 148
B.2 Technical specifications considered by the conceptual framework 149

List of Figures

1.1 The three complementary levels and the cross-cutting continuity function connected by the thesis. Data stewardship maintains relations across data-space participation, organisational decisions, evidence, and technical records; the diagram does not represent a validated end-to-end system. 7

4.1 AI Ethics Canvas: overview and legend 68
4.2 AI Ethics Canvas: overview and legend 69
4.3 AI Ethics Canvas: overview and legend 71
4.4 AI Ethics Canvas: overview and legend 73
4.5 AI Ethics Canvas: overview and legend 74
4.6 AI Ethics Canvas: overview and legend 76
4.7 AI Ethics Canvas: overview and legend 77

5.1 The single five-stage conceptual framework 89

Chapter 1

Introduction

My doctoral path began with a background in law, within a research environment shaped mainly by computer science. This context made visible the gap between identifying a legal issue and translating it into operational terms within a technical and organisational setting. A rule can guide legal analysis while still leaving open what information should be collected, who should decide, how the decision should be documented, and when it should be reconsidered.

The difficulty is not resolved by a more precise technical vocabulary. Operationalisation is also a matter of organisation and change management. In a public administration, legal, domain, managerial, and technical perspectives must be coordinated; responsibilities assigned; evidence maintained; and decisions revisited as the purpose, data, system, or applicable law changes. Technical representations can support this work, but they cannot determine the authoritative interpretation or the adequacy of the organisational process.

In this setting, the AI Ethics Canvas functions as a hybrid organisational tool. Its questions create touchpoints at which participants can articulate assumptions, information needs, disagreements, and responsibilities throughout an AI project. Information collected at those touchpoints may enrich project documentation and metadata used in later compliance and assessment activities. The Canvas does not itself generate metadata, establish compliance, or operate as a required stage of the conceptual framework.

The analysis proceeds across material capacity, organisational deliberation, and technical traceability. Its central premise is that EU data and AI regulation becomes operational only when legal meaning remains connected to responsible actors, supporting evidence, and the technical records through which decisions are implemented and reviewed. The implications of that premise become clearer when followed through the journey of a dataset.

1.1 A dataset's journey

The problem becomes concrete in an illustrative case. A regional health authority holds a dataset that is to be made discoverable through a European data space. Before another organisation can request access, the relevant public actors must identify the resource, prepare and maintain its catalogue description, document its provenance and quality, state the applicable conditions, and allocate responsibility for keeping that information current. Public administrations therefore participate in data spaces not only as prospective data users, but also as data holders, providers, and governance actors.

A public research organisation may subsequently discover the dataset and propose a specific downstream use, such as developing or validating an AI system for a hospital. The catalogue can make the dataset visible, and an access procedure can govern the request, but neither step settles the proposed operation. Its purpose, actors, data categories, transformations, system classification, evidence, safeguards, and assessment date may all affect whether and how it can proceed under the relevant regulations [1, 2, 3].

The journey also divides work between technical components and people. On the technical side, identifiers, metadata, policy expressions, validation rules, access controls, and logs can carry a reviewed decision into a workflow and reveal when a recorded condition is no longer satisfied. On the organisational side, the prospective data users—including the researchers who formulate the purpose and intend to perform the downstream operation—must provide information about the proposed use, the planned transformations, and the system to be developed or validated. They interact with the functions that already maintain data-governance, legal, compliance, security, and data protection processes within the organisations involved. Depending on the organisation and the applicable rules, these functions may include a Data Protection Officer, where one is designated, as well as legal counsel, compliance officers, information-security personnel, data managers, and other actors responsible for the relevant governance process.

These participants do not perform the same task or exercise the same authority. They assemble and maintain the dataset description, provenance, quality information, permissions, evidence, responsible actors, and lifecycle state; identify missing or outdated information; and refer legal, ethical, security, or institutional questions to the actor competent to address them. The technical system can make the resulting decisions actionable and inspectable, but it does not acquire the authority of the person or institution responsible for them.

This scenario is illustrative rather than empirical and anticipates no legal outcome. It provides a concrete point of entry into the problem studied in the thesis: how legal requirements, technical representations, organisational responsibilities, and later reassessment can remain connected as a dataset moves from discovery to a proposed use.

1.2 The Problem

In this thesis, a data space denotes a federated environment rather than a single repository. Independent participants can make data discoverable and available under shared governance, interoperability, and trust arrangements, while data and control may remain distributed among their holders [4, 5]. The 2020 European strategy for data placed Common European Data Spaces (CEDS) at the centre of a broader plan to build a single market for data. It envisaged interoperable spaces in strategic sectors, including health, mobility, energy, agriculture, public administration, and research, as a means of reducing legal and technical fragmentation and enabling reuse in accordance with EU law [6, 7].

The preceding example unfolds across precisely these organisational and technical boundaries. Data-space governance and shared infrastructure can support discovery, access negotiation, policy expression, transfer, and bounded forms of enforcement. Yet each subsequent operation must be assessed in its own context: who intends to use the data, what they propose to do with them, and the legal and operational circumstances surrounding that use may all change as the system develops over time. Access begins this process without completing it in legal or organisational terms.

The thesis addresses the interface between legal interpretation and technical representation. Legal interpretation establishes who has authority, to whom a rule applies, under which circumstances, and with what effect. Technical representation structures and exchanges the information needed to test conformance, track changing dependencies, and connect a reviewed decision to a control or workflow. The translation fails when a legally meaningful distinction is flattened into a field, when a technical check is presented as a legal conclusion, or when no person or institution remains responsible for maintaining and reviewing the record.

These difficulties form part of the wider question of digital sovereignty, which has become increasingly important as AI turns control over data, computing capacity, platforms, and infrastructure into a source of economic and political power [8, 9]. The European Union has sought to establish a distinctive position in a technological landscape dominated industrially by the United States and China, relying substantially on rights-based regulation to protect individual autonomy and project European values. A comparison of European and Chinese AI and data regulation nevertheless shows why regulatory capacity alone cannot deliver individual digital sovereignty. Formal rights may give people greater legal control over their data, while their practical choices remain constrained by dependence on non-European service providers and by the limited availability of alternative industrial and infrastructural capacity within Europe. The problem therefore lies in the gap between a formally recognised capacity to choose and the material means needed to exercise that choice, rather than in regulation itself [9].

A parallel pressure concerns the digital commons on which AI development extensively relies. These commons encompass not only openly accessible collections, but also the com-

munities, institutions, commitment, and infrastructure through which knowledge is created, curated, documented, and kept useful [10, 11]. Generative AI may strengthen these resources, for example by supporting translation or metadata enrichment, but current extractive practices also create risks of erosion. Intensive crawling transfers infrastructural and environmental costs to the organisations that host open resources; reduced traffic may weaken incentives to contribute new knowledge [12]; resistance to indiscriminate extraction may encourage the closure of previously open collections [13]; and synthetic content may increase the burden of moderation and quality control. Openness alone does not account for these dependencies or ensure that the values and preferences of contributing communities survive when their material is repurposed for AI training [11, 14].

This pressure intersects directly with the rationale for data spaces. A data space cannot replace a digital commons, because shared infrastructure does not by itself create the community, reciprocity, or continuing commitment on which a commons depends. It can, however, provide a governed setting in which valuable data are made discoverable and reusable while their provenance, conditions of use, quality, responsible participants, and supporting evidence remain visible. In this sense, data spaces may contribute to Europe's search for greater practical capacity over data, but only if their governance recognises the people and institutions that produce and maintain those resources rather than treating data as an inexhaustible input. A workable solution must connect the legal rule and the technical representation to the people who supply evidence, preserve context, exercise judgment, and respond to change.

The problem draws on several bodies of work. Legal scholarship examines how rules are interpreted and applied, while research on metadata and information systems considers how relevant information can be structured and exchanged. Work on data spaces and human-centred governance adds the organisational context in which technical controls, professional judgment, and responsibility operate. Each perspective illuminates part of the dataset's journey.

Viewed together, these literatures suggest that the central difficulty may lie less in the absence of individual methods than in the connections among them. This thesis investigates that possibility by asking how a reviewed legal requirement can inform technical representation and organisational practice without losing the context and judgment on which it depends.

1.3 Human-centred compliance

The expression "human-centred compliance" is used in this thesis to describe a process, not a certification. The process is human-centred because the people who know the public purpose, legal context, data, technical system, affected communities, and institutional responsibilities take part in formulating and reviewing the relevant questions. It is compliance-oriented

because those questions are connected to current sources, documented decisions, supporting evidence, responsible actors, and conditions for reassessment. It remains a process because no workshop, checklist, metadata graph, or validation rule can conclusively establish that a changing socio-technical project is lawful and ethically acceptable.

This definition does not oppose people to automation. A public administration may properly use catalogue services, workflow engines, access controls, policy languages, shape validation, provenance records, and monitoring. Some propositions can be checked mechanically once their terms and inputs have been authoritatively fixed. The human-centred boundary concerns what those operations are allowed to mean. A machine may confirm that a required field is present or that a declared date has not passed; it cannot infer that an ambiguous legal provision applies, that the declaration is true, that the evidence is sufficient, or that the public value of a project justifies its effects. Human involvement is meaningful only where participants receive adequate information, have an opportunity to challenge assumptions, and can direct or escalate the decision rather than merely endorse an output [15].

The same distinction applies to the word "compliance". The thesis does not define compliance as correspondence between a legal sentence and a database field. Legal requirements are interpreted in relation to facts, institutional competence, temporal effect, and other sources. Organisational decisions translate that interpretation into responsibilities, design choices, and evidence. Metadata can then preserve selected assertions and make bounded checks possible. The accuracy and authority of each handoff must remain open to review. Human-centred compliance is therefore best understood as the continuing capacity to ask the right questions, retain the reasons and evidence for provisional answers, and reopen those answers when a relevant dependency changes.

Public administrations make this problem especially visible. They pursue public purposes under legal mandates, rely on procurement and external providers, process information about communities, and remain subject to administrative-law duties that are not exhausted by technical performance. An AI project may involve domain officials, data holders, developers, legal and data-protection functions, information-security personnel, procurement officers, managers, and representatives of affected groups. Those actors use different vocabularies and encounter the project at different times. A structured artifact can create a shared deliberative space, while source-linked records can preserve selected outcomes after the meeting. Neither function is useful if the organisation lacks the resources, infrastructure, or authority needed to act on what it learns.

Human-centredness also extends beyond the professionals assembled by the administration. The people and communities affected by an AI-supported service may know that a category, proxy, data gap, or procedural assumption produces consequences that are invisible in a technical specification. Their involvement cannot be reduced to a generic stakeholder field or to a consultation box marked as complete. The appropriate form of participation will

vary with the project and may be constrained by confidentiality, procurement, security, or statutory decision procedures. The relevant principle is that the organisation should identify whose knowledge and interests are material, record how they were considered, and make the limits of participation visible. A metadata record can identify a consultation document or responsible body; it cannot establish that participation was representative or that objections were adequately answered.

CEDS intensify this coordination problem because information and authority are distributed. A catalogue operator may know how a dataset is described, a data holder may know how it was produced, an access body may control one route of reuse, and the receiving administration may alone know the downstream purpose and system design. A supplier may operate the model while another public body remains accountable for the service. No participant necessarily possesses the whole legal and technical picture. Human-centred compliance therefore requires that exchanges between participants preserve provenance and qualifications instead of assuming that trust in one participant makes every downstream assertion reliable.

Traceability is valuable in this setting, but it is not synonymous with transparency. A chain may be internally traceable and still be inaccessible to affected people, based on weak evidence, or organised around an unreasonable interpretation. Conversely, not every sensitive record can be published. The thesis uses traceability in the narrower sense of maintaining inspectable relations among a source, a proposition, a decision, supporting material, a responsible actor, and a later representation. Questions of public disclosure, access to documents, trade secrets, personal data, and procedural rights remain governed by their applicable regimes. This narrower definition allows the technical work to be assessed without claiming that a source-linked graph resolves the wider politics of transparency.

Nor does a human-centred approach imply that every decision must be reopened continuously. Reassessment consumes scarce public resources and can itself make administration unstable. The practical need is to identify material triggers: a new legal version, a changed purpose, a different dataset or provider, an altered model, movement to a new lifecycle stage, evidence that has expired, or a significant objection that the earlier review did not address. Metadata and workflow tools can signal some of these events. The administration must determine their significance, document whether a prior decision remains adequate, and allocate the work proportionately.

Within this thesis, data stewardship is the central organisational continuity function. It is not treated as a new statutory office, a universal job description, or a source of legal decision-making authority. The function may be assigned to one or more data stewards, distributed through an institutional network, or embedded in existing organisational roles. It keeps data descriptions, provenance, quality assertions, use conditions, evidence, ownership, and review states aligned; identifies missing or stale information; and routes material questions to the competent actor [16, 17, 18]. The need for this function already arises where data,

infrastructure, evidence, and decision authority are distributed across organisations. This thesis argues that its importance will increase as Common European Data Spaces move into operational use. Legal counsel, a Data Protection Officer, a domain expert, or an authorised manager retains their respective decision authority. The stewardship function maintains connections among their decisions and the technical records on which later action depends.

The GDPR provides a relevant but limited precedent. Accountability and data protection by design require continuing organisational arrangements, documentation, and responsibility rather than a legal review undertaken only after technical development [3]. Studies of implementation show that this work can require changes to systems, routines, expertise, and cooperation, with burdens and practices varying by organisation [19, 20]. The present thesis extends the organisational insight, not the doctrinal reach of the GDPR: CEDS and public-sector AI involve personal-data questions in some cases, but also concern access institutions, intellectual property, dataset quality, sectoral regulation, public authority, and shared infrastructures.

1.4 Three complementary capacities

The thesis organises the problem through three complementary capacities. The first concerns the strategic and infrastructural conditions under which European actors can exercise meaningful choices over data and AI. The second concerns the organisational ability of a public administration to deliberate across professional boundaries and revisit decisions during a project. The third concerns the need for a structured, source-linked method that translates selected legal concepts into metadata needs and examines whether and how existing standards may represent them, while preserving possible gaps and questions reserved for competent judgment. These are levels of analysis, not stages of an implemented product.

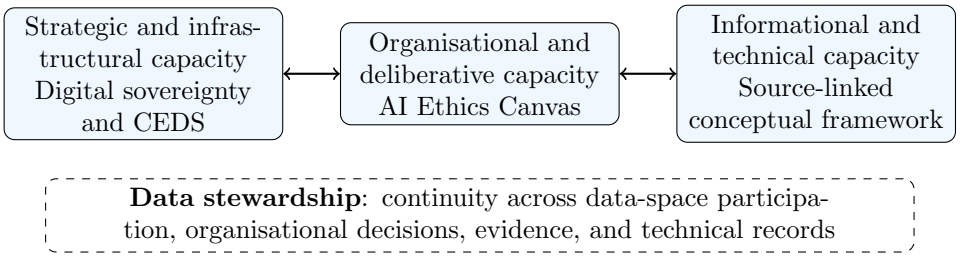


Figure 1.1: The three complementary levels and the cross-cutting continuity function connected by the thesis. Data stewardship maintains relations across data-space participation, organisational decisions, evidence, and technical records; the diagram does not represent a validated end-to-end system.

1.4.1 Strategic and infrastructural capacity

Digital sovereignty provides the macro-level context. The term is contested because sovereignty may be attributed to individuals, organisations, states, or the European Union, and because the object of control may be personal data, industrial data, platforms, software, networks, cloud services, or computing capacity. Absolute autonomy is neither a realistic description of interconnected digital systems nor the objective examined here. The relevant issue is whether formally recognised choices can be exercised through available institutions and infrastructures [21, 22].

The comparative legal analysis examines this issue through European and Chinese approaches to data and AI regulation. It finds that rights and regulatory ambition may coexist with limited individual power when material alternatives are absent [9]. The thesis extends that analysis by asking what CEDS contribute to Europe's capacity over data in the AI economy. Data spaces may improve discoverability, interoperability, governance, and reuse within European arrangements. They do not manufacture advanced chips, create cloud capacity, eliminate platform dependence, or guarantee that public administrations possess the people and processes needed to use the infrastructure well. CEDS are therefore treated as a partial response within a wider industrial and institutional strategy, not as evidence that digital sovereignty has been achieved.

A separate update brings the analysis to 2026. The AI Continent Action Plan and the Data Union Strategy connect data availability with computing infrastructure, skills, adoption, regulatory simplification, and the scaling of data spaces [23, 24]. These documents establish policy direction and reported initiatives. They do not by themselves demonstrate outcomes, independence, or operational maturity. This evidence boundary is important to the thesis's general argument: an announced capacity, a legal entitlement, and a capacity exercised in practice are different claims.

1.4.2 Organisational and deliberative capacity

The AI Ethics Canvas moves the analysis inside the public administration. AI projects distribute knowledge across people who may not share a common method for identifying ethical, legal, technical, and societal concerns. A lawyer may not know which data transformation is planned; a developer may not know which public-law justification or consultation process is required; a domain officer may understand the service but not the limitations of a model; an affected community may recognise harms that are invisible in the project specification. The organisational question is how to bring these perspectives into a sufficiently structured conversation and how to revisit it as the project changes.

The AI Ethics Canvas addresses this organisational capacity through an artifact grounded in action research and value-sensitive design that organises questions across three participant roles and five areas. Twelve interviews and a facilitated workshop provide direct formative

evidence about its development and use. The workshop involved 18 participants, two broader project cases, and three scenarios. Survey and focus-group material is used for a different evidentiary purpose: it helps explain the context and needs from which the work emerged, but is not merged with the interviews and workshop as if all observations measured the same outcome.

The thesis assigns the Canvas three possible moments of use: before development begins, iteratively when a material choice changes, and at a final review before a project crosses an institutional threshold. This lifecycle positioning is a thesis-level proposal derived from the artifact's deliberative function. It has not been validated through a longitudinal deployment, and "final review" does not mean certification. The Canvas release gives teams a reusable shared object; the decision to proceed remains with the competent administration and must draw on legal, technical, procurement, security, data-protection, domain, and democratic processes appropriate to the case.

1.4.3 Informational and technical capacity

The Canvas creates touchpoints at which technical and non-technical participants can identify purposes, assumptions, data sources, responsibilities, evidence, and unresolved questions. When recorded, this information may enrich project documentation and metadata and support subsequent compliance and assessment activities. These touchpoints do not prescribe a metadata schema or make the conceptual framework a necessary continuation of a Canvas session. Metadata standards remain relevant because CEDS connect organisations that retain their own systems and responsibilities and therefore need shared ways to identify resources and communicate descriptions, provenance, quality, policies, and other assertions across boundaries [4, 25].

The informational contribution develops the conceptual framework as a deliberately narrow translation method. It begins with the authoritative legal source rather than with the available vocabulary. A selected source passage is reconstructed as a source-linked legal concept or requirement; the information necessary to communicate or support that concept is expressed as a metadata need; the need is compared with exact, versioned elements of a standard; and the result records what is represented and what is not. The chain is:

$$\begin{aligned}\text{legal source} &\longrightarrow \text{source-linked legal concept or requirement} \\ &\longrightarrow \text{metadata need} \longrightarrow \text{exact standard element} \longrightarrow \text{representational outcome.}\end{aligned} \tag{1.1}$$

The missing connection is not a direct equivalence between a legal requirement and a metadata field. The five-stage framework provides a structured, source-linked, and revisable method for relating the reviewed requirement, the information needed to assess it, the standard-based representation selected for that information, and the supporting evidence

available in the concrete configuration. Standards define how information may be represented or structurally checked; they are not, by themselves, evidence that an assertion is true or that a requirement has been satisfied. Data stewardship maintains these relations over time, while the competent legal, technical, or institutional actor evaluates the evidence and determines its significance.

Responsibility, evidence, versioning, and reassessment apply throughout this chain. They are not additional layers and do not turn it into a general compliance architecture. The representational outcome may show direct support by one element, composite support through several elements within or across standards, partial or evidential support, a gap requiring profiling or extension, or a judgment that should not be represented as metadata. DCAT-AP, ODRL, PROV-O, DQV, DPV, Croissant, and related specifications offer complementary semantics for different parts of the problem [26, 27, 28, 29, 30, 31]. Their existence does not prove the truth of a statement, the sufficiency of evidence, the applicability of a rule, or compliance.

The relationship among the three capacities is reciprocal. CEDS can provide a governed environment in which a public administration finds and accesses data, but the administration needs a deliberative process to formulate a defensible project. The Canvas can expose a need to document provenance, quality, purpose, authority, restrictions, or affected interests, but a separate source-linked analysis is required before any of those concepts are represented. Metadata can preserve selected assertions and reveal change, but the organisation must interpret the signal and decide what follows. The thesis connects these functions while resisting the stronger assertion that their combination has already been implemented or validated.

1.5 Scope, method, and evidence boundaries

The thesis is a monograph partially based on publications. This format preserves the questions, methods, co-authorship, and evidentiary boundaries of the incorporated work while the introduction and conclusion construct a single thesis-level argument. Chapters 3 and 4 begin with disclosures identifying the relevant publications and their provenance. New 2026 updates and thesis-specific extensions are distinguished in the chapter text. The disclosures prevent the monographic synthesis from obscuring which claims were produced with co-authors and which are advanced here for the first time.

The substantive chapters do not share one evidentiary design. The digital-sovereignty chapter uses comparative legal and policy analysis to identify a capacity problem. Its 2026 extension uses official policy material to update the European context, while treating policy announcements as evidence of institutional direction rather than demonstrated effects. The Canvas chapter uses action research, value-sensitive design, interviews, and a workshop to report formative artifact development. Chapter 5 makes a conceptual contribution: it de-

velops a five-stage, source-linked framework informed by a bounded examination of selected legal and technical materials. The conclusion may interpret relationships among these contributions, but it cannot make one chapter's evidence validate another chapter's method or artifact.

The conceptual framework concerns selected EU rules relevant to data-space participation, access, reuse, and downstream operations. Data Act Article 33(1)(a) provides a bounded application because it states an explicit relationship among dataset description, restrictions, licences, collection methodology, quality, and uncertainty [32]. Two aspects are followed through the five stages: a collection-methodology relation that can be represented through exact provenance elements, and descriptive sufficiency, which remains a non-metadata judgment. The EHDS provides a shorter sectoral illustration of why the analysis must be repeated when the relevant actors, objects, procedures, and legal effects change [1]. These illustrations make the intended operation and boundaries of the framework visible; they do not establish its completeness or measure how much European regulation is represented through metadata.

At the technical stage, the unit of comparison is an exact element in a dated or versioned specification rather than a standard named in the abstract. This is a methodological requirement of the framework. The published semantics of the element are examined against the identified information need, and both its possible contribution and its limitation remain explicit. A provenance property does not establish the lawfulness of the recorded activity; a policy expression does not establish the issuer's authority; and a quality assertion does not prove the adequacy of the underlying method. The framework therefore examines possible representational fit without converting standard use, structural conformance, or supporting evidence into a legal conclusion.

The substantive legal analysis is fixed at 1 April 2026. Retrieval and subsequent verification dates are recorded separately and do not redefine this legal cut-off. Later legal developments fall outside that substantive baseline. Technical specifications are considered in the exact versions identified in the source inventory, which may have different publication and retrieval dates. Dates of entry into force and application remain provision-specific and are distinguished from the analytical cut-off.

Several matters remain outside scope. The thesis does not offer legal advice for a particular administration, reconstruct all Member-State implementing law, test a production CEDS transaction, audit a deployed AI system, or compare the technical performance of models. It does not establish the external legal validity of each authorial proposition or the organisational effectiveness of the proposed connections. The Canvas's formative evaluation does not validate the conceptual framework; the conceptual framework does not validate the Canvas; and a CEDS policy programme does not prove that either tool will be adopted. These exclusions make the central contribution narrower but testable.

1.6 Research questions

The main research question is:

> How can the complexity of EU data and AI regulation be translated into organisational and metadata-based tools that support public administrations engaging with Common European Data Spaces?

Three subquestions organise the inquiry:

1. How does the relationship between digital sovereignty, material capacity, and Common European Data Spaces shape the conditions under which European public administrations can exercise meaningful agency over data and AI infrastructures?
2. What role can the AI Ethics Canvas play in the initial design, iterative reconsideration, and final review of public-sector AI projects, and what are the limits of the evidence supporting that role?
3. How can selected EU legal concepts and requirements be translated into source-linked metadata needs, compared with exact elements of existing standards, and classified without presenting representation as legal compliance?

The verb "support" controls the interpretation of all four questions. Support may mean making a dependency visible, structuring a multidisciplinary discussion, retaining evidence, carrying an assertion between systems, performing a bounded validation, or identifying that human judgment remains necessary. It does not mean replacing the legal, democratic, managerial, or professional authority of the public administration. The questions ask how regulatory complexity can be made more operational and traceable, not how compliance can be automated.

1.7 Contribution of the Thesis

The original contribution of this thesis lies in connecting three levels at which European data and AI regulation must become operational. The legal analysis of digital sovereignty identifies the gap between regulatory ambition and the material capacity required to exercise meaningful control over data and infrastructures. The AI Ethics Canvas addresses the organisational level by giving public administrations an iterative setting in which technical and non-technical actors can formulate and revisit ethical, legal, and social questions throughout an AI project. The conceptual framework addresses the informational level by proposing a structured, source-linked method for translating selected EU legal concepts into metadata needs and examining their possible representation through exact elements of existing standards. Taken together, the thesis argues that human-centred compliance in Common

European Data Spaces depends on the interaction of institutional capacity, structured deliberation, organisational continuity, and technical traceability; none of these elements, taken alone, establishes legal compliance.

Across the organisational and informational levels, the thesis identifies data stewardship as the central continuity function. Federated data arrangements distribute datasets, metadata, evidence, infrastructure, and decision authority across participants and over time. Stewardship maintains the connections among the governed object, applicable conditions, supporting evidence, responsible contributions, decisions, versions, and review states. The function may be assigned to one or more data stewards, distributed through a network, or embedded in existing roles; it does not create a statutory office or transfer legal and institutional authority. Its importance already follows from the coordination required in distributed data practices and is expected to increase as Common European Data Spaces move into operational use.

The second cross-cutting contribution is a source-linked, maintainable, and revisable connection between reviewed legal requirements, standard-based technical representations, and supporting evidence. This is not an equivalence between a legal requirement and a metadata field. Standards define how information may be represented or structurally checked; evidence supports the resulting assertions; and competent actors determine legal applicability, evidential sufficiency, and institutional consequence. The two bounded Data Act micro-applications make these transitions and their boundaries inspectable without establishing the framework's completeness, effectiveness, or representational coverage. The comparative legal analysis, formative Canvas evidence, and conceptual framework therefore retain their separate evidentiary bases. Future research may examine selected interactions among them without presupposing a single integrated workflow.

Reproducibility is pursued in ways appropriate to each component. Publication disclosures identify the provenance and co-authorship of the incorporated scholarship. The Canvas repository provides the canonical version 1.0.0 under the CC BY-SA 4.0 licence [33]. For the conceptual framework, precise references to the selected legal sources and technical specifications, an explicit five-stage method, and the bounded Data Act application make its analytical transitions open to inspection and challenge. Reproducibility does not turn a legal interpretation into an official reading or a reusable canvas into proof of organisational effectiveness.

1.8 Structure of the Thesis

The six chapters follow the movement from a European capacity problem to organisational deliberation and then to informational traceability. The present chapter introduces the dataset journey, defines human-centred compliance, states the research questions, and explains the partially publication-based design. Chapter 2, entitled *State of the Art*, reviews

the intersecting work on digital sovereignty and data spaces, downstream legal authority, public-sector AI, ethics tools, metadata standards, traceability, technical enforcement, and meaningful human involvement. It identifies a shared distance between legal norms, organisational decisions, and technical representations.

Chapter 3, *Digital Sovereignty, Regulatory Capacity, and Common European Data Spaces*, presents a comparison of European and Chinese approaches through comparative legal analysis and states its bounded result. A separate thesis-specific section adds a 2026 update on data, cloud, compute, and European AI policy. The final sections examine CEDS as a partial response to the capacity deficit and explain why infrastructures still depend on administrative ability to organise responsible projects.

Chapter 4, *The AI Ethics Canvas: Designing and Reviewing Public-Sector AI*, presents the problem, design method, artifact, and formative evidence. It distinguishes interviews and workshop evidence from surveys and focus groups used to establish context. A thesis-specific extension positions the Canvas at initial, iterative, and final-review moments and identifies information and evidence touchpoints that may support later project review and metadata enrichment.

Chapter 5, *From EU Legal Concepts to Metadata Standards*, develops the source-linked five-stage framework connecting an authoritative legal source, a reviewed legal concept or requirement, a metadata need, an exact standard element, and a representational outcome. Responsibility, evidence, versioning, and reassessment remain transversal, while data stewardship provides the organisational continuity needed to maintain the resulting relations over time. Two bounded micro-applications to Data Act Article 33(1)(a) reach, respectively, exact provenance elements and a judgment that should not be reduced to metadata; the EHDS illustrates why the method must be reapplied in a sector-specific setting. A governed operational corpus, an application profile, and an assessment-supporting knowledge graph are presented as possible developments of the framework rather than as completed results of the thesis.

Chapter 6, *Synthesis and Conclusions*, answers the main question and the three sub-questions, distinguishes demonstrated findings from possible interactions, and returns to the dataset journey. It identifies directions for empirical evaluation of the contributions, separately or in combination, in concrete public-administration settings.

The dataset journey remains the narrative thread because it exposes three kinds of capacity that must cooperate. CEDS create conditions in which public administrations can supply, govern, discover, and access data; the administration must determine which project or operation should be pursued and under which responsibilities; and metadata can preserve selected descriptions, evidence, and decisions. Human-centred compliance lies in maintaining these relations through data stewardship without confusing availability with authority, deliberation with approval, or representation with legality.

Chapter 2

State of the Art

Chapter 1 distinguishes a governed route to data from the authority to perform a particular downstream operation (e.g., AI training). The illustrative health-data scenario makes the remaining problem concrete: law, metadata standards, data-space architectures, and organisational governance each illuminate part of the route, but none connects the full progression from a current legal source to an attributable decision, a technical response, supporting evidence, and later reassessment. This gap frames the review, which considers what each body of work contributes and why the relations among them remain unresolved.

The analysis begins with CEDS because data access reveals how formal rights depend on material, organisational, and infrastructural capacity. From there, it follows the downstream operation through conditional legal applicability, documentation, metadata, policy representation, and technical enforcement before turning to the people and institutions that interpret evidence, maintain governance records, and remain responsible for decisions. Across these bodies of work, the same integration gap emerges: the necessary components exist, but a source-linked governance and traceability structure is still needed to coordinate them without collapsing legal judgment into technical conformance.

2.1 Review Method

The research problem crosses law, data-space research, software engineering, standardisation, and human–computer interaction. Because these fields use different concepts, methods, and forms of authority, the evidence base was constructed as a structured scoping review rather than presented as a systematic or exhaustive review. Its purpose is to identify the principal approaches, test what each can support, and expose the unresolved points of integration that the thesis must still address.

The review combined scholarly databases, publisher and institutional repositories, official EU sources, standards bodies, and technical documentation. Five core searches covered

CEDS and AI, data-space governance, AI-training-data governance, legal requirements engineering, and human oversight; three supplementary searches examined implementation barriers, automated-compliance limits, and overreliance on human reviewers. Backward and forward citation tracing added related, competing, and critical work. A final search-and-screening audit updated the literature and status records.

The database audit screened bounded, relevance-ordered samples: 97 unique Scopus records and 95 unique OpenAlex records, with at least seven detected overlaps, yielding 185 unique records across the two samples. These are audit denominators rather than a claim that complete result sets were screened or that retaining a record at the title-and-abstract stage made it evidence for the chapter. OpenAlex and Scopus expose different searchable fields and ranking logic. Coverage is therefore limited by the bounded samples and the differences between the databases.

Screening first identified contributions relevant to the research strands and then verified the authoritative record and the passages needed for each claim. Superseded versions were replaced by authoritative ones; insufficiently described records remained uncertain. Sources were used only within their evidentiary function: primary law establishes enacted text and status; scholarship supports interpretation, critique, and empirical findings; standards establish their own semantics and conformance models; and architecture or product documentation establishes declared scope and capability. Claims of practical effectiveness therefore require separate empirical evidence.

Table 2.1 reports the historical retrieval yields and bounded samples recorded in the database audit of 17 July 2026. No date, subject-area, document-type, or language filter was applied. Scopus searches used the title, abstract, and keyword fields; OpenAlex used its indexed search fields and query interpretation. Both samples followed relevance ordering, but their searchable fields and ranking procedures differ. The retrieval yields therefore describe the searches performed and should not be treated as equivalent measures of coverage.

Table 2.1: Database retrieval yields and captured screening samples, 17 July 2026

Search group	**Scopus**		**OpenAlex**	
	Returned	**Sampled**	**Returned**	**Sampled**
CEDS and AI	11	11	548	19
Data-space governance	354	20	13,724	18
AI-training-data governance	45	20	4,882	18
Legal requirements engineering	20	20	2,206	18
Human oversight, refined query	347	20	25,767	18
Implementation barriers	672	3	26,843	3
Automated-compliance limits	137	3	18,840	3
Oversight overreliance and failures	1,801	3	66,794	3
Captured sample before deduplication	—	**100**	—	**100**

Scopus sampling covered the first twenty relevance-ordered records for each core query, or the complete result set where smaller, and the first three records for each contrary-evidence query. OpenAlex allocated nineteen records to the first core query, eighteen to each remaining core query, and three to each contrary-evidence query. Within database deduplication left 97 unique Scopus records and 95 unique OpenAlex records. Matching DOIs and titles identified at least seven cross-database overlaps, giving a recorded combined sample of 185 records after the detected overlaps were removed.

The complete database-specific search strings, eligibility criteria, aggregate retrieval yields, screening procedure, and deduplication method are reported in Appendix A. Record-level screening logs are not included in the thesis.

2.2 Data Spaces and Downstream Authority

The European data-space programme should be understood as a set of governed and federated arrangements rather than as a single database, marketplace, platform, or legal object. The European Strategy for Data linked common European data spaces to the development of a single market for data, with particular emphasis on the governance of data access, use, and value creation [6]. The Commission's current programme extends this work across a range of sectors and domains [7]. These policy sources indicate the programme's ambition and institutional direction, but they neither establish a general right of access nor demonstrate that every proposed data space is operational.

The term also has an earlier technical genealogy. Franklin et al. [34] proposed dataspaces as an abstraction for managing numerous heterogeneous and interrelated sources without requiring complete integration in advance, while Halevy et al. [35] developed the corresponding idea of gradual, "pay-as-you-go" integration in dataspace support systems. This lineage explains the emphasis on coexistence, discovery, relationships, and incremental integration. It does not by itself supply the legal, institutional, or sectoral governance that distinguishes the contemporary European data-space programme.

The academic literature helps to clarify the Commission's strategy. Scerri et al. [4] describe data spaces through resources, services, contracts, ontologies, governance, and provenance, while identifying legal uncertainty and purpose control as persistent challenges. The systematic survey by Bacco et al. [5] treats data spaces as distributed, data-centred arrangements organised around governed transactions and reports continuing interoperability and complexity barriers. Together, these sources support a functional definition: a data space combines participants, governance rules, catalogues, identity and trust services, interoperability arrangements, contractual or policy instruments, and technical means for exchange or processing.

These terms overlap, but they do not describe the same object. A platform refers mainly to the technical or organisational structure through which participants interact. A mar-

ketplace is more specific, since it also supports functions such as matching, transactions, or payments. An ecosystem is broader and includes the actors, resources, institutions, and services that develop around data sharing. A data-intermediation service under the Data Governance Act (DGA) is different again, because it identifies a legally defined role rather than a general model of organisation.

The concept of a commons adds a further normative dimension. It usually implies some form of collective stewardship, participation, reciprocity, or orientation towards a common benefit [10]. Related work on digital commons widens the object of governance to the communities and labour that create and maintain resources, their licences and reuse preferences, the infrastructures on which access depends, and associated economic and environmental costs [14]. These considerations help distinguish a commons from other forms of data organisation and caution against treating catalogue availability as a complete account of reuse. The paper develops an interdisciplinary agenda and open questions; it neither creates legal obligations nor demonstrates that a particular CEDS governance model or metadata mapping is effective.

The meaning of "data sovereignty" also depends on the claimant, value, and institutional context [21]. Ryan et al. [22] distinguish individual, organisational, and state claims within European data-space policy and reject an image of absolute, unitary control. In our comparative analysis of EU and Chinese regulation, we further argue that formal rights or declarations of control may coexist with limited material, organisational, or infrastructural capacity [9]. Applied to CEDS, these contributions require sovereignty claims to identify who exercises control, over which object, for which value, through which capability, and within which institutional setting. They provide a conceptual discipline for analysing CEDS rather than evidence of their operational performance.

2.2.1 Sectoral implementation initiatives

Sectoral initiatives give the data-space programme different institutional and technical forms [7]. The following examples distinguish preparatory work, deployment objectives, and reported components; they do not establish uniform operational maturity or independently evaluated effectiveness.

In agriculture, AgriDataSpace delivered a farmer-centred governance and technical-architecture roadmap for the Common European Agricultural Data Space [36]. In mobility, deployEMDS supports local implementation projects and common building blocks for sharing transport data and developing services across cities and regions [37]. In manufacturing, SM4RTENANCE addresses maintenance services through federated asset data, including reported integration work connecting Asset Administration Shell digital twins with the Eclipse Dataspace Connector [38]. These initiatives link interoperability to sector-specific participants, purposes, and services.

Other domains foreground different governance requirements. The HealthData@EU Pilot

informed recommendations for health data access bodies, metadata catalogues, and application procedures for secondary use [39]. In cultural heritage, Europeana's 2024–2025 implementation summary reports identifier and API developments supporting reuse of digitised collections [40]. For public administration, the Public Procurement Data Space describes ontology-based integration of European and national procurement sources alongside analytics services [41]. Together, these examples show how discovery, access governance, semantic integration, and analysis are combined in different institutional settings.

Energy provides a further set of implementation examples. ENERSHARE, OMEGA-X, DATA CELLAR, EDDIE, and SYNERGIES explore interoperable architectures, metering interfaces, forecasting, and digital twins [42, 43, 44, 45, 46]. They were funded under Horizon Europe topic `HORIZON-CL5-2021-D3-01-01` [47]. INSIEME's subsequent Digital Europe vision addresses the integration of existing energy infrastructures through shared governance and interoperability arrangements [48]. The project reports distinguish released or beta components from functions still envisaged.

Data-space functions consequently extend beyond exchange. The International Data Spaces Association (IDSA) Rulebook includes participants offering algorithms, code, or specialised processing infrastructure under contractual and usage-policy arrangements [49]. Such arrangements can support governed processing, derived data products, analytics, and AI-related services. Their availability and performance depend on the implementation and resources provided.

Technical specifications and coordination activities support this work. The Dataspace Protocol, initially developed under IDSA governance and subsequently under the Eclipse Dataspace Working Group, provides publication, agreement-negotiation, and access protocols in its stable release 2025-1 [50]. BRIDGE's October 2025 DERA version 3.1 report documents standards mapping and interoperability tests between EDDIE and OMEGA-X [51, pp. 7–9, 18–20]. These contributions support technical coordination without themselves establishing substantive legal compliance.

AI factories supply a distinct capacity: access to AI-optimised supercomputing and related support [52]. Data spaces may incorporate processing resources without supplying all the capacity required for a particular AI workload. For public administrations, the sectoral comparison therefore reinforces the need to identify the actors, operations, input and output conditions, and evidence required for reassessment. Neither an announced infrastructure nor a reported component establishes the authority or adequacy of a particular downstream use.

The literature and implementation records consequently support a bounded account of data-space capability. Data spaces can make identities, offers, policies, permissions, provenance, and transfers visible, and can organise defined exchanges and processing services among participants. Their governance and technical services do not convert "trusted", "sovereign", or "compliant" access into a legal status that automatically follows the data. The next question is therefore whether the relevant access route also authorises the proposed

downstream operation, including processing performed within or through the data space.

2.3 Access and Downstream Use

For analytical purposes, this chapter distinguishes at least six stages: technical access; admission under a governance rulebook; an instrument-specific statutory access or reuse permission; contractual or licence-based permission; lawful access for the relevant transaction; and authority for the downstream processing operation. These stages identify different grounds on which data may become available and subsequently be used. A participant may have the credentials required to retrieve data and satisfy the conditions of a data-space rulebook, yet still lack a statutory, contractual, or other legal basis for a particular act of reuse. Conversely, lawful access to a resource does not settle every legal question raised by copying, combining, retaining, or using it to train an AI system. The distinction is therefore analytical rather than necessarily sequential: a particular arrangement may combine several stages, but evidence at one stage does not substitute for the assessment required at another.

The GDPR provides a clear example of why lawfulness must be assessed in relation to the specific operation under examination. The analysis begins with the processing itself. Where personal data are used to train an AI system, the controller must first comply with the principles set out in Article 5 and identify an appropriate legal basis under Article 6. Where the training data include special categories of personal data, Article 9 introduces a further layer of conditions. The assessment cannot therefore stop at the general purpose of developing an AI model; it must consider the type of data used, the purpose pursued, the actors involved, and the way in which the processing is carried out.

The same reasoning applies when data collected for one purpose are later reused for training. In that case, Article 6(4) may require an assessment of whether the new purpose is compatible with the original one. Article 89 does not remove this requirement by creating a general permission for research. Rather, it establishes safeguards and, in defined circumstances, structures for derogation. The mere description of an activity as research is therefore insufficient to establish its lawfulness [3, 53].

EDPB Opinion 28/2024 confirms that these questions remain dependent on the facts and on the stage of the AI lifecycle. Whether a model can be considered anonymous must be assessed in relation to the particular model and the means reasonably available to identify individuals. Where legitimate interests are relied upon, the analysis must still distinguish the purpose pursued, the necessity of the processing, and the balance between the interests of the controller and the rights of the data subjects. The Opinion also makes clear that a later actor cannot simply disregard unlawfulness arising during an earlier stage of development [54]. Against this background, scholarship suggesting that legitimate interests may support certain AI-training operations should be understood as identifying a possible legal basis in defined configurations, rather than as establishing a general basis for AI training [55].

The DGA regulates specified protected public-sector reuse, intermediation services, and data altruism, but expressly preserves data-protection law and does not provide a general personal-data legal basis [56]. The Data Act likewise operates within specified connected-product, related-service, actor, and relationship contexts. It contains data-space interoperability provisions, but those provisions do not convert the Regulation into a horizontal CEDS access statute or authorise unrelated downstream processing [32]. These instruments are important precisely because their boundaries show why treating "data access law" as one cumulative block automatically attached to every data-space transaction would be misleading.

Copyright text-and-data-mining rules make the distinction equally visible. DSM Directive Article 3 requires Member States to provide an exception for reproductions and extractions made by research organisations and cultural heritage institutions to conduct TDM for scientific research on works or other subject matter to which they have lawful access. Its beneficiary, purpose, operation, and access conditions are cumulative. Copies made under that route must be stored with an appropriate level of security and may be retained for scientific research, including verification, while rightholder measures for network or database security and integrity may not go beyond what is necessary. Article 7(1) separately makes contractual provisions contrary to the Article 3 exception unenforceable. These elements describe a bounded TDM route rather than a general research or AI-training permission [57, 58].

Article 4 provides a different route. It covers TDM reproductions and extractions of lawfully accessible works or other subject matter without limiting the beneficiary to the Article 3 categories, but applies only where the relevant use has not been expressly reserved by the rightholder in an appropriate manner. For content made publicly available online, the Directive identifies machine-readable means as the relevant example. Article 4 copies may be retained only for as long as necessary for TDM, and Article 4 does not affect Article 3. The two routes must remain separate: an Article 4 reservation cannot silently be imported into Article 3, and failure to qualify for Article 3 does not itself establish Article 4 applicability.

Neither route resolves the broader lawfulness of AI training. The proposed operation must first qualify as TDM and involve a reproduction or extraction governed by one of the referenced exclusive rights; mere facts or unprotected data do not require copyright authorisation merely because they are processed computationally. Conversely, lawful access supplies a condition without resolving every protected act, transformation, retention period, disclosure, or other applicable regime. A machine-readable reservation likewise records an attributable assertion without itself proving rightholder identity, scope, appropriate form, temporal effect, or enforceability. Margoni and Kretschmer [59] identify continuing difficulties concerning lawful access, opt-outs, harmonisation, and incentives. The relevant protected subject matter, rights, acts, Member-State implementation, and any authoritative judicial interpretation must be fixed before a concrete national conclusion can be drawn.

Taken together, these sources establish the chapter's principal distinction: access is a necessary gate, but authority for downstream use remains a separate assessment. Conditional applicability and role allocation follow from that distinction.

2.4 Conditional Applicability

The applicable regulatory framework cannot be determined by listing relevant instruments in advance. Each instrument defines its own scope through particular objects, actors, purposes, exclusions, territorial conditions, and dates of application. The analysis must begin by testing the facts against the conditions of each regime separately. Only after those assessments have been completed is it possible to conclude that several regimes apply to the same operation. Their simultaneous application should be treated as the outcome of the analysis, not as its starting assumption. Even where the resulting obligations overlap or create implementation difficulties, this does not necessarily amount to a legal conflict. Such a conflict arises only where two applicable rules cannot both be satisfied under the same factual circumstances.

The AI Act provides a central example of why data-governance obligations must be derived from the classification and intended purpose of the system under examination. Article 10 principally concerns high-risk AI systems that use techniques involving the training of models with data. For those systems, training, validation, and testing datasets are subject to governance and quality requirements defined in relation to the system's intended purpose. The provision has a narrower application where a high-risk system does not use such training techniques, since Article 10(6) applies the relevant requirements only to its testing datasets. Article 10 should therefore not be treated as a general code governing every dataset used for AI development. Its application first requires the system to fall within the high-risk category and then depends on the techniques and datasets used in its development [2].

The same conditionality prevents the label "research" from operating as an automatic exemption. Our analysis of the AI Act's research provisions shows that the relevant result can depend on the activity's purpose, whether a system is placed on the market or put into service, the actor's role, downstream modification, and the lifecycle stage at which the assessment is made [60, 61]. For the running scenario, a public research organisation must record more than its institutional identity. The intended operation, development and deployment pathway, system classification, responsible actors, and events that may reopen the assessment must be recorded and kept connected to the project as it develops.

Earlier literature remains useful in identifying the main concerns associated with AI training data, including data quality, personal-data protection, copyright, discrimination, and innovation. Since some of this work examined the Commission's proposal, its conclusions should be read alongside the changes introduced during the legislative process and the final text of the Regulation [62]. More recent analysis of the relationship between the AI Act and

the GDPR similarly shows that the two instruments overlap at particular points, including Article 10, but do not form a single set of requirements that applies unchanged to every AI operation [63]. The applicability of each regime must instead be assessed separately before their combined effect can be established.

The conceptual framework should accordingly treat classification as a variable rather than as a permanent characteristic of the technology. For a defined system and operation, the assessment may conclude that the system is high-risk, not high-risk, not yet determined, or subject to a contested classification. The last two descriptions are analytical states rather than terms used by the AI Act. Reassessment becomes necessary where the intended purpose, integration context, or other facts relevant to Article 6 change. In particular, Article 25 recognises that modifying the intended purpose of a system already placed on the market or put into service may cause a previously non-high-risk system, including one based on general-purpose AI, to become high-risk [2].

General-purpose AI models are excluded from the principal analysis because the data-space configurations examined here concern the reuse of data for research and other defined downstream tasks. The relevant legal and governance questions therefore arise mainly in relation to context-specific processing operations, identified purposes, and specialised systems developed or deployed within a particular domain. GPAI models raise a different set of questions: they are developed for broader and less determinate uses and are governed through a distinct model-level framework under Chapter V of the AI Act. A full treatment would require separate consideration of provider obligations, downstream information duties (including the Commission template for public summaries of training content), copyright compliance, systemic-risk classification, and the applicable codes of practice [2, 64], extending the inquiry beyond its present focus on operation-specific data governance. This exclusion does not preclude a GPAI model from forming part of a downstream system considered in the analysis. In that configuration, the assessment remains directed to the intended purpose, integration context, and regulatory classification of the resulting system.

Regulatory and organisational role vocabularies also differ. One participant can simultaneously be a data-space provider or consumer, a GDPR controller or processor, a contractual licensee, an applicant under a sectoral access regime, and an AI Act provider or deployer. Each role attaches to a particular instrument, operation, object, and period, which prevents treating the roles as interchangeable. A responsible representation should therefore record the asserted role together with the actor, operation, object, source, rationale, reviewer, uncertainty, and validity interval.

Instrument-defined and contractual roles should remain distinct from the organisational functions through which an assessment is carried out and maintained. A statutory or contractual role identifies the addressee of a duty or the party to an agreement under a particular source. It does not determine who, within or across the participating organisations, approves a purpose, accepts residual uncertainty, or authorises a transition from one stage of the data

or system lifecycle to another [65, 66, 67].

In this context, data stewardship provides the continuing organisational link between legal and governance determinations and the data practices to which they relate. This understanding is consistent with the Open Data Institute's account of responsible data stewardship as an iterative and systemic practice that must respond to changing contexts, consequences, and mitigation measures [16]. The concept of stewardship identifies an organisational function rather than a statutory role; final decision-making authority depends on a separate allocation. Through this function, organisations coordinate how they define and maintain data purposes, conditions of access and reuse, quality information, provenance, restrictions, responsibilities, and associated evidence throughout the lifecycle. The function helps to keep these elements aligned as data are prepared, shared, reused, combined, or transferred to a downstream process, and to ensure that material changes are returned for assessment rather than absorbed as routine technical modifications.

Recent governance models make this function broader than conventional data administration. Data spaces, public data trusts, collaborative research infrastructures, and human-centred public-service ecosystems distribute control, use, and responsibility across several organisations. O'Hara's analysis of data trusts is instructive because it separates law, ethics, architecture, and governance: a trust arrangement may make actors and processing more accountable, but cannot make an unlawful operation lawful or replace the authority of the relevant legal roles [68]. The World Economic Forum's Helsinki-oriented account likewise treats public-sector data governance as a relationship among institutions, residents, communities, and service providers requiring oversight, transparency, accountability, interoperability, and attention to affected interests [69]. In such settings, stewardship is the operational capacity that preserves those relationships in the data record rather than a generic claim that the organisation is trustworthy.

Stewardship matters in federated public-sector settings because data, decision authority, infrastructure, and consequences are distributed across organisations. OECD evidence presents it as a dedicated role, an institutional network, or a capability diffused among public officials, while distinguishing strategic and technical responsibilities [17]. In the European B2G context, a Commission-appointed expert group recommended a recognised steward function to coordinate data collaborations and connect external partners with internal quality, metadata, legal, and ethical processes [70]. This policy recommendation was not a general legal obligation, but supports the narrower proposition that federated exchange needs identifiable continuity and coordination.

Data governance literature locates governance in the allocation of decision rights, accountabilities, policies, standards, and monitoring, while distinguishing those choices from their day-to-day execution [71]. A recent public-sector operating model makes the allocation more concrete: data owners remain accountable for major decisions, stewards carry out operational governance and quality work, and custodians implement technical requirements

in the systems they operate [18]. For the purposes of this thesis, this separation is adopted analytically, without importing the United Kingdom model as a universal taxonomy or allowing an internal assignment to create legal competence. Strategic authority concerns the approval of purposes, risk positions, exceptions, and lifecycle transitions. Technical custody concerns the operation of infrastructure, access controls, storage, and other implementation measures, without necessarily including legal or governance authority. Stewardship sits between these functions: it maintains the governance record, coordinates the actors responsible for its different parts, and identifies when a change in data, purpose, context, or recipient requires escalation. This allocation is consistent with organisational research showing that responsibility must be embedded in workflows and matched by suitable authority, resources, coordination, and escalation capacity if it is to be effective [65, 66, 67].

Stewardship should therefore be treated as a resourced organisational capability, not as an incidental task added to the person who collected the data or operates the platform. The original FAIR principles frame findability, accessibility, interoperability, and reusability for both people and machines and emphasise machine-actionable metadata in data-intensive research [25]. FAIRness can make a resource easier to discover, understand, and reuse under stated conditions; it does not establish that a particular downstream operation is legally authorised. Work on FAIR data stewardship accordingly presents lifecycle care and machine-actionability as continuing professional work and argues for data expertise embedded within data-intensive teams [72]. Research on data-steward and data-librarian functions similarly identifies responsibility for data-management planning, metadata and standard definitions, quality rules, issue coordination, preservation, access, reuse, and communication across the lifecycle [73]. The exact allocation will vary by organisation: one steward may coordinate a bounded domain, while a public administration or data space may require a multidisciplinary stewardship office linking domain experts, legal and ethical functions, security, data protection, technical custodians, data owners, and external providers.

Requirements-engineering research supports the need to make this transformation explicit and traceable. Breaux and Norton [74] frame legal accountability through traceability, completeness, validity, auditability, and continuity, while rejecting guaranteed completeness. A systematic map of GDPR requirements-engineering studies identifies recurring difficulties concerning the interpretation of evolving legal language, differences in the issues addressed, and the still limited involvement of interdisciplinary perspectives [75]. The broader systematic mapping study by Kosenkov et al. [76] reaches a compatible conclusion: among 280 retained studies on regulatory requirements engineering, only a minority involved both legal experts and software engineers or connected requirements work to other software-lifecycle processes. These findings do not prove that an interdisciplinary team will reach the correct conclusion. They do, however, weaken the credibility of models in which legal interpretation is implicit, source transformation is unreviewed, or uncertainty is erased, and they support separate legal, technical, and integration review gates.

Time forms part of the assessment and helps determine whether a dataset remains usable. The information recorded in the metadata package (including purpose, provenance, permitted uses, restrictions, rights reservations, quality characteristics, recipients, and applicable conditions) must be read in relation to the moment at which a downstream activity is carried out. A dataset that was suitable for a particular use when first made available may no longer support the same conclusion at the time of reuse, including when it is later selected for AI training.

This is because both the dataset and the conditions attached to it may change over time. New versions may alter its content, purpose, structure, quality, or provenance; legal and contractual conditions may be amended; permissions may expire or be withdrawn; and new restrictions or obligations may arise. Metadata should therefore be versioned together with the dataset and linked to the period in which each statement was valid. This makes it possible for downstream actors to determine which conditions applied to the specific version they received and whether those conditions remain valid when the processing begins. A governance record should consequently include not only the substance of an assessment, but also its date, the dataset and metadata version to which it refers, and the events that require it to be reopened.

2.5 EHDS as a Sectoral Reference

The EHDS provides a concrete illustration of how an access regime distributes decisions across an institutional sequence. The Regulation distinguishes permitted and prohibited secondary-use purposes, applications and permits, responsible public bodies, secure processing conditions, monitoring, and staged application. Article 53(1)(e) includes the development, training, testing, and evaluation of algorithms and AI systems among the permitted purpose categories, while Article 54 identifies prohibited uses [1, arts. 53–55, 57–60, 67–69 and 73]. These provisions connect an asserted purpose to an application, a public decision, conditions of access, a processing environment, monitoring, and evidence. They illustrate the governance chain examined by the thesis without determining whether a particular applicant is entitled to a permit or whether a separate downstream operation is lawful.

Proposal-era commentary clarifies why this sequence is a governance problem rather than merely a technical exchange process. Terzis and Santamaria Echeverria [77] analyse the EHDS as a combination of legal, institutional, and technical-infrastructural change and identify concerns about the alignment, accountability, transparency, and representation of Health Data Access Bodies. Their analysis concerns the proposal and cannot establish the effects of the enacted Regulation. Its continuing relevance lies in showing that interoperability and access-body design allocate institutional power and responsibility alongside technical functions.

More recent implementation research makes these dependencies more concrete without

establishing that the enacted system is already operational. Svingel et al. [39] derive recommendations for Health Data Access Bodies from workshops, interviews, questionnaires, and experience gathered across technical activities and use cases. They identify national infrastructure, common application forms, catalogue practices, harmonised validation, provenance documentation, and multidisciplinary capacity as interdependent conditions of implementation. Their findings also show that alignment with common data models can involve information loss and that provenance requires continuing institutional effort. The significance for the present work lies in the governance of the mapping and access process: common specifications can coordinate discovery and exchange, but national capacity, validation choices, responsibility, and retained source information still affect what the resulting record means. The EHDS thus provides a limited illustration of the transition from an access decision to the documentation, evidence, and institutional responsibility examined below.

2.6 Documentation and Traceability

Dataset and model documentation research provides structured means of exposing decisions that would otherwise be lost. Data statements focus on population and language context relevant to NLP generalisation [78]. Datasheets organise questions about motivation, composition, collection, preprocessing, use, distribution, and maintenance across the dataset lifecycle [79]. Model cards instead focus on intended use, evaluation conditions, performance, and limitations of a model [80]. Data Cards add an audience- and workflow-oriented approach to dataset documentation [81]. These artifacts are complementary because they describe different objects, stages, and audiences; together, they show how lifecycle decisions may be recorded without assuming that one document can serve every governance purpose.

Whole-product documentation extends this prior work beyond separate dataset and model records. The AI Product Cards proposed by Celepija et al. [82] combine descriptive and machine-readable documentation for an AI product's problem framing, data, models, operations, deployment artifacts, and implementation references, with particular attention to adaptation and reuse. This work demonstrates that system-level documentation can connect components that separate cards leave fragmented. The remaining problem concerns the governance relations among those components: current authority, reviewed interpretation, conditional applicability, instrument-specific roles, decision authority, evidence ownership, uncertainty, disagreement, and reassessment still need to be represented explicitly. The thesis builds on whole-product documentation by addressing those relations rather than by claiming to introduce holistic documentation as such.

Documentation quality determines how far these artifacts can support accountability. A completed record may be partial, stale, copied, or inaccurate, and may omit downstream transformations, rights decisions, responsible roles, or review status. Hutchinson et al. [83] accordingly position documentation within a wider accountability infrastructure, while com-

parative studies report heterogeneous formats, incomplete fields, and weak standardisation in dataset-documentation practice (see also [84, 85, 86]). The large-scale audit by Longpre et al. [87] further shows that licence metadata and attribution involve substantial human and legal work. The retrospective documentation of BookCorpus similarly shows how sparse early records can leave later users reconstructing composition, duplication, representational skew, and rights-related concerns only after widespread reuse [88]. Documentation may therefore provide reviewable evidence about a dataset or system; ownership, enforceability, and downstream lawfulness remain separate determinations.

Catalogue and ML-metadata standards offer reusable structure. DCAT 3 and its European application profile, DCAT-AP, describe catalogues, datasets, dataset series, distributions, data services, versions, agents, access information, and relevant legal metadata [26, 89]. Croissant addresses the internal structure of ML datasets more directly, including files, record sets, fields, and training-related splits [31]. The standards are therefore complementary: they solve substantial parts of the discovery, exchange, and dataset-description problem at different levels of detail.

Application profiles and crosswalks provide two related ways of organising metadata reuse. An application profile adapts existing vocabularies to a defined community or purpose by selecting relevant terms and specifying how they should be used. It may be expressed through human-readable documentation, machine-actionable structures, and validation rules [90]. A crosswalk performs a different function: it records how elements or concepts in one schema correspond to those used in another.

These correspondences should not be assumed to be complete or exact. The comparison of fourteen research-data crosswalks by Wu et al. [91] shows that descriptive elements were mapped more consistently than rights information, structural metadata, and controlled vocabularies. A mapping may therefore remain useful even where it does not preserve every distinction in the original schema. Its value depends on making visible what has been mapped, in which direction, under what assumptions, through which transformation, and with reference to which source and version.

The schema- and ontology-matching literature reinforces this qualification. Rahm and Bernstein [92] organise matching approaches according to the information and combination strategies used to propose correspondences. More recent empirical work reports material uncertainty and error in automated ontology matching and finds greater maturity in semi-supervised than fully automated approaches [93]. These contributions justify computational support for candidate identification, but not automatic acceptance of a legal-to-metadata relation. In the present setting, a proposed correspondence must still be reviewed against the source proposition, the intended metadata capability, the relevant version of the standard, and the use-case context.

Current application profiles further illustrate the value and limits of specialisation. MLDCAT-AP 3.1.0, published by SEMIC as a Candidate Recommendation, extends DCAT-

AP with descriptions of machine-learning models, datasets, evaluations, quality measures, publications, and provenance relations [94]. At the review date, HealthDCAT-AP Release 7, dated May 2026, was the European Commission Draft Specification for catalogue description in the EHDS context [95]. Its practical relevance is already visible in catalogue preparation: the German Federal Institute for Drugs and Medical Devices states that its national health-data metadata catalogue will be based on HealthDCAT-AP [96]. These developments establish direct precedent for representing legally relevant EHDS dataset characteristics in metadata. They do not, however, make catalogue conformance a decision that a permit should be issued, an asserted right is valid, or a downstream training operation is lawful.

AIDOC-AP extends the application-profile approach from catalogue description to AI-system technical documentation. The resource models systems, datasets, models, lifecycle activities and responsible agents, and links these objects to an ontology of the AI Act's Annex IV documentation requirements [97]. It is particularly close to the present thesis because it uses an application profile to connect regulatory documentation needs to machine-readable lifecycle objects. Its object and evidentiary status remain different: AIDOC-AP addresses the technical-documentation obligation associated with Article 11 and Annex IV rather than the authority to obtain and reuse CEDS data for a defined training operation, and its journal submission was under major revision at the review date. It therefore provides a recent design comparator, not peer-reviewed evidence that the profile is complete or operationally effective.

The proposed approach relies on several vocabularies because no single one covers every part of the assessment. For legal-source identification, the framework draws on the Publications Office Core ELI OWL distribution dated 17 July 2020, which declares version 1.3, while ELI-EP 1.5 remains a separate European Parliament application profile [98, 99]. A later official Publications Office OWL distribution declares Core ELI version 1.5; the framework's present use of Core ELI is confined to version 1.3 and does not assess that later release [100]. ELI can therefore help identify the legal source, but it does not express the reviewed proposition derived from that source.

The remaining layers require different vocabularies. DPV provides reusable terms for purposes, processing activities, data categories, legal bases, roles, measures, rights, and risks, together with versioned resources and jurisdiction-specific extensions [30, 101]. DQV addresses a different question by representing quality dimensions, measurements, policies, annotations, and provenance [29]. It does not determine whether data are suitable in the abstract, since that judgment remains dependent on the intended use. These vocabularies are therefore complementary: ELI identifies the source, DPV describes the governance context, and DQV records the evidence relevant to data quality. Combining them does not mean reducing these different functions to a single compliance field.

The Data Use Ontology (DUO) supplies a stronger production precedent for machine-readable conditions on sensitive biomedical datasets. DUO standardises terms derived from

consent and data-sharing policies, attaches allowable-use conditions to datasets, and can be matched against a research purpose and a researcher's machine-readable permissions. The reported implementation included the Broad Data Use Oversight System, major international repositories, and more than 200,000 annotated datasets at the time of publication [102]. This experience shows that metadata can materially support discovery, access requests, and data-access committee decisions. It also exposes a boundary relevant to this thesis: DUO deliberately limits logical expressivity to preserve interoperability and does not encode every legal, contractual, ethical, or contextual combination. It standardises a reviewed data-use vocabulary; it does not determine the full applicability of EU data-space, data-protection, copyright, or AI rules to a later operation.

Provenance standards supply a complementary lineage layer. PROV-O represents entities, activities, agents, derivation, attribution, association, generation, and invalidation [28]. PROV-AQ nevertheless separates access to provenance from trust in its authority or correctness [103]. A provenance graph can therefore reconstruct an asserted transformation chain without authenticating its inputs, proving that an omitted event did not occur, or determining whether a represented operation was lawful.

Taken together, the literature leaves one relationship insufficiently developed: how a versioned data resource remains connected to the governance context in which it is proposed, assessed, approved, transformed, used, shared, retained, or withdrawn. At the same time, it provides the basis for such a profile. Existing concepts for datasets, distributions, record sets, and provenance can be reused, while purpose, responsibility, applicability decisions, controls, evidence, and change remain explicitly linked to the relevant data version and stage of use.

A technical representation can give an assertion a defined structure, while authentication and legal interpretation remain separate tasks. A catalogue can correctly express a publisher's claim without establishing that the publisher owns every right asserted, and a structurally valid ML description can remain incomplete for a particular governance decision. Consequently, a mapping method must keep source, interpretation, applicability facts, requirement, role, decision or control, evidence, and change distinct, while allowing them to refer to shared dataset and version identifiers.

Maintaining these distinctions avoids treating the legal source as if it produced one universal requirement, the requirement as if it prescribed a single control, or the existence of a control or supporting evidence as proof of compliance. Traceability performs a different function: it reconstructs the basis of a conclusion, the decisions through which it was reached, and the actor responsible for maintaining it, without replacing the legal assessment on which that conclusion depends.

2.7 Formalisation and Legal Judgment

Compliance-by-design research brings together legal knowledge representation, requirements engineering, policy languages, control mapping, and continuous compliance. Across these approaches, a recurring distinction remains necessary: a formal language can record a legal or governance proposition, but it does not determine whether that proposition is legally correct or applicable to the facts under examination. LegalRuleML, for example, can represent obligations, permissions, temporal conditions, jurisdictional scope, and links to legal sources. These capabilities make the reasoning structure visible, but they do not determine which interpretation of the source should be adopted [104].

ODRL illustrates the same distinction at the level of data-use policies. It can express that a particular party is permitted, prohibited, or required to perform an action in relation to an asset and under defined constraints. The policy can therefore be represented in a machine-readable form, but the language does not by itself verify the underlying facts, monitor performance, enforce the stated conditions, resolve conflicts between policies, or establish their legal effect [27]. In both cases, formal representation supports traceability and implementation, while interpretation, applicability, and legal validity remain separate assessments.

Implemented systems demonstrate both potential and scope limits. The SPECIAL policy language developed by Bonatti et al. [105] supports executable checking for a deliberately bounded consent and business-policy subset. The authors do not encode the whole GDPR, and surrounding provisions, temporal context, input truth, and prior interpretation remain outside the checked subset. Successful execution therefore shows that the encoded policy behaved as specified under the supplied facts and semantics; it does not validate omitted law or the interpretation on which the policy rests.

The closest general conceptual precedent is evidence-based compliance engineering. Hulstijn [106] links legislation, formal legal requirements, cases, and evidence documents through forward and backward traceability. The proposal treats legal interpretation and evidence as first-class elements, uses change to trigger updates and renewed verification, and rejects the reduction of legal compliance to process conformance. This substantially overlaps with the source–requirement–evidence–change chain used here. The remaining difference is the object of translation: the proposal maintains compliance invariants over case data, whereas this thesis proposes an intermediate, standard-neutral step for expressing selected legal concepts as metadata needs and examining their possible representational fit across heterogeneous dataset, provenance, quality, privacy, policy, and ML-documentation standards.

The Trustworthy AI Requirements (TAIR) ontology provides the closest precedent for regulation-to-standards mapping. Hernandez et al. [107] represent normative requirements and concepts from the AI Act and international standards in an open knowledge graph, record implementation and tracking responsibilities, and use mappings to examine how far standards conformance supports regulatory compliance. Because TAIR already maps

regulatory requirements to standards, this thesis makes no broad claim that it introduces repeatable legal-requirements-to-standards mapping. The present work addresses a different and narrower transformation: conditional legal propositions concerning CEDS participation, access, reuse, and downstream operations are expressed as standard-neutral metadata needs before candidate elements of relevant specifications are examined. The bounded Data Act application illustrates how the framework can reach exact technical elements for one circumscribed need while preserving a related question as a non-metadata judgment, without reducing the analysis to a single correspondence or coverage score.

Recent implementations move further from representation towards executable assurance. Nweke and Yeng [108] connect clauses to controls, executable checks, versioned releases, evidence contracts, runtime decisions, and OSCAL assessment packages, and evaluate a prototype on two public identity workloads. In a version 1 Open Research Europe article, Sabaliauskaitė et al. [109] instead structure civilian and defence obligations for a dual-use AI system in RDF, configure them by lifecycle and deployment context, and query and validate the resulting knowledge graph with SPARQL and SHACL. These systems demonstrate that clause-to-control traceability, explicit evidence gaps, change-impact analysis, and deployment-specific compliance graphs can be implemented. Their evaluations remain bounded to the encoded obligations, supplied facts, threat models, and selected system contexts. Neither addresses the source-linked intermediate step proposed here for relating legal requirements to metadata needs across dataset standards nor resolves the authority for obtaining and using CEDS data in a particular downstream operation.

Legal-requirements research makes the interpretive dependency more explicit. Breaux and Antón [110] show how regulatory statements can be transformed into operational rules only through structured interpretation, while later mapping and traceability work continues to report ambiguity, vocabulary alignment, change, and validation as persistent problems [111, 112]. Formalisation can expose assumptions and improve consistency; it cannot eliminate the professional judgment by which source text is selected, scoped, and converted into a requirement.

Traceability research supplies a more mature, though still bounded, evidence base. The review of pre-requirements traceability by Mucha et al. [113] finds many prototype and single-case evaluations. Ruiz et al. [114] report practitioner concerns about cost, maintenance, collaboration, and manual effort. The controlled experiment by Mäder and Egyed [115] shows that trace support can improve performance in bounded software-maintenance tasks, while leaving link-creation and maintenance costs unresolved. These results support traceability as an aid to review and maintenance, not as proof that a legal trace is complete or correct.

2.8 Limits of Technical Enforcement

Data-space specifications contribute to a traceability framework by allocating distinct functions across the data-exchange process, rather than by providing a complete governance solution. The DSSC Blueprint connects governance arrangements to technical building blocks, showing how rulebooks, participant roles, catalogues, and access and usage policies may structure the conditions under which data are offered and exchanged [116]. IDS-RAM develops this functional separation by distinguishing policy information, decision, and enforcement, while requiring abstract policies to be mapped to the relevant implementation environment before they can govern a specific operation [117]. The Dataspace Protocol addresses a narrower operational layer by specifying catalogue exchange, contract negotiation, and transfer-control states [50]. Taken together, these specifications provide reusable mechanisms for coordinating and recording data-space interactions, but they do not determine whether a policy is legally valid, substantively appropriate, or effective once data leave the controlled environment.

Research on data intermediaries shows that trustworthy data exchange depends not only on technical architecture but also on the explicit allocation of purpose, legitimate use, responsibility, oversight, enforcement, and accountability [118]. Data-space specifications do not assign these functions to a universally defined data steward. They nevertheless create an organisational need for a function that maintains the connection among governance decisions, participant roles, policy conditions, evidence, and operational changes. In this thesis, stewardship is proposed as that continuity function, while decision authority remains with the actor competent under the applicable legal and institutional arrangement.

Gaia-X provides a related production-oriented trust and conformity architecture. Its compliance scheme separates participant declarations, machine-readable credentials, automated validation, issuer and signature checks, and assessment by recognised actors, while retaining versioned rules and evidence associated with claims [119]. This arrangement closely resembles the thesis's separation of asserted metadata, supporting evidence, technical validation, and institutional authority. It also confirms why these layers cannot be merged: a valid credential or successful shape check establishes conformance to the encoded scheme under its trust model, not compliance with every law applicable to a particular downstream use.

The relevance of these architectural specifications extends beyond voluntary data-space design. At the horizontal EU level, Article 33 of the Data Act establishes essential interoperability requirements for participants in data spaces, including machine-readable descriptions of dataset content, use restrictions, licences, collection methods, quality and uncertainty, data structures, vocabularies, access interfaces, and applicable policies [32, art. 33]. The provision is relevant here because it gives legal significance to several architectural concerns without prescribing a complete implementation model. ETSI's subsequent analyses examine how existing specifications might support those requirements and where additional standard-

isation remains necessary [120, 121]. The reports thus illustrate the connection between a horizontal legal interoperability requirement and the technical work needed to implement it, while also showing that the mapping remains incomplete and subject to review.

The Interoperable Europe Act provides an adjacent public-sector framework for trans-European digital public services, distinguishing legal, organisational, semantic, and technical dimensions of interoperability. The present analysis uses it only as contextual material rather than as a rule governing CEDS generally [122].

These boundaries follow from the layer at which each specification operates. Equivalent ODRL text can be interpreted by different evaluator functions, trusted-context feeds, and enforcement environments. Protocol conformance can show that messages and state transitions behave as specified while leaving the legal validity of the policy untouched. Activities outside instrumented components (e.g., copied files, screenshots, model weights, unauthorised derivatives, or human actions) still depend on organisational measures, secure environments, monitoring, contracts, and remedies.

Implementation studies reinforce this boundary. Usage-control research demonstrates technically meaningful enforcement and monitoring mechanisms in distributed systems, but also depends on observable events, trusted components, and defined enforcement domains [123]. Practice-oriented studies of data sovereignty and sectoral data-space implementation report integration costs, policy translation problems, heterogeneous infrastructures, and organisational coordination as continuing barriers [124, 125]. A 2023 survey of four open-source IDS-compliant connectors likewise identified differences in architecture, data models, usage-control support, and documentation, illustrating the distance between a declared connector capability and interoperable operational enforcement [126]. These findings support enforcement at defined points rather than an assumption of continuous control after data leave an instrumented environment.

Commercial platforms also pursue policy-to-control integration. OneTrust describes regulatory intelligence, risk evaluation, policy interpretation, controls, runtime enforcement, monitoring, evidence, and decision lineage [127]. Its documentation does not expose the source selection, version controls, mappings, error rates, or independent evaluation needed to assess semantic and legal accuracy. It therefore demonstrates production interest without replacing an inspectable legal-to-metadata derivation.

For the purposes of this work, these architectures and standards provide reusable vocabulary and modular components. DCAT-AP and Croissant supply candidate elements for catalogue discovery and ML-dataset structure, while a policy-expression or enforcement mechanism must be selected for the relevant environment. Exact components can be compared with a source-linked metadata need while preserving the function and scope of a catalogue record, policy, negotiation state, permit condition, processing environment, or log [26, 31, 116].

This functional view also separates strategic data governance from operational data stew-

ardship. The Open Data Institute describes responsible stewardship as an iterative and systemic process in which contexts and mitigations evolve and the effects of data practices may extend beyond one organisation's control [16]. Governance determines why a data package exists, which conditions apply, who holds decision authority, what evidence is required, and which criteria permit a lifecycle transition. Stewardship maintains those decisions through coordination, metadata, versioning, access administration, transformations, quality checks, evidence capture, and change records. A trace may connect a source and requirement to a governance decision, responsible role, control, evidence, review, and lifecycle gate without assuming that every requirement has a technical implementation. Because this allocation depends on whether organisational actors can inspect and challenge the record, the analysis now turns to meaningful human review.

For public organisations, stewardship preserves continuity across decision makers, departments, suppliers, data, and evidence. It records the applicable version, definitions, provenance, quality limits, access conditions, responsible actors, evidence, review date, and reopening trigger. Missing, disputed, or changed conditions should be escalated rather than converted into a favourable technical status. Stewardship thus maintains reviewed assertions without establishing their legal or factual truth.

2.9 Authority in Human Oversight

Formalisation can make requirements, decisions, and evidence more consistent and inspectable, but it does not by itself ensure that the people reviewing them can recognise, contest, and act on their implications. Human involvement should therefore be understood as an organisational capability rather than as the mere presence of a person at the end of an automated process. Morley et al. [128] emphasise that ethical principles require contextual interpretation and organisational embedding, while Mittelstadt [129] shows that principles alone do not supply the professional norms, institutions, and methods needed for their application. Selbst et al. [130] further demonstrate how technical representations can obscure the institutional and social context in which a system operates. Together, these contributions indicate that meaningful review depends on how responsibilities, information, and authority are organised around the reviewer.

Empirical studies identify the practical conditions under which this organisational capability may be developed. Madaio et al. [65] show that governance instruments must be aligned with existing workflows, roles, organisational culture, and supporting infrastructure, although their study does not establish that checklists improve downstream fairness or compliance. Rakova et al. [66] identify distributed accountability and weak coordination between organisational functions, together with needs for monitoring and authority to stop or escalate a process. Similarly, Holstein et al. [131] report that practitioner requirements remain dependent on the system, professional setting, and stage of development. A review

of practical AI-ethics approaches likewise reports a heterogeneous landscape and limited evidence concerning validity, reliability, efficacy, and adoption, supporting context-sensitive evaluation rather than treating completion of a tool as ethical assurance [132]. The relevant question therefore concerns whether a human reviewer has access to the necessary evidence, understands the applicable context, and has sufficient authority to challenge or reopen a decision; nominal involvement alone cannot answer it.

The AI Ethics Canvas incorporated in this thesis is one response to this organisational problem. Its development treats stakeholder engagement as an iterative activity and draws on direct formative evidence from twelve qualitative interviews and a workshop with 18 participants, two broader project cases, and three scenarios [133]. Survey and focus-group material provides context about public-sector needs rather than another measure of artifact performance. The Canvas illustrates how a structured tool may support multidisciplinary questioning without certifying ethics, technical quality, or legal compliance.

Reviewability concerns whether a decision can be reconstructed and challenged across the system lifecycle, rather than whether a final explanation or approval has been recorded. Cobbe et al. [67] connect review to the institutional context, data, design choices, testing, deployment, effects, and opportunities for challenge. Amershi et al. [134] similarly treat human–AI interaction as a process that evolves through use, error, and system change, while Ehsan et al. [135] emphasise the value of recording who acted, what was decided, and why. In human-centred compliance, these records must remain connected to their factual and legal basis so that earlier decisions can be reviewed rather than followed as unquestioned precedent.

Human review does not, by itself, ensure that an automated error will be identified or corrected. Explanations may increase acceptance of both accurate and inaccurate recommendations, while users of interpretability tools may remain confident despite overlooking known problems [136, 137]. Interventions designed to reduce overreliance can improve deliberation, but may also introduce usability and trust costs [138]. More broadly, human–AI combinations tend to outperform humans acting alone but do not consistently outperform the better-performing human or AI component [139]. These findings caution against treating the presence of a reviewer as evidence that oversight is effective.

The effectiveness of oversight therefore depends on the conditions under which the reviewer operates. Research on automation bias documents both omissions and incorrect actions under decision support, while empirical studies show that nominal human involvement may reproduce discriminatory outcomes or encourage deference to erroneous recommendations [140, 141, 142, 143, 144]. Reviewing government policies, Green [15] warns that unsupported oversight requirements may create false assurance and diffuse accountability. By contrast, reviewers with relevant professional knowledge and access to contextual information may identify errors invisible from the system output alone [145]. Meaningful review consequently requires appropriate evidence, competence, independence, time, decision au-

thority, and a defined route for challenge or escalation.

2.10 Bounded Evaluation

The same distinction between assistance and judgment also constrains evaluation. In legal-NLP research, the expression "ground truth" should be used with caution because a reference set normally represents decisions made by annotators or experts under a defined task, annotation scheme, and factual context. It provides a basis for comparison within those conditions, but it does not necessarily establish the only legally valid interpretation.

Our LREC-COLING study provides a bounded example of this distinction. It adapted categories from the expert-annotated Contract Understanding Atticus Dataset (CUAD) to Italian Terms-and-Conditions, producing an annotation scheme covering fourteen categories of potentially relevant contractual clauses over 3,626 clauses from 41 contracts; the classification experiments addressed the eleven most frequent categories [146, 147]. Unlike prior work on supervised detection of potentially unfair online Terms of Service clauses, the study classified a broader set of clause types rather than legal unfairness [148]. The reference annotations made retrieval and classification measurable within the defined corpus. They did not determine the interpretation, enforceability, or legal effect of the identified clauses, and the study does not validate the legislative decomposition or metadata-capability derivation performed in this thesis.

A similar qualification applies to automated compliance evaluation. Cappelli et al. [149] assess an exploratory AI Act checker through a bounded set of cases and reference answers constructed by legal experts. The study reports promising performance within that setting and uses the remaining discrepancies to identify aspects requiring further refinement. Its reference answers provide a transparent evaluation benchmark rather than an objective resolution of every possible legal question. This distinction is reinforced by Dijck et al. [150], who show that agreement in the annotation of EU legislation varies according to the task, annotators, metric, and revision stage. Expert agreement can therefore indicate the clarity and reproducibility of an assessment procedure, while the legal validity of a contested proposition remains dependent on the relevant authority, facts, and interpretation.

Legal-informatics work on explainable case reasoning adds another relevant distinction: an explanation can expose the argument structure and factors used by a system without proving that its legal conclusion is correct or transferable to a different doctrinal task [151]. Evaluation must consequently separate fidelity to the implemented reasoning process from validity of the underlying legal interpretation. This distinction also determines how far evaluation evidence can support the research gap and contribution stated below.

2.11 The Integration Gap

Substantial components for the accountable governance of data used in AI training already exist, and some are integrated. DUO and HealthDCAT-AP connect dataset descriptions to domain-specific access or reuse information; TAIR connects AI Act requirements to standards; evidence-based compliance engineering connects legislation, requirements, cases, evidence, and change; and compliance-as-code, Gaia-X, IDS-based architectures, Eclipse Dataspace Components (EDC) [152], and commercial platforms connect policies to technical checks, credentials, enforcement, and audit records. These precedents mean that the thesis cannot claim priority for machine-readable data-use conditions, regulation-to-standards mapping, clause-to-control traceability, or policy-to-evidence integration. Their different objects and boundaries nevertheless leave a narrower integration question unresolved.

Table 2.2 condenses that result. Each strand supplies a necessary part of the route but leaves unresolved work for another technical or organisational layer.

Table 2.2: Contributions of the reviewed approaches and remaining integration tasks

Approach	**Contribution**	**Remaining integration task**	**Human responsibility**
Legal analysis and requirements engineering	Identify sources, roles, modalities, conditions, exceptions, and traceable propositions.	Translate a reviewed proposition into information and action without losing context or uncertainty.	Adopt and revisit the interpretation for the factual configuration.
Metadata and documentation	Describe datasets, provenance, quality, policies, purposes, and evidence in reusable form.	Establish authority, truth, sufficiency, and the legal effect of the recorded assertion.	Maintain the record and assess the evidence behind it.
Data-space architecture and compliance mechanisms	Support catalogues, negotiation, validation, bounded checks, enforcement, logs, and credentials.	Connect controls to current source-linked reasoning and activity outside the observable enforcement domain.	Authorise controls, manage exceptions, and act on violations or change.
Participatory and human-centred governance	Surface assumptions, values, responsibilities, and contextual knowledge across professional roles.	Preserve decisions and reasons through later technical and organisational transitions.	Deliberate, challenge, escalate, and retain accountable decision authority.

Across these strands, the gap lies in the distance between three kinds of object, despite the availability of regulation, ethics methods, metadata vocabularies, and enforcement components. A legal norm acquires practical meaning through interpretation in a factual and institutional setting. An organisation must turn that interpretation, together with public values and operational knowledge, into a decision for which someone has responsibility. A technical representation may then carry selected information and support a bounded action. The reviewed literature shows mature work at each point, but also recurring loss when one object is treated as equivalent to the next.

Strategically, sovereignty scholarship shows that regulatory ambition and formal control do not create material capacity by themselves; CEDS may provide shared governance and

infrastructure for data, but the literature does not support an inference that they remove cloud, compute, platform, skills, or organisational dependencies. Research on public-sector AI and ethics tools, meanwhile, shows the value of multidisciplinary deliberation while reporting limited evidence that completion of an artifact changes later outcomes. At the informational level, metadata, requirements engineering, policy languages, and compliance architectures provide reusable representations and controls, but leave the authority, truth, sufficiency, and legal effect of the assertions they process to separate assessment.

The thesis is therefore positioned as an integrative but bounded contribution. It connects the capacity problem identified through digital sovereignty, the structured deliberation supported by the Canvas, and a source-linked five-stage translation from legal source to representational outcome. It does not claim priority for regulation-to-standards mapping, machine-readable data-use conditions, participatory ethics tools, or policy-to-evidence integration. Its narrower claim is that public administrations engaging with CEDS need these capacities to remain distinct and connected: access must not be confused with downstream authority; deliberation must not be confused with approval; and a metadata match must not be confused with compliance.

This positioning also determines the evidentiary boundary. The thesis brings together two bodies of work with their own methods and one conceptual contribution. The digital-sovereignty analysis diagnoses a relationship between formal rights and material capacity; the Canvas study reports formative development of an organisational artifact; and Chapter 5 proposes a source-linked conceptual framework whose intended operation is illustrated through two bounded micro-applications to one Data Act provision. Their relationship is a thesis-level synthesis rather than evidence that an end-to-end workflow has been implemented. Empirical studies may examine their feasibility and organisational value separately or in selected combinations, without treating integration as a necessary outcome.

Chapter 3

Digital Sovereignty, Regulatory Capacity, and Common European Data Spaces

This chapter is based on work previously published as: Riccardo Nanni, Pietro G. Bizzaro and Maurizio Napolitano, "The False Promise of Individual Digital Sovereignty in Europe: Comparing Artificial Intelligence and Data Regulations in China and the European Union," *Policy & Internet* 16, no. 4 (2024): 711–726, https://doi.org/10.1002/poi3.424. The chapter has been revised and extended to connect the published legal analysis to Common European Data Spaces and to the wider argument of this thesis.

3.1 Sovereignty and Capacity

European debates about digital sovereignty begin with an intuitively simple concern: actors should be able to make meaningful choices about the digital systems, services, and data on which they depend. The simplicity disappears as soon as the subject of sovereignty is specified. The term may refer to the authority of a state, the autonomy of the European Union, the control claimed by an organisation over data and infrastructure, or the agency of an individual in relation to a platform. These subjects possess different powers, confront different dependencies, and act through different institutions. A state may regulate a provider; an organisation may negotiate an access arrangement; an individual may exercise a data-protection right. None of these acts is interchangeable with the others, even though all may be described as exercises of digital sovereignty [8, 21, 153].

We examined one particularly demanding formulation: the European promise of *individual* digital sovereignty. That promise links sovereignty to the possibility for people to

understand and influence what happens to their data and how digital systems affect them. It draws normative force from individual freedom and self-determination. In the European policy discourse under examination, rights of access, information, erasure, portability, and withdrawal were presented not merely as isolated protections but as elements of a wider capacity to choose in digital environments. Platform regulation and rules for artificial intelligence added protections against forms of market or informational power that could make individual choice nominal rather than real [9].

The analysis asked whether that individual promise could be attained through the European Union's regulatory approach. Its answer was deliberately narrower than a general judgment on European digital policy. Regulation is indispensable for defining rights, allocating duties, and creating avenues for enforcement. The difficulty is that legal entitlement and practical agency do not coincide. A person may formally be entitled to change provider while facing a market with few functionally equivalent alternatives. A public body may be legally permitted to procure a different service while lacking compatible infrastructure, staff, or bargaining power. A company may obtain a right to access data but lack the technical or organisational capacity to make the data usable. Sovereignty in this sense depends on the relationship between a recognised freedom and the material conditions under which it can be exercised.

This distinction provides the point of departure for the whole thesis. The comparison diagnoses a mismatch between a regulatory promise and the industrial capacity needed to make individual choice meaningful. This chapter retains that finding, updates its institutional context to July 2026, and develops a new synthesis from it. Common European Data Spaces (CEDS) can be read as a partial attempt to build some of the missing capacity at the level of data access, governance, and interoperability. They are relevant to digital sovereignty precisely because they seek to change the conditions under which European actors can discover and exchange data. They do not, however, supply every component of technological autonomy, and they do not convert data availability into a lawful, technically adequate, or socially legitimate AI project.

Because these parts rely on different evidence, three layers must remain distinct. Sections 3.2–3.4.1 present the comparative question, method, legal analysis, and conclusions. Section 3.4.1 reports a bounded documentary update based on EU legal and policy materials available by July 2026. Sections 3.5–3.6 offer the synthesis developed for this thesis: they connect regulatory capacity, CEDS, the AI competition, and the operational position of public administrations. The first part does not empirically assess CEDS, AI factories, data labs, cloud infrastructure, or public-administration projects. Those matters enter only in the update and synthesis, and the status of each claim is marked accordingly.

3.1.1 From territorial authority to individual agency

Sovereignty traditionally names a claim to final authority within a territory. Digital networks complicate that image because the relevant resources and decisions are distributed. A service offered in one Member State may depend on a provider established in another jurisdiction, computing facilities in several regions, a model trained on data assembled elsewhere, and technical standards maintained by transnational communities. Control over one layer does not imply control over the others. The capacity to adopt a rule, for example, differs from the capacity to operate the infrastructure to which the rule applies. The ability to host data differs from the ability to design processors, train models, or offer competitive services. Digital sovereignty is consequently better understood as a continuing effort to reduce, structure, or make dependencies contestable than as a condition of complete independence.

European policy has given this effort a characteristic rights-based orientation. The Union exercises substantial regulatory authority through its internal market and has used that authority to establish rules for personal-data processing, platform conduct, access to data, and artificial intelligence. This approach protects interests that markets or technical architectures would not necessarily secure by themselves. The General Data Protection Regulation (GDPR), for example, establishes enforceable rights and duties concerning personal-data processing, including transparency, access, erasure, portability, and conditions for valid consent [3]. The Data Governance Act (DGA) and Data Act organise different aspects of data sharing, reuse, and access without displacing the GDPR where personal data are involved [32, 56]. The AI Act adds a risk-based regulatory framework for specified AI practices and systems [2]. Together these instruments make legal accountability a central component of the European digital model.

Individual digital sovereignty nevertheless asks more of this body of law than the existence of rights. A right can alter bargaining relations, require a provider to disclose information, or permit an individual to challenge an unlawful practice. Yet a formal option becomes an effective choice only if the person can understand it, bear the cost of exercising it, and find an acceptable alternative. The same is true at an institutional level. A public administration may possess formal authority to use a dataset or develop an AI service, while remaining dependent on suppliers for infrastructure, integration, maintenance, and expertise. The legal power to specify contractual conditions can reduce that dependence, but it cannot by itself create technical substitutes or internal competence.

The argument concerns whether regulation provides a complete set of means for the stated objective. Regulation can make power visible and contestable, constrain conduct, and create a common baseline for market participation. It may also influence practices outside the Union when providers adapt products to retain access to the internal market. These are material effects. The 2024 analysis therefore did not claim that European digital regulation was merely rhetorical or that individual rights lacked value. It questioned the stronger proposition that a predominantly regulatory strategy could, without sufficient industrial

and infrastructural alternatives, deliver the level of individual autonomy associated with sovereignty [9, 154].

The distinction also guards against two opposite forms of reductionism. Legal reductionism treats the adoption of a rule as evidence that the social or technological condition named by it has been achieved; technological reductionism assumes that ownership of infrastructure or the availability of a domestic provider is sufficient to protect individual freedom. A domestic platform can still misuse data; a European cloud service can still be insecure; a locally developed AI system can still discriminate or operate without a valid legal basis. Digital sovereignty requires both normative limits and practical capacities. The dispute concerns their configuration, not a choice between law and technology.

3.1.2 Data, platforms, and infrastructure

The comparison organised digital-sovereignty measures into three connected realms: data, platforms, and infrastructure. The division was analytical rather than absolute. Platform services depend on infrastructure and process data; data rules shape how platforms operate; infrastructure determines which services and exchanges are technically possible. The tripartition nevertheless made a difference in emphasis visible. The European approach examined in 2024 was particularly developed in the data and platform realms. Infrastructure was addressed through cybersecurity measures and programmes, but no single infrastructure-oriented legal instrument played a role equivalent to that of the GDPR in data protection or the emerging platform rulebook [9].

The three realms correspond to different forms of capacity. Data regulation determines who may collect, access, share, port, or reuse data and under what conditions. Platform regulation constrains intermediaries whose market position, recommender systems, or AI services influence what people and organisations can see and do. Infrastructure policy concerns the physical and logical resources on which those activities depend: connectivity, cloud and edge services, data centres, high-performance computing, processors, identity and trust services, and interoperable exchange mechanisms. A regulatory success at one level may be neutralised by dependence at another. Data portability, for instance, is more useful when an alternative provider can ingest the exported data with tolerable switching costs. A data-sharing right contributes more to innovation when recipients have the compute, expertise, and organisational processes needed to use the data responsibly.

Data spaces occupy the intersection of the three realms. They require rules about access and use, services for discovering and exchanging data, and shared technical components and standards. They may reduce dependence on isolated bilateral arrangements and support a European market for data. Yet they are not equivalent to a general cloud, semiconductor, or AI-industrial policy. Nor are they a reservoir from which any actor can take data for any purpose. Their contribution lies in creating governed possibilities for exchange while preserving the responsibilities and conditions attached to particular data and uses [4, 5, 22].

3.2 Question, Comparison, and Method

The research question was: *Is individual digital sovereignty attainable for the European Union?* We approached the question by comparing selected European and Chinese rules on data and AI governance. China was not treated as a normative model for Europe, and textual resemblance was not equated with equivalent protection. The comparison served a critical purpose. If legal mechanisms presented as features of a distinctively European promise of individual empowerment also appeared, in relevant respects, in a polity with a different constitutional structure and a security-centred account of Internet sovereignty, the distinctiveness and sufficiency of the European promise required closer examination [9].

The European corpus consisted of the GDPR, Data Governance Act, Data Act, and the then-final AI Act text. The Chinese corpus centred on the Personal Information Protection Law (PIPL) and Data Security Law (DSL) [155, 156], and regulations on algorithmic recommendation services [157], with the Cybersecurity Law [158], critical-information-infrastructure rules [159], cross-border data-transfer measures [160], and rules on deep synthesis and generative AI [161, 162] used to explain the surrounding architecture. The selection followed the substantive focus of the question: data governance and AI-mediated platforms were the principal fields in which claims about individual autonomy could be compared. Infrastructure measures were mapped as a third realm but were not examined with the same depth. That asymmetry became an explicit limitation and a direction for later work.

We conducted a qualitative thematic analysis of the selected legal materials. Themes were developed during the reading of the texts rather than imposed as a closed codebook in advance. The analysis identified the protected or regulated subject, the actors on whom duties were placed, the powers assigned to public authorities, mechanisms for consent or user choice, rules for data minimisation and cross-border transfer, and obligations concerning algorithmic services. It then situated those findings in each polity's digital-sovereignty discourse. The method was interpretive and comparative. It did not calculate compliance rates, evaluate enforcement files, interview affected individuals, or measure market outcomes.

The legal comparison was also temporally bounded. At the time of analysis, the Data Act and DGA had little or no application history from which their wider consequences could be established, and the AI Act's implementation was still ahead. We therefore compared enacted or final legal texts and the policy claims made around them. The analysis could identify similarities, differences, allocations of authority, and the relationship between legal design and the sovereignty discourse. It could not demonstrate how often an individual successfully exercised a right, whether a provider changed conduct because of a specific enforcement action, or whether the legislation would eventually stimulate European industrial alternatives.

These method limits matter for the present chapter. The conclusion was a diagnosis about the insufficiency of a regulatory strategy for a demanding promise, not an econo-

metric finding that regulation had no effect. The 2026 update below addresses a different evidentiary period and does not evaluate later initiatives through the earlier corpus. It asks whether subsequent policy has begun to address the missing material dimensions and how CEDS should be positioned in that wider response. The update uses institutional documents and enacted legislation; it does not claim to measure operational performance.

3.2.1 Rationale for the EU–China Comparison

At first sight, the European Union and China appear too different for a useful comparison. The Union is a supranational legal order composed of democratic Member States, while China is a centralised party-state. Their institutions, judicial safeguards, relationships between public power and private actors, and conceptions of political freedom are not equivalent. A provision granting an individual a right against a private data handler cannot be assumed to have the same practical meaning where avenues of redress, state accountability, and institutional independence differ. We made these differences part of the analysis rather than treating them as noise [9].

The comparison was nevertheless analytically productive. Both polities regarded data, algorithms, and digital infrastructure as strategic and sought to constrain large technology companies and govern cross-border data flows. Each also used legislation with forms of extraterritorial or cross-border reach to influence conduct beyond a single domestic transaction. Their legal texts further converged around recognisable mechanisms: limitations on excessive personal-data collection, requirements connected with consent, user-facing controls, duties placed on service providers, and supervisory authority over consequential algorithmic services. The comparison could therefore ask what was specifically European about the individual-sovereignty promise and what depended on legal instruments increasingly used elsewhere.

China's discourse placed greater emphasis on Internet sovereignty, information security, public order, and state authority over infrastructure and information flows. European discourse more often framed digital sovereignty through fundamental rights, individual agency, fair markets, and strategic autonomy. Similar regulatory mechanisms could thus be embedded in profoundly different political projects. This was not evidence that the projects were normatively interchangeable. It demonstrated that the presence of a familiar legal mechanism does not, on its own, establish the political condition that a sovereignty discourse attributes to it.

The comparison also sharpened the role of industrial structure. China combined extensive regulatory control with major domestic platform and technology firms. The European Union possessed a powerful market and significant rule-making influence but fewer European providers at the scale of the dominant American and Chinese platforms considered in the analysis. The practical menu available to European users was consequently shaped not only by European rights but also by services, devices, and infrastructures developed elsewhere.

A right to choose has a different material quality when the available choices reproduce the same external dependency.

3.2.2 Boundaries of comparability

Four boundaries constrain the comparison. Similarity in statutory vocabulary does not imply similarity in constitutional protection. We did not rank the two systems as equivalent guardians of individual autonomy. European data-protection law operates within a system of fundamental rights and judicial remedies that cannot be inferred from an isolated comparison with the PIPL. State access to data and accountability of public authorities remain critical differentiators.

The instruments also differed in object and structure. The GDPR is not simply the European version of the PIPL, and the AI Act is not a direct counterpart to Chinese rules on recommendation algorithms or generative services. The comparison concerned functions and themes—such as user control, duties for providers, transparency, supervision, and data-transfer restrictions—rather than one-to-one doctrinal equivalence.

The analysis also preserved institutional variation within each polity. European implementation involves EU institutions, national authorities, courts, and organisations operating in varied markets. Chinese digital governance also involves interacting public and private interests and a changing relationship between industrial policy, security, and platform regulation. The comparison used the central legal and policy texts to examine a discourse; it did not eliminate institutional variation.

Causal effects lay outside the analysis. It could not determine that a particular rule produced a specific market structure or that industrial weakness resulted from the regulatory strategy. Its narrower claim was that the means visible in the European approach were insufficient for the objective as framed. Formal protections may be necessary and influential while remaining unable, on their own, to create the alternatives on which meaningful choice depends.

3.3 Comparative Legal Analysis

3.3.1 Regulatory architectures and the allocation of authority

The comparison reconstructed how rules in each polity related to the three realms of data, platforms, and infrastructure. In China, the Cybersecurity Law supplied a broad reference point connecting network operation, personal-information practices, and infrastructure security. Later rules developed more specific obligations for personal information, data classified as important or sensitive, cross-border transfers, recommendation algorithms, deep-synthesis services, and generative AI. The resulting architecture placed substantial supervisory and

filing powers in public authorities and linked the governance of algorithms and data to cybersecurity and social-order concerns [9, 163].

In the European Union, the architecture was more distributed. The GDPR established a horizontal framework for personal-data processing. The DGA addressed, among other matters, the reuse of certain protected public-sector data, data-intermediation services, and data altruism. The Data Act allocated rights and duties around access to and use of data, especially data generated by connected products and related services, and addressed unfair terms and switching between data-processing services. The AI Act established harmonised requirements and prohibitions organised around AI-related risks. Adjacent platform and cybersecurity instruments completed the broader policy landscape, though they were outside the main corpus [2, 3, 32, 56, 164].

This distributed architecture reflected the Union's sectoral competences and legislative development. It also made the sovereignty claim cumulative: individual empowerment was not located in one provision but assembled through rights, market rules, transparency duties, risk controls, and data-access mechanisms. The legal analysis therefore avoided asking whether any single regulation "created" sovereignty. It examined whether the package, understood alongside the Union's industrial position, could sustain the stronger political promise.

Public authority appeared differently in the two arrangements. The Chinese instruments under examination assigned administrative bodies extensive power to supervise sensitive algorithmic services, require filings, evaluate certain cross-border transfers, and connect provider duties with national cybersecurity objectives. European instruments divided authority among supervisory bodies, market regulators, national competent authorities, the Commission, courts, and the individuals or organisations entitled to exercise specific rights. This institutional difference affects both accountability and speed of intervention. It also shows why the comparison cannot be reduced to a table of similar clauses.

3.3.2 Platforms, algorithms, and the position of the user

On the Chinese side, the rules on algorithmic recommendation services combined administrative supervision with user-facing controls. Providers of services with public-opinion properties or social-mobilisation capacity were subject to filing and assessment arrangements. Users were given options concerning recommendation services based on personal characteristics and the tags used to generate recommendations. Related measures addressed the labelling of synthetic content, duties concerning the use of biometric information, and responsibilities of providers and users of generative AI services. These mechanisms linked algorithmic governance to both individual interests and state oversight [9, 157, 161, 162].

The European AI Act approached the problem through a horizontal, risk-based framework rather than service-specific rules directed primarily at recommendation platforms. For systems classified as high-risk, the Act requires risk management, data governance, docu-

mentation, transparency to deployers, human oversight, and post-market processes, with duties distributed among providers, deployers, and other actors. Its human-oversight provisions are relevant to individual autonomy because they require the system to be designed and provided in a way that enables assigned persons, as appropriate, to understand limitations, monitor operation, recognise automation bias, interpret outputs, disregard or reverse outputs, and intervene or stop the system [2].

These requirements do not grant every affected person direct control over every AI system. They allocate responsibilities within a regulated value chain and seek to preserve human agency where specified risks justify intervention. This distinction is important for the sovereignty claim. The "individual" in individual digital sovereignty may be a data subject, user, consumer, professional operator, citizen, or person affected by an automated process. Each role has different information, authority, and remedies. Treating them as a single sovereign user would obscure the organisational arrangements through which oversight is actually exercised.

The comparison found both convergence and divergence. Both systems imposed obligations on providers and sought to constrain algorithmic systems capable of influencing individuals. Both connected transparency or user choice with public supervision. Yet the objects, institutional settings, and underlying political purposes differed. European law foregrounded risk, fundamental rights, and accountable market placement; the Chinese rules examined were more directly integrated with state supervision of platforms, information, and social mobilisation. The presence of user-facing controls in both systems undermined any claim that such controls alone made the European project uniquely sovereignty-enhancing.

The comparison also exposed a limit of transparency. Information about a system can support agency, but only if the recipient can interpret and use it. Instructions provided to a deployer do not necessarily make an affected person capable of contesting a decision. Disclosure of an algorithmic filing to an authority does not necessarily create public understanding. A choice to disable personalisation may still leave the user dependent on the same platform. Transparency, oversight, and choice are therefore components of capacity, not substitutes for the economic and institutional conditions in which a choice becomes effective.

3.3.3 Personal data, data flows, and extraterritorial influence

The data comparison centred on the GDPR and PIPL before situating them within the wider European and Chinese regimes. At the level of selected legal mechanisms, the analysis identified notable points of resemblance. Both laws extend in defined circumstances to processing or handling connected with persons located within their respective territories even when the relevant actor is outside. Both include purpose-related constraints, data-minimisation principles, and forms of individual control over personal-information handling. Both place duties on the organisations that process or handle personal information. These similarities made data protection an especially useful field in which to test claims about a

uniquely European path [3, 9, 155].

The resemblance stopped well short of equivalence. The GDPR is embedded in the EU system of fundamental rights, independent supervision, and judicial protection, and it specifies a developed set of data-subject rights and controller and processor duties. European law also includes the right to erasure and data portability within a broader doctrinal and institutional history. The PIPL's relationship to state authorities, national security, and available forms of redress differs materially. Cross-border transfer mechanisms in the two systems likewise express different institutional priorities: European law asks, through several transfer routes, how protection travels with personal data, while the Chinese measures under examination place strong emphasis on security assessment and state supervision. Comparison of selected clauses cannot erase these differences.

The European data package also pursued objectives beyond personal-data protection. The DGA established procedures intended to support trusted forms of reuse and intermediation, including routes for the reuse of certain public-sector data protected on grounds other than ordinary open-data availability. The Data Act sought to make more data generated by connected products and related services accessible to users and, on request and under applicable conditions, third parties. It also addressed business-to-government access in exceptional circumstances and contractual and technical barriers in data markets. These instruments attempted to reconcile the protective logic of data regulation with an economic objective: enabling more data to circulate and generate value [32, 56].

That reconciliation matters for digital sovereignty. Protection without usable alternatives may leave individuals and organisations dependent on incumbent services; circulation without protection may transfer power to actors best positioned to aggregate and exploit data. The European strategy sought to create a data market governed by rights and trust. Its success therefore depended not only on the legal possibility of sharing but on infrastructure, interoperability, business models, enforcement, and the capacity of European organisations to participate. Those dependencies were acknowledged in the policy context but could not be assessed from the legal texts alone.

We also examined extraterritorial influence as a source of European power. Access to the EU market gives global firms reasons to adapt practices to European law, and the GDPR has affected privacy policies and compliance structures beyond the Union [165, 166]. Regulatory reach is thus a genuine capacity. Yet regulatory influence should not be conflated with the availability of European-controlled alternatives. A provider may comply with European rules while the underlying service, cloud stack, model ecosystem, or device market remains concentrated elsewhere. Individuals can receive stronger protections without acquiring a materially broader choice of infrastructure or provider.

Cross-border data flows make this tension particularly visible. The Union has an interest in data moving across borders for trade, research, and innovation while protecting fundamental rights and strategically sensitive information. China likewise regulates outbound

transfers, though through a different political and legal framework. Neither absolute localisation nor unrestricted movement supplies a complete sovereignty model. Localisation can increase control over one physical or jurisdictional aspect while leaving software, hardware, expertise, and contractual dependencies unchanged. Openness can support innovation while creating exposure to external rules or infrastructures. The relevant question is whether dependencies are governed, alternatives exist, and affected actors can challenge the terms under which data move.

3.4 Comparative Findings

Three connected findings emerged. Individual empowerment was a distinctive and recurrent element of the European *discourse* on digital sovereignty, but several of the legal mechanisms associated with that discourse were not uniquely European. The Chinese laws and measures examined also constrained private data handlers and algorithmic services, recognised specified user interests, and extended public oversight. The surrounding constitutional systems and political purposes differed, yet the overlap weakened an argument that the existence of those mechanisms by itself realised a specifically European form of sovereignty.

The Union's market and regulatory capacity were substantial but incomplete. European rules could influence global providers that wished to serve the internal market. They could create enforceable protections and reshape business processes. They could not, through command alone, produce a competitive European platform, cloud service, processor supply, or AI ecosystem. Where users continued to depend on a small group of external providers, the practical exercise of choice remained constrained even when legal protection improved.

The remaining gap could not be captured by asking whether European legislation was strong or weak. The central issue was fit between instrument and objective. Rights-based regulation is well suited to setting conditions of lawful conduct and protecting individuals against specified harms. It is not by itself an industrial policy, an infrastructure investment, a procurement capability, or a workforce-development programme. The promise of individual digital sovereignty became false when success in the regulatory domain was treated as sufficient evidence of autonomy across those other domains [9].

We therefore answered the research question cautiously: individual digital sovereignty could not be attained through a purely regulatory effort unaccompanied by strong material and industrial capacity. This formulation leaves room for regulation to be necessary, effective within its domain, and globally influential. It also leaves room for later infrastructure, investment, and data policies to alter the diagnosis. What it rejects is the inference from formal right to realised sovereignty.

3.4.1 Results and Evidentiary Boundaries

The phrase "false promise" is easily misunderstood if separated from the method used here. It does not mean that every European commitment to autonomy was made in bad faith. Nor does it mean that GDPR rights were fictitious, that European regulators had no influence, or that Chinese and European protections were equivalent. It identifies a structural overclaim: an individual cannot obtain meaningful sovereignty from regulatory protection alone where the material set of available choices remains externally constrained.

The result should not be generalised beyond the fields examined. The corpus comprised selected data and AI rules; the analysis did not evaluate every dimension of digital policy, defence, telecommunications, semiconductors, research, competition, or public procurement. It did not assess whether the Union might exercise greater influence in another technological sector. It also did not establish the future effects of the Data Act, DGA, or AI Act. Those instruments were too recent for such a conclusion when the research was performed.

The legal analysis likewise did not measure enforcement quality. A statute can be interpreted and applied differently over time, and rights may have uneven practical accessibility across cases and Member States. The comparison was capable of showing what the selected texts required, permitted, or empowered and how those elements related to the policy discourse. It could not determine whether every duty was observed, whether every remedy was effective, or whether individuals experienced greater autonomy as a result.

Most importantly for this thesis, Common European Data Spaces were outside the scope of the comparison. European data-sharing and infrastructure initiatives appeared only as part of the surrounding landscape. The analysis did not inspect a data-space deployment, test interoperability, examine participant behaviour, compare sectoral governance models, or analyse whether a CEDS supplied data used in an AI project. Any proposition below about the potential role of CEDS is therefore an extension grounded in later institutional documents and the literature on data spaces, not an empirical result of the 2024 comparison.

That boundary also clarifies the basis for the 2026 update. The analysis identified infrastructure as an under-examined dimension and called for attention to the material conditions of sovereignty. Between its research cut-off and July 2026, the European policy agenda increasingly connected data spaces, computing facilities, cloud capacity, skills, and AI adoption. The update examines how far that policy development responds to the diagnosed capacity gap without treating announced investment or deployed components as proof of achieved sovereignty.

The institutional position described in this section is a documentary update with a cut-off date of 20 July 2026. It is not part of the evidence used in the 2024 comparison. By this date, the GDPR was a mature component of the European legal order, the DGA and Data Act had moved from legislative novelty towards application, and the AI Act was progressing through its staged implementation. At the same time, Commission policy documents described the AI challenge in terms that extended beyond regulation. The AI Continent Action Plan

organised intervention around computing infrastructure, access to data, adoption, skills, and implementation of the AI Act. The later Data Union Strategy explicitly linked data availability and data spaces to AI development. This broader portfolio is significant because it recognises, at policy level, the same basic limitation identified through the comparison: rule-making must be accompanied by resources and institutions capable of acting under the rules [23, 24].

This recognition does not show that the capacity gap has been closed. Policy communications state objectives, identify initiatives, and report institutional milestones. They do not independently establish the amount of usable compute available to a particular organisation, the cost of access, the performance of a data-space service, the competitive position of European providers, or the outcomes of an AI deployment. The update therefore uses careful verbs. The Commission *proposes*, *funds*, *reports*, or *plans*; the thesis does not infer from those acts that sovereignty has been achieved. The evidentiary question is more modest: has the policy architecture begun to address data, cloud, compute, and adoption as distinct capacities rather than assuming that regulation alone will produce them?

3.4.2 Data, cloud, and compute in the AI competition

AI makes the interdependence of regulatory and material capacity unusually clear. Developing and deploying a system may require domain-relevant data, computing resources, models and software, energy and connectivity, specialist skills, capital, and an organisation able to integrate these inputs into a service. Their relative importance varies by use case. A public body fine-tuning a modest model for a bounded administrative task is not in the same position as a firm training a frontier general-purpose model. It would therefore be misleading to reduce the "AI race" to a single ranking or to assume that every European actor must reproduce the same stack. The relevant sovereignty question is whether strategically important actors have credible, governed options for obtaining the resources their task requires.

The AI Continent Action Plan is a material change of emphasis in this respect. It treats large-scale computing, cloud and data-centre capacity, data access, sectoral adoption, talent, and regulatory implementation as connected pillars. By April 2026, the Commission reported nineteen AI factories associated with the European supercomputing network and thirteen AI Factory antennas intended to extend regional access, while AI gigafactories remained a developing initiative. These are official progress claims rather than an independent performance evaluation. Their relevance lies in the architecture they describe: computing is treated as a capacity to be provisioned and made accessible, not as an automatic consequence of a harmonised legal market [167].

The distinction between an AI factory and a data space is important. A computing facility supplies or coordinates computational resources and related support. A data space organises governed data sharing among participants. A cloud or data-centre policy concerns another layer again: the availability and operation of infrastructure on which storage, ex-

change, training, and deployment may depend. Connecting these components can reduce friction, but their functions should not be collapsed. A catalogue entry cannot train a model; a supercomputer does not establish the right to use a dataset; a cloud contract does not guarantee semantic interoperability; and an AI Act conformity process does not create the organisational purpose or public legitimacy of a project.

The Data Union Strategy makes the data side of this portfolio more explicit. It identifies access to high-quality data for AI, simplification or clarification of the data-rule environment, and the Union's position in international data flows as related priorities. It also proposes data labs as service facilities connecting data holders and data spaces with AI ecosystems. In this account, the problem is no longer merely that data should be legally shareable. Data must be discoverable, prepared, curated, and usable in an environment capable of preserving applicable protections and conditions [24].

This policy shift supports, but does not settle, the argument of the thesis. It supports it because it treats data quality, interoperability, infrastructure, expertise, and legal support as complementary. It does not settle it because the effectiveness of those arrangements depends on implementation: which datasets are actually available; who can access them; what quality and provenance information accompanies them; how restrictions are interpreted; what services are affordable; and whether organisations can integrate the data into a responsible workflow. These are questions for evaluation at the level of a data space and a concrete use, not conclusions that can be read from the strategy itself.

3.4.3 Common European Data Spaces as a partial response

CEDS emerged from the 2020 European strategy for data as sectoral or domain-oriented environments intended to make more data available across the Union while allowing data holders and other participants to retain control under common rules [6, 7]. They are not, as a class, one central EU database. The data-space model is federated and socio-technical: multiple participants use shared governance arrangements, semantic resources, catalogues, identity and trust mechanisms, contractual or policy instruments, and technical components to discover and exchange data. Implementations vary across domains and maturity levels, so "CEDS" denotes a policy family rather than a single uniform operational system [4, 5, 116].

This model can contribute to digital sovereignty in at least four bounded ways. Catalogues can improve discoverability by allowing organisations that would otherwise rely on private knowledge or bilateral contacts to find described data offerings. Common models, vocabularies, interfaces, and exchange procedures can improve interoperability. Participant roles, policies, contracts, and trust services can structure control over access and transfer. Collectively developed conditions can support an ecosystem in which European public and private actors share data without surrendering every exchange to a vertically integrated platform.

Each contribution is conditional. Discoverability depends on records that are maintained and sufficiently informative. Common vocabularies support semantic exchange only when participants implement them consistently and preserve the relevant context. Policy languages can structure control only within architectures able to interpret and enforce them at defined points. A federation creates alternatives only if enough data holders, users, intermediaries, and service providers participate on viable terms. The Data Spaces Support Centre Blueprint is valuable here because it presents governance, business, legal, semantic, and technical building blocks as connected parts of an ecosystem rather than as a single technical product [116].

The Data Act gives some of these concerns legal salience. Article 33 sets essential requirements for interoperability concerning participants in data spaces that offer data or data services to other participants. Among other matters, it addresses descriptions of content, restrictions, licences, collection methodology, quality, uncertainty, data structures, vocabularies, and access interfaces [32, art. 33]. The provision illustrates a more general relationship: a legal requirement can create an information need, while standards and metadata elements may provide part of its representation. Neither the statutory requirement nor the metadata record alone proves that the underlying data are accurate, that every downstream use is authorised, or that a particular AI project is acceptable.

From the perspective of the 2024 diagnosis, CEDS therefore represent more than another layer of regulation. They seek to create shared operational capacity for data exchange. They can reduce information and coordination costs, support European providers and public institutions, and keep conditions attached to data across organisational boundaries. In that limited sense they respond to the gap between a right to act and the means to act. Their contribution remains partial because data exchange is only one component of the industrial and infrastructural setting in which digital choice is exercised.

3.4.4 Limits of Data Spaces

A careful account of CEDS must resist a new version of the same false inference criticised above. The earlier inference moved from regulation to sovereignty. The new one would move from the existence of a data-space initiative to usable and legitimate AI capacity. Several steps intervene.

Catalogue availability leaves authority unresolved. A dataset may be visible in a catalogue while access remains restricted, or it may be accessible for one purpose without being available for AI training. Personal-data rules, intellectual-property rights, trade secrets, contractual conditions, sectoral legislation, and the terms of a permit or access decision may all affect the proposed operation. An earlier data-space transaction cannot be treated as a general licence for every downstream use. The GDPR continues to apply where personal data fall within its scope, and the AI Act may add requirements according to the system, actor, and lifecycle stage [2, 3].

Governed access still requires a separate assessment of technical fitness. Data obtained through a governed route may be incomplete, biased, poorly documented, outdated, incompatible with the intended method, or unsuitable for the population on which a model will operate. Metadata can communicate reviewed information about provenance, quality, format, and restrictions. It cannot make the dataset adequate by declaration. Fitness is a contextual judgment supported by evidence and testing.

Computing capacity remains a separate requirement. A data space can coordinate access to valuable sectoral data without supplying the processors, memory, storage throughput, cloud environment, or energy needed for a particular training or inference workload. Connections with AI factories and data labs may help bridge that divide, but the components remain institutionally and technically distinct. Their integration has to be designed and evaluated.

Organisational competence forms another condition. Even where data and compute are available, an organisation needs a defensible purpose, procurement and security capability, people able to assess legal and ethical implications, domain expertise, data and model governance, mechanisms for stakeholder involvement, and authority to stop or revise a project. These capacities are especially important in public administration, where an AI system may affect access to services, resource allocation, inspection, planning, or communication with citizens.

European location or participation can coexist with external dependency. A European data space may rely on components, cloud services, processors, or software controlled by non-European firms. Such dependence is not automatically undesirable; openness and international supply chains can support resilience and innovation. The sovereignty question is whether dependencies are understood, governed, substitutable where necessary, and compatible with public objectives. Complete autarky is neither the operational model of CEDS nor a useful benchmark for their success.

These limits do not diminish the value of data spaces. They locate that value accurately. CEDS can provide a governed data layer within a larger AI ecosystem. They can make the conditions, participants, and evidence of exchange more visible. They can help public and private actors collaborate without requiring all data to be transferred into one platform. They cannot determine the legal validity, ethical legitimacy, technical quality, or social desirability of every project built with the exchanged data.

3.5 Three Capacities in Data Governance

The original synthesis developed in this thesis connects the comparative diagnosis and the 2026 update through three capacities. European institutions seek to shape the strategic and infrastructural conditions under which data, cloud services, computing facilities, and AI ecosystems are available. Public administrations and other participants need organisational

and deliberative capacity to formulate purposes, identify affected groups, allocate responsibility, and revisit assumptions throughout a project. At the informational and technical level, reviewed legal and governance conditions must remain connected to datasets and evidence in forms that can travel across systems. These capacities are complementary; they are not stages in a validated end-to-end compliance system.

The connection can be expressed as a series of questions rather than a claim of technological completion. Does an organisation have a credible route to the data, infrastructure, and services needed for its task? Can the relevant people identify why the project should exist, whose interests it affects, and who has authority to decide? Can the legal source, reviewed requirement, metadata need, standard element, and representational limit be traced? A positive answer to one question cannot stand in for the others. Material capacity without deliberation can scale a harmful project; deliberation without usable information can become detached from implementation; technical traceability without institutional authority can record a decision that nobody is empowered to challenge.

This formulation also clarifies the meaning of human-centred compliance used in the thesis. It denotes a process in which tools and machine-readable records support people who hold different forms of expertise and responsibility. It does not denote automated certification that a system or dataset is lawful. Legal interpretation, factual assessment, ethical judgment, and public accountability remain human and institutional tasks. Technical artifacts help make their inputs, outputs, and changes inspectable.

CEDS provide a particularly revealing setting because the three capacities meet there. A data space is part of a strategic European effort to create data infrastructure. Its operation depends on organisations capable of joining, publishing or requesting data, interpreting rulebooks, maintaining metadata, and managing access. Its interoperability depends on shared representations and technical services. If any of these elements is absent, the formal possibility of exchange may remain unused or be used without sufficient context.

The synthesis should not be read as a finding that CEDS already integrate these capacities successfully. The thesis has not performed a cross-data-space maturity assessment, measured adoption, or followed a project longitudinally from catalogue discovery through model deployment. It develops one part of the relationship conceptually and illustrates its intended operation through two bounded micro-applications to a Data Act provision. Organisational deliberation and the conceptual framework perform different functions and rest on different evidentiary bases.

3.5.1 Public Administrations in Data Governance

Public administrations occupy more than one position in this setting. They may hold or produce data that are made available through open-data channels, protected-reuse procedures, sectoral data spaces, or specific agreements. They may request data from other public or private participants. They may procure or develop AI systems, operate shared services, su-

pervise regulated actors, and represent the public interest in a data-space governance body. The roles can coexist within one institution but should not be conflated.

As a data provider, an administration needs to know what it is making available, under which authority and conditions, with what quality and provenance information, and who will maintain the description. As a data user, it needs to determine whether access covers the intended operation, whether the data are fit for the task, what restrictions persist, and what evidence must be retained. As the designer or deployer of an AI service, it needs to connect the dataset decision to a public purpose, system requirements, affected communities, procurement and technical choices, human oversight, monitoring, and avenues for contestation.

This dual position makes public administration an important beneficiary of CEDS and a demanding test of their governance. Better access to cross-sectoral data can support research, planning, service delivery, and policy evaluation. At the same time, public bodies exercise authority and may affect people who cannot simply opt out of an administrative relationship. The availability of data and compute must therefore be accompanied by procedures that expose assumptions and assign responsibility before technical development narrows the range of feasible choices.

The sovereignty argument is visible at this operational scale. A public body is not autonomous merely because a regulation permits access or because a data space offers a technical endpoint. It needs internal or reliably governed external capability to understand the data, negotiate or comply with conditions, evaluate suppliers, and maintain the service. Conversely, building all technology internally is not a necessary condition of autonomy. A well-governed dependency can preserve meaningful choice when terms are transparent, alternatives or exit routes are credible, expertise is retained, and public authority can challenge the provider. Sovereignty is therefore expressed through the quality and contestability of relationships, not only through ownership.

3.6 Organisational Capacity and Deliberation

The analysis now reaches a problem that infrastructure policy cannot answer. Suppose that a public administration can discover a relevant dataset in a CEDS, has a plausible access route, and can obtain computing support. It still must decide whether an AI project is appropriate, which problem it is meant to solve, what data and analytical methods are defensible, whose knowledge is missing, what social and territorial effects may arise, and which requirements must shape design and procurement. These questions arise before development, change as the project develops, and return when the system is reviewed.

Ordinary project documents often distribute those questions across legal advice, technical specifications, procurement files, meeting notes, risk assessments, and informal knowledge. That fragmentation can reproduce the capacity gap in organisational form. Legal rights and

technical resources exist, but the actors who must use them do not share a stable setting in which to make assumptions visible and revisit them. A final compliance check is too late if the decisive choices about purpose, data, architecture, and affected groups were fixed much earlier.

Administrative capacity therefore also requires a structured setting in which technical and non-technical participants can discuss data, algorithms, analytical methods, sociocultural considerations, and functional requirements. The AI Ethics Canvas provides one such iterative setting during project framing, design, development, and review. Its deliberative and documentary function is to expose dependencies and create information touchpoints at which participants identify decisions, uncertainties, responsibilities, and evidence requiring further work [133].

The Canvas addresses the organisational dimension of the capacity gap: it structures multidisciplinary deliberation, exposes missing evidence and responsibilities, and creates opportunities to revisit project assumptions. It does not supply industrial alternatives, infrastructure, or computing capacity. Its proposed contribution concerns how an administration frames and reviews a project within the material conditions available to it; the formative evidence examined in Chapter 4 does not establish that it closes the wider capacity gap.

These touchpoints can enrich project records and support later compliance and assessment activities, but they do not establish legal compliance, prove technical quality, or guarantee representative participation. The formative evidence concerning the Canvas is distinct from the legal and policy evidence used in the sovereignty analysis. The conceptual framework has a separate function and evidentiary basis: it specifies how the possible technical representation of selected legal concepts and requirements should be examined. The two functions are autonomous and do not form a pipeline; competent human and institutional judgment remains necessary in both.

3.7 Conclusion

The comparative analysis established a carefully delimited proposition: the European promise of individual digital sovereignty could not be delivered by regulatory effort alone while meaningful choice remained constrained by limited industrial and infrastructural alternatives. The EU–China comparison showed that several regulatory mechanisms associated with individual empowerment were not uniquely European, while also preserving the profound constitutional and institutional differences between the two systems. The result concerned the fit between a promise and its means; it was not a finding that regulation lacked value.

The 2026 update shows a broader European capacity agenda in formation. Data spaces, data labs, AI factories, cloud and computing initiatives, skills, and adoption policies address

material and organisational conditions identified as missing above. CEDS are a particularly important partial response because they seek to create governed, interoperable routes through which European actors can share data. Their existence does not prove that data are usable, that downstream AI training is authorised, or that external technological dependencies have disappeared.

For public administrations, the remaining task is operational. They must be able to turn access and infrastructure into projects whose purposes, assumptions, responsibilities, evidence, and effects can be examined and revised. This operational need concerns the structured deliberation through which a public-sector AI project is framed, documented, and reviewed.

Chapter 4

The AI Ethics Canvas: Designing and Reviewing Public-Sector AI

Publication disclosure. This chapter is based on work accepted for publication as: Riccardo Nanni, Pietro Giovanni Bizzaro, Albana Celepija, Munazza Usmani and Maurizio Napolitano, "AI Ethics Canvas: A Graphical Tool to Design and Deploy Ethical Artificial Intelligence for Public Administrations," *Data & Policy* (forthcoming, 2026).

The canonical artifact discussed in this chapter is AI Ethics Canvas release 1.0.0, available at `https://github.com/dclfbk/canvasaiethics` under the Creative Commons Attribution–ShareAlike 4.0 International licence (CC BY-SA 4.0) [33]. Earlier versions, including the version used during the 2024 development activities, are treated as stages in the artifact's evolution and not as the canonical version.

4.1 From regulatory principles to an organisational practice

Public administrations are increasingly expected to decide whether, where, and how artificial intelligence can support policy formation, service delivery, and internal administration. These decisions occur under conditions that differ from the development of a stand-alone commercial product. A public authority must relate the proposed system to a statutory mandate, public values, administrative procedure, procurement arrangements, the rights and expectations of citizens, and the responsibilities of officials who remain answerable for the resulting action. It must also work across organisational boundaries. Data may be main-

tained by one department, software developed by a public company or an external supplier, models adapted by researchers, and outputs interpreted by domain officials who did not participate in development. The ethical problem is therefore not confined to a model or a dataset. It arises in the relationships among a purpose, an institution, several professional communities, affected people, and a changing technical system.

The European framework provides an extensive set of principles and legal obligations with which these organisations must engage. The General Data Protection Regulation (GDPR), the Data Governance Act, the Data Act, the Open Data Directive, and the Artificial Intelligence Act govern different parts of the institutional and technical setting [2, 3, 32, 56, 168]. Their scopes, addressees, concepts, and effects are not interchangeable. Ethical frameworks add requirements and aspirations concerning human agency, fairness, transparency, privacy, technical robustness, social well-being, and accountability. Yet the presence of rules and principles does not by itself tell a multidisciplinary project team what to discuss first, who should answer a particular question, which assumptions remain unresolved, or how a concern should be carried into a design decision.

This gap between a principle and an organisational practice is well established. Reviews of AI-ethics methods show a heterogeneous field of guidelines, checklists, impact assessments, documentation formats, and participatory tools, often with limited evidence about their sustained use and downstream effects [128, 132]. Principles require contextual interpretation, institutional authority, and resources; they cannot guarantee ethical outcomes on their own [129]. A technical abstraction may also hide the social system in which an apparently neutral model is embedded [130]. For a public administration, identifying the applicable principle is only the starting point; the practical task is to enable people with different vocabularies and responsibilities to make its implications discussable before a design hardens into infrastructure.

The AI Ethics Canvas responds to this organisational problem. It is a visual, question-based artifact intended to structure deliberation among the people who propose, build, adapt, procure, and use an AI system in a public-administration setting [133]. It does not convert legislation into a checklist whose completion proves conformity. Nor does it calculate whether a model is sufficiently accurate, fair, robust, secure, or explainable. Its narrower function is to create a common working surface on which participants can identify the proposed use, expose assumptions, assign questions to relevant professional perspectives, record uncertainties, and turn ethical and legal concerns into matters that a project must address.

The Canvas operates at the organisational level. It supports a public administration in framing an AI project, involving the relevant professional perspectives, surfacing uncertainties, and identifying information that requires further assessment. It does not determine whether a project is lawful or ethically acceptable.

4.2 A Canvas for Public-Sector AI

4.2.1 AI Adoption as a Socio-Technical Decision

An AI project in a public administration rarely begins with a fully specified technical requirement. It may begin with a policy objective, an operational bottleneck, a new source of data, a supplier proposal, or a general expectation that AI should improve efficiency. Before selecting a model, the administration must decide what problem is being addressed and whether an AI system is necessary or proportionate to that problem. It must identify who will use the output, whether the output supports or determines a decision, which people may be affected, what information the system requires, and what institutional act gives the project its authority. An early ethical discussion is therefore also a requirements discussion.

The organisational setting complicates that discussion. Expertise about an administrative domain does not entail expertise in machine learning. Conversely, a technologist may understand training and deployment while lacking authority to determine the public purpose, acceptable error distribution, or consequences of a false result. Suppliers may control implementation details that are difficult for the administration to inspect. Officials may use generative AI tools informally while conceiving of AI mainly through chatbots, even though predictive, classificatory, vision, and decision-support systems raise different questions. The AIxPA work confirmed this asymmetry: technical development was substantially supported by organisations outside the core provincial administration, while officials retained domain knowledge and responsibility for administrative use [133].

These are not merely communication difficulties. They shape which risks become visible. If a data scientist frames the problem alone, organisational history or unequal effects on service users may remain outside the model boundary. If a domain official frames it alone, limitations of the data or uncertainty of the output may be underestimated. If a procurement unit receives only a high-level statement that a tool should be "ethical", a supplier cannot translate that statement into testable requirements without further choices. Research on responsible-AI practice similarly shows that accountability is distributed across roles and that practitioners need coordination, authority, monitoring, and escalation routes rather than principles detached from workflow [66, 131].

The Canvas makes this distribution visible by requiring several perspectives to work on the same use case. Its graphic format lowers the threshold for entering a technical conversation, while the questions prevent a visually accessible artifact from becoming an empty brainstorming exercise. A question may reveal that the answer is unknown, contested, dependent on evidence, or assigned to no one. That is a useful result. The purpose of the session is not to produce affirmative answers but to expose what the organisation would otherwise discover late, after procurement, training, integration, or deployment.

4.2.2 Canvases and Structured Deliberation

Canvas methods were popularised by the Business Model Canvas as visual ways of organising interdependent elements of a business model [169]. Their attraction lies in spatial compression: participants can see several components together, move between them, and record connections without first producing a long report. Later canvases applied this logic to service design, data governance, open data, public value, and responsible technology. The AI Ethics Canvas belongs to this design tradition and can be compared with tools such as the Data Ethics Canvas and public-AI canvases [133, 170, 171].

The difference from a compliance checklist is fundamental. A checklist commonly seeks confirmation that defined actions have been performed. The Canvas instead asks questions whose answers may change the design, require investigation, or remain open. Some questions concern legal obligations, but the Canvas does not determine their applicability or interpret them conclusively. Some concern ethical values that cannot be reduced to binary status. Others concern technical performance and must be answered through testing rather than discussion. The artifact therefore supports a conversation that can generate requirements and evidence needs; it does not replace the specialised activities to which those needs refer.

4.2.3 Value-sensitive design and action research

The design approach combined value-sensitive design with action research. Value-sensitive design treats human values as relevant from the inception of a technology rather than as a final constraint applied to a completed system. Conceptual investigation identifies values and possible tensions; empirical investigation examines how stakeholders understand and experience them; technical investigation considers how design choices may support or undermine them [172]. For the Canvas, this orientation meant that fairness, human agency, transparency, privacy, responsibility, and public value were not placed in a final "ethics" box. They were distributed across questions concerning the data, the algorithm, interpretation of outputs, institutional context, and functional requirements.

Action research supplied the relationship with practice. The researchers did not observe a public administration from a distance and then deliver a finished instrument. They worked within AIxPA, a project involving the Autonomous Province of Trento (Provincia autonoma di Trento, PAT), to develop an artifact that could be criticised and revised by the people expected to encounter AI in their work. This orientation is interventionist and iterative: a provisional artifact structures discussion; participants reveal problems in its content, vocabulary, and usability; the artifact is revised; later encounters test whether those revisions improve the conversation [173]. Co-design research on responsible-AI instruments likewise emphasises alignment with roles, organisational culture, infrastructure, and existing work practices [65].

Action research and value-sensitive design informed the participatory and iterative orien-

tation of the development process. The word "validation" was used during the development process to describe these cycles. In this thesis, the more precise expression is *formative evaluation*. The interviews and workshop examined clarity, practical relevance, and the kinds of reflection the Canvas prompted. They did not establish construct validity through a controlled instrument, compare project outcomes with and without the Canvas, or follow systems longitudinally after deployment. This terminological discipline preserves the value of the empirical work while avoiding a claim that the resulting AI systems were thereby ethical, lawful, or technically sound.

4.3 Research and design process

4.3.1 Institutional Setting

The AIxPA project explored how AI could support the provincial public-administration system. PAT offered a useful action-research setting because it combines significant administrative autonomy and domain responsibilities with the organisational constraints familiar to many public bodies: heterogeneous departments, legacy procedures, limited in-house AI-development capacity, and reliance on public or research partners for specialised technical work. The object of intervention was not one finished algorithm. It was the organisational process through which candidate AI uses could be framed, discussed, documented, and translated into requirements.

The primary participants were officials, directors, developers, and researchers rather than a representative sample of residents. That choice followed the initial purpose of the artifact: to support teams responsible for proposing and implementing AI for administrative work. It does not imply that affected communities are unimportant. On the contrary, the socio-cultural part of the Canvas asks project teams to identify differentiated impacts and consider consultation. The design study nevertheless evaluated the artifact principally with organisational users. Future evaluations must involve people affected by the relevant service when their knowledge is necessary to identify harms, exclusions, or competing understandings of public value.

The unit around which the Canvas is completed is a defined use case. Participants should state the administrative problem, intended benefit, contemplated users, affected people, operational environment, data sources, and anticipated decision before answering the thematic questions. This preliminary step prevents "AI adoption" from becoming the objective in itself. A team may conclude that a rules-based process, improved data management, staff allocation, or a non-technical intervention better serves the stated purpose. Recording that conclusion is consistent with the Canvas. The tool is intended to discipline an adoption decision, not to maximise the number of AI projects.

The three macro-roles reflect the working classification adopted within AIxPA to distinguish contributions to the design, implementation, and administrative use of AI systems.

This project-level distinction informed the Canvas's role legend and intended working group. It was not developed as an exhaustive classification of the actors involved in AI or data governance. One person may contribute through several macro-roles, and scientific or engineering functions may be supplied by organisations outside the administration.

The workshop illustrates this distinction. Across the Family, Extreme Events, and Overtourism scenarios, researchers, technical personnel, and public officials contributed to completing the Canvas, while facilitation and observation were recorded as separate workshop functions. These participation functions should not be confused with legal roles or institutional decision authority.

Other actors, including data stewards, Data Protection Officers, legal advisers, procurement and security personnel, and representatives of affected groups, need not be assigned to one of the three macro-roles. Their contributions may be obtained through targeted consultation, participation, or specialist review according to the questions raised. When the Canvas is used beyond AIxPA, its role classification may be supplemented or adapted to the concrete setting. Such adaptations should identify the working participants, distinguish expertise from decision authority, and record relevant perspectives that remain absent. The original categories provide an entry point for deliberation rather than a fixed boundary on whose knowledge may inform it.

4.3.2 Documentary and legal foundations

The initial design drew on two complementary analyses. One reviewed nine influential ethical and practical frameworks: European trustworthy-AI guidance [174]; Italian deontological rules for processing for statistical or scientific purposes [175]; European data-protection and privacy ethics guidance [176]; European research-ethics guidance [177]; the Open Data Institute's Data Ethics Canvas [170]; the Toronto Declaration [178]; an algorithmic-impact-assessment case study from the Ada Lovelace Institute [179]; Ethics as a Service [180]; and the NIST AI Risk Management Framework [181]. The objective was not to merge them into one universal theory. It was to identify recurring areas that a public-administration project team should be prompted to discuss [133].

A complementary legal analysis examined five EU legal instruments relevant to the project context: the GDPR, Open Data Directive, Data Governance Act, Data Act, and AI Act. These sources address different objects and cannot be collapsed into generic "EU compliance". The analysis was therefore used to formulate prompts, not to assert that every provision applies to every project. For example, a question about personal data can reveal the need to determine purpose, legal basis, necessity, data minimisation, or impact-assessment duties under the GDPR. A question about AI-system risk can prompt a separate legal classification under the AI Act. Questions about data access and reuse can identify issues arising under data and open-data instruments. The legally competent actor must still select the current source, establish applicability, interpret the requirement, and document

the decision.

Together, the two analyses yielded the five thematic areas: data, algorithms, methods of analysis, socio-cultural elements, and functional requirements. This organisation deliberately moves from resources and technical construction to use, institutional context, and implementation. It also avoids placing law in an isolated section. Legal and ethical concerns recur throughout the lifecycle: authority to process data affects the data area; risk classification and prohibited practices affect the algorithm area; explanation and human oversight affect methods of analysis; equality and public communication affect socio-cultural elements; and documentation, procurement, security, and monitoring affect functional requirements.

4.3.3 Iterative drafting and the 2024 artifact

The documentary analysis produced an initial draft, which was reviewed by key project informants before the interview cycle. The visual grammar assigned questions to three macro-roles by colour and repeated the distinction through shapes, reducing reliance on colour perception alone. Introductory pages explained how to use the artifact. The English version remained broadly applicable to EU public administrations, while the Italian materials provided language and contextual explanation suited to the local administration.

The version circulating during 2024 should be read as a research instrument in this iterative process. It supplied a concrete object that participants could criticise; it was not the final authority against which later versions are judged. Interviews led to changes in wording, explanatory material, and links with documentation practices. Later project work and the workshop generated further observations. Release 1.0.0 consolidates that evolution and is the reference artifact for citation and reuse [33, 182].

4.3.4 Ethics, confidentiality, and analytical boundaries

The interview material was pseudonymised. Role categories, dates, organisational profiles, and the interview questionnaire are documented as part of the study, and the findings are presented without individual attribution [133]. The repository makes the analysed framework and legal-source inventory inspectable. Workshop artifacts included completed canvases, notes, and post-its. Their interpretation in this chapter is limited to patterns documented at scenario and cross-scenario level in the project record.

The analysis does not treat agreement by participants as proof of validity. A participant can find a question understandable without the question being complete. A group can identify a risk without implementing an effective mitigation. Saturation in a formative interview process can indicate diminishing new suggestions for the artifact, but it does not establish that every public-administration domain has been covered. The empirical claims are accordingly bounded: the work shows how the Canvas was developed and what it enabled participants in this setting to articulate.

4.4 Canvas Structure

The canonical Canvas is a five-part visual document accompanied by instructions. Before entering those parts, the team describes the use case and asks whether AI is necessary for the objective. The thematic questions then identify which professional perspective should lead or contribute. A negative or unknown answer is not automatically a failure. It may require a design change, an explanation of non-applicability, collection of evidence, allocation of responsibility, specialist advice, or an explicit decision not to proceed. The figures in this section reproduce the seven pages of the English-language Canvas available in the project repository [33].

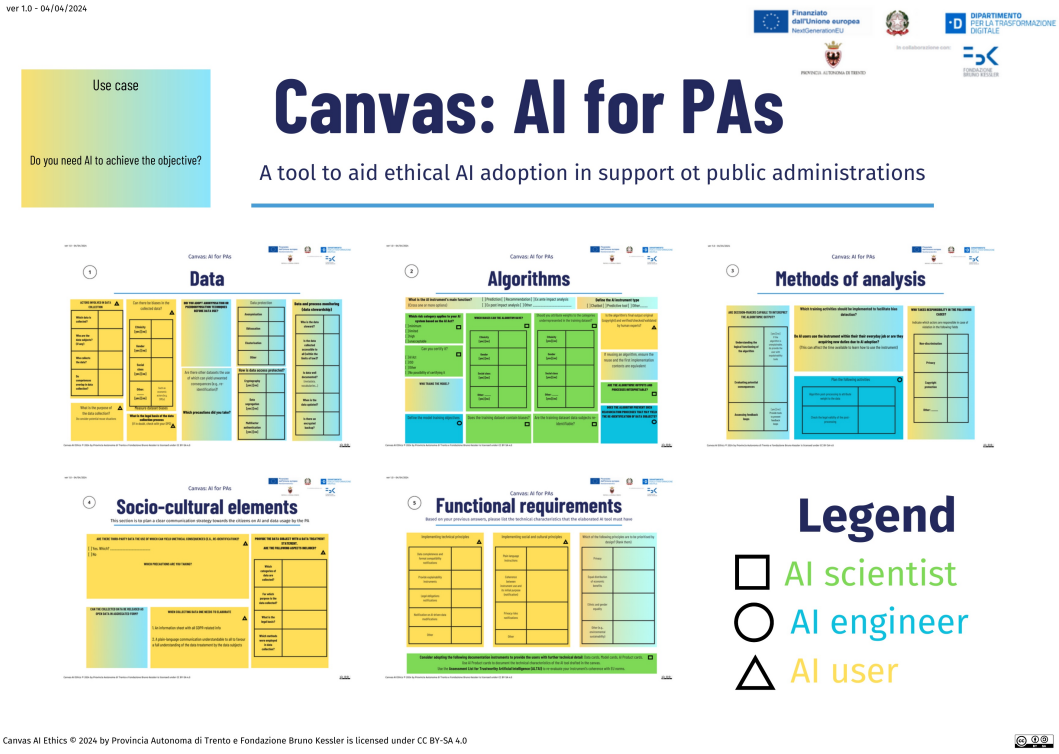


Figure 4.1: AI Ethics Canvas: use-case overview and role legend.

The Canvas is most useful when answers are expressed as propositions rather than ticks. Instead of recording "privacy: yes", a group should identify which data are involved, why they are required, what transformation is proposed, who determined the measure, where the supporting analysis can be found, and when the decision must be reviewed. This richer practice is not fully imposed by the graphic artifact itself; it is a facilitation and documentation discipline developed further in the thesis extension below.

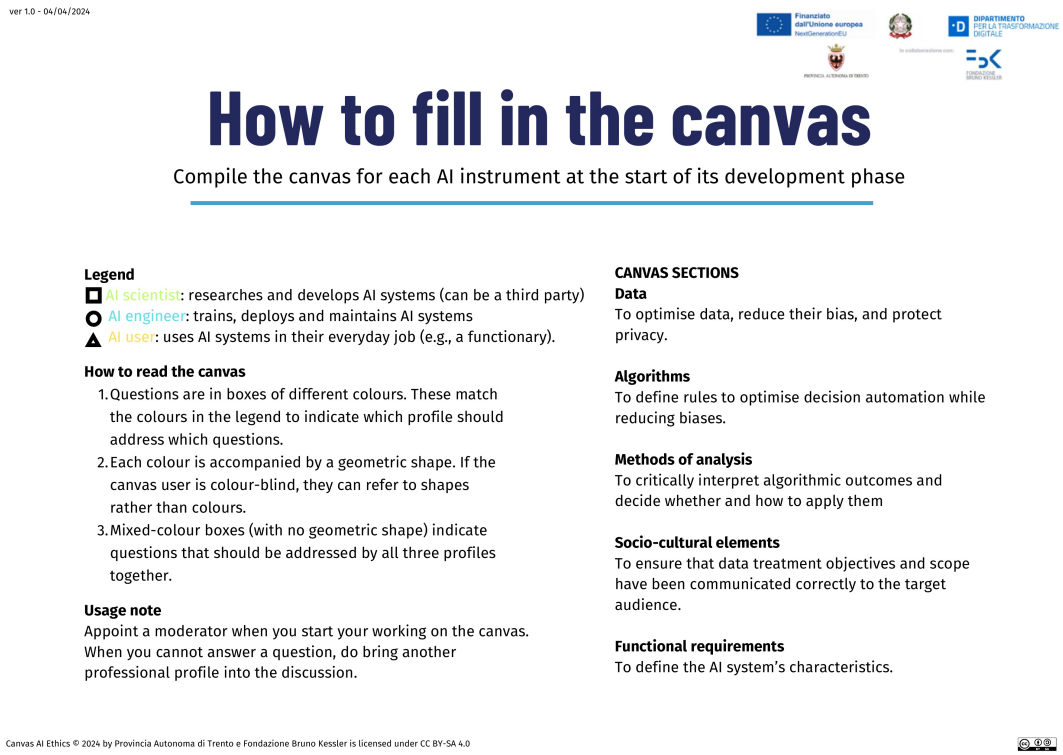


Figure 4.2: AI Ethics Canvas: use-case overview and role legend.

4.4.1 Three complementary roles

The *AI scientist* is the macro-role associated with creating or substantially developing algorithms and models. It includes researchers and software specialists who experiment with model architectures, training approaches, evaluation designs, or new technical methods. This role can explain what is being optimised, which assumptions are built into a modelling choice, what uncertainty is observed, and which tests have been performed. The title does not grant authority to determine the administrative purpose or acceptable impact.

The *AI engineer* is the macro-role associated with making the system operational. It includes data engineers, machine-learning engineers, and personnel responsible for fine-tuning, integration, deployment, maintenance, monitoring, and technical controls. This role connects an experimental model with data pipelines, interfaces, security, versioning, and an operating environment. The engineer is often best placed to identify what can be logged, which dependency may change, and how a requirement could be implemented. Again, technical control does not confer the institutional competence to decide whether the resulting use is lawful or desirable.

The *AI user* is the macro-role associated with the public-administration use of the system. It includes officials, managers, policymakers, and domain specialists who understand the administrative task, interpret outputs, or make decisions supported by the system. This role specifies what information is useful, which errors matter in practice, when an output must be challenged, and how a system affects a procedure or service. An AI user is not a passive endpoint. Domain expertise is necessary to judge whether a model's categories and recommendations correspond to institutional reality.

The roles are analytical categories, not a universal staffing chart. One person may combine more than one role; several organisations may share one role; and an external supplier may perform important scientific and engineering functions. Nor are the categories a substitute for legal or governance roles such as controller, provider, deployer, procurement authority, Data Protection Officer, or final decision-maker. A Canvas session should identify the actual people and organisations behind the macro-roles and record where decision authority lies. This is especially important when outsourcing creates a gap between the actor who can explain a technical choice and the authority answerable for administrative use.

4.4.2 Data

The data area begins with the resources on which the system depends. Participants identify training, validation, testing, operational-input, and reference data rather than treating "the dataset" as one undifferentiated object. They consider provenance, collection method, authority to obtain and use the data, update frequency, stewardship, quality, representativeness, documentation, and retention. Where personal data are present, the discussion must trigger a specialised assessment of purpose, necessity, legal basis, data minimisation, trans-

parency, security, and the possibility of anonymisation or pseudonymisation. The Canvas raises these matters but does not perform a data-protection impact assessment.

Questions about linkage are particularly important in a data-space setting. A dataset that appears non-identifying in isolation may create re-identification or inference risks when combined with other sources. Access through a governed infrastructure does not establish authority for every later operation. The group should therefore identify anticipated joins, derived variables, transformations, and secondary uses. Unknown future combinations cannot be resolved by a one-off session, but the project can record constraints and define a trigger for reassessment.

Stewardship connects these questions over time. Someone must know which version was used, whether an update altered the population, who can correct an error, and which documentation accompanies the data. Dataset documentation and provenance records can support this work [72, 79, 84]. They remain claims requiring evidence and maintenance; a complete card cannot guarantee that the underlying data are accurate or lawfully used.

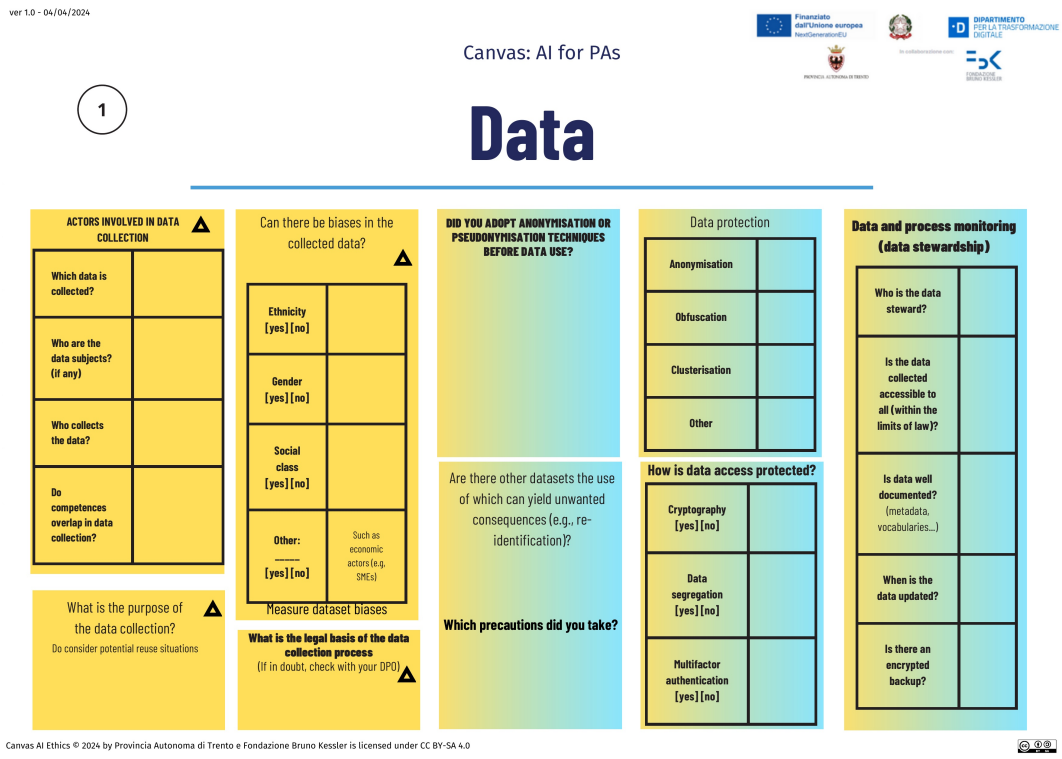


Figure 4.3: AI Ethics Canvas: use-case overview and role legend.

4.4.3 Algorithms

The algorithms area addresses the technical mechanism and the purpose for which it was selected. Participants specify whether the contemplated system performs classification, prediction, ranking, recommendation, generation, computer vision, optimisation, or another function. The distinction matters because a general label such as "chatbot" can obscure the model, retrieval sources, prompts, output controls, and downstream action. The 2024–2025 interviews revealed that non-technical participants often associated AI principally with conversational tools, prompting clearer differentiation in the artifact.

The group examines how objectives, labels, thresholds, weights, and optimisation choices relate to the administrative task. A technically reasonable objective may be normatively incomplete: optimising speed can distribute errors unfairly; predicting a historical decision can reproduce an earlier policy; and maximising aggregate accuracy can conceal poor performance for a small but materially affected group. The Canvas therefore connects algorithm design with the public purpose and asks who has authority to approve the trade-off.

This area also prompts identification of the AI Act classification and relevant obligations. It does not ask project participants to guess a risk label by analogy. The output should be a request for a documented, source-linked legal assessment based on the current system, purpose, role, and deployment context [2]. If the model or intended purpose changes, the classification and associated analysis may need to be reopened.

4.4.4 Methods of analysis

The methods-of-analysis area concerns the interpretation and use of outputs. A model may produce a score, label, generated text, alert, map, or recommendation; the administration must decide how that output enters a procedure. Participants identify who sees it, what explanation and contextual information accompany it, whether it can be contested, how uncertainty is communicated, and who takes the consequential decision. This moves the discussion from model performance to the human–AI system.

Human oversight is not satisfied by placing an official after the model. Empirical research documents automation bias, selective reliance, and the possibility that explanations increase acceptance without improving error detection [15, 140, 144]. Effective review requires competence, sufficient time, access to underlying evidence, independence, and authority to reject or escalate an output. The Canvas can ask who reviews and what happens when they disagree; the organisation must create the conditions that make the answer credible.

Responsibility questions belong here because the significance of an output depends on the action taken with it. The group should distinguish development, operation, use, approval, and appeal responsibilities. It should define an escalation route for an unexpected result, a suspected discriminatory pattern, data drift, a security incident, or a complaint. Reviewability then requires a record of who acted, what was considered, and why [67, 135].

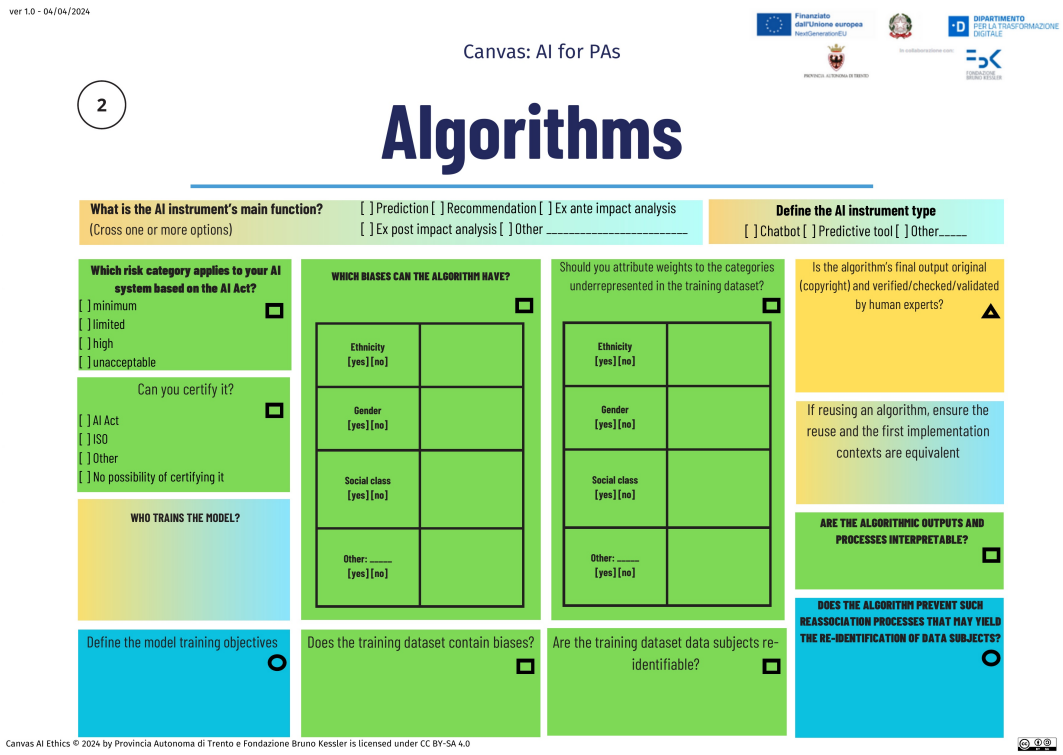


Figure 4.4: AI Ethics Canvas: use-case overview and role legend.

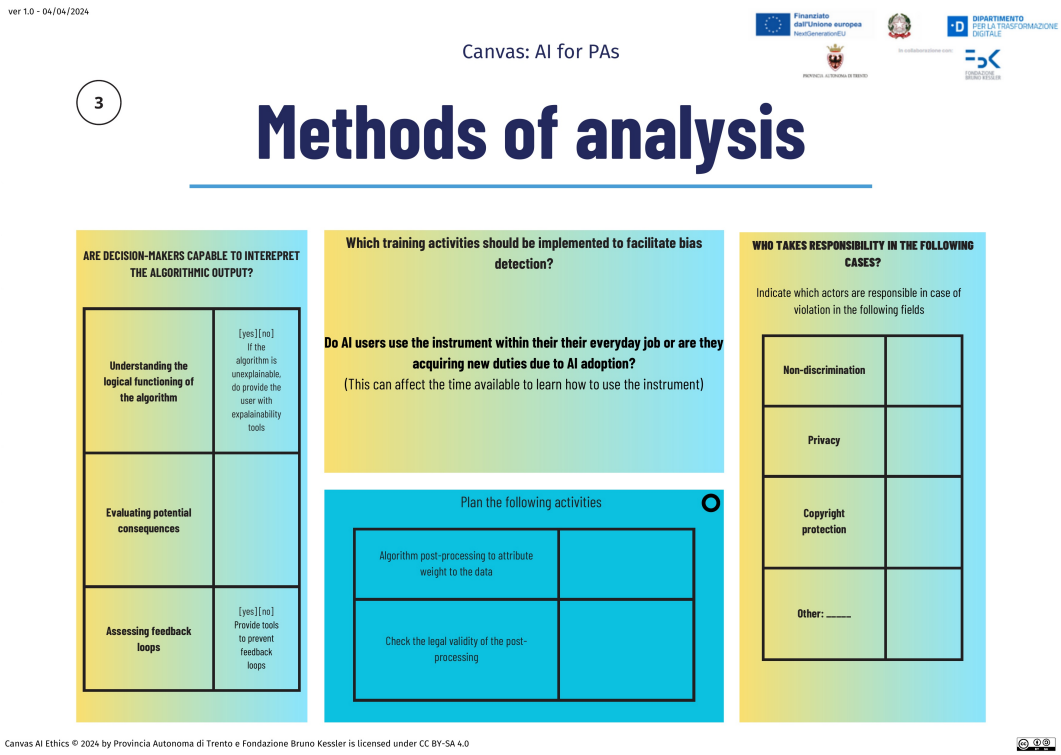


Figure 4.5: AI Ethics Canvas: use-case overview and role legend.

4.4.5 Socio-cultural elements

The socio-cultural area places the system within the community and institutional environment it affects. Participants identify groups who may benefit, bear risk, be excluded, or experience the service differently. They examine whether categories and language reproduce stereotypes, whether access depends on digital resources, whether disability or language needs are recognised, and whether an apparent efficiency gain shifts burdens to citizens or front-line staff. Fairness is consequently not reduced to parity between preselected groups. It includes the question of how the groups and outcome were defined.

Public communication is another component. People should be able to understand when an AI system materially contributes to a service or decision, which institution remains responsible, what the system cannot do, and how a concern can be raised. Transparency should provide usable information rather than promotional reassurance. The Canvas can identify audiences and information needs, after which legal notices, service descriptions, user interfaces, and communication plans must be designed for the context.

This area also considers whether collected or derived data can be made available for reuse, including in aggregated or open form. Openness is not presumed to be ethically preferable in every case. A publication decision must consider privacy, re-identification, group harms, intellectual-property interests, security, quality, and the capacity of future users to interpret the data. In a Common European Data Space, the same discipline applies to controlled sharing: discoverability and access conditions do not remove the need to consider downstream purpose and context.

4.4.6 Functional requirements

The functional-requirements area converts earlier discussion into features, constraints, organisational measures, and documentation tasks. It asks what the system and surrounding process must do to address the identified concern. Examples include role-based access, data-version recording, a human confirmation step, display of uncertainty, logging of a recommendation and override, monitoring by subgroup, an appeal channel, a defined retention period, supplier documentation, or a scheduled review. A broad value becomes actionable only after such decisions are made.

Not every response is a software feature. A requirement may concern staff training, procurement, a governance committee, domain validation, communication, a contractual right to audit, or a decision not to automate. Distinguishing technical from organisational requirements prevents a project from assigning every ethical problem to developers. It also helps procurement personnel state what evidence a supplier must provide and what remains the administration's responsibility.

Documentation instruments become relevant at this stage. Datasheets, model cards, data cards, system cards, and AI Product Cards can record structured information at different

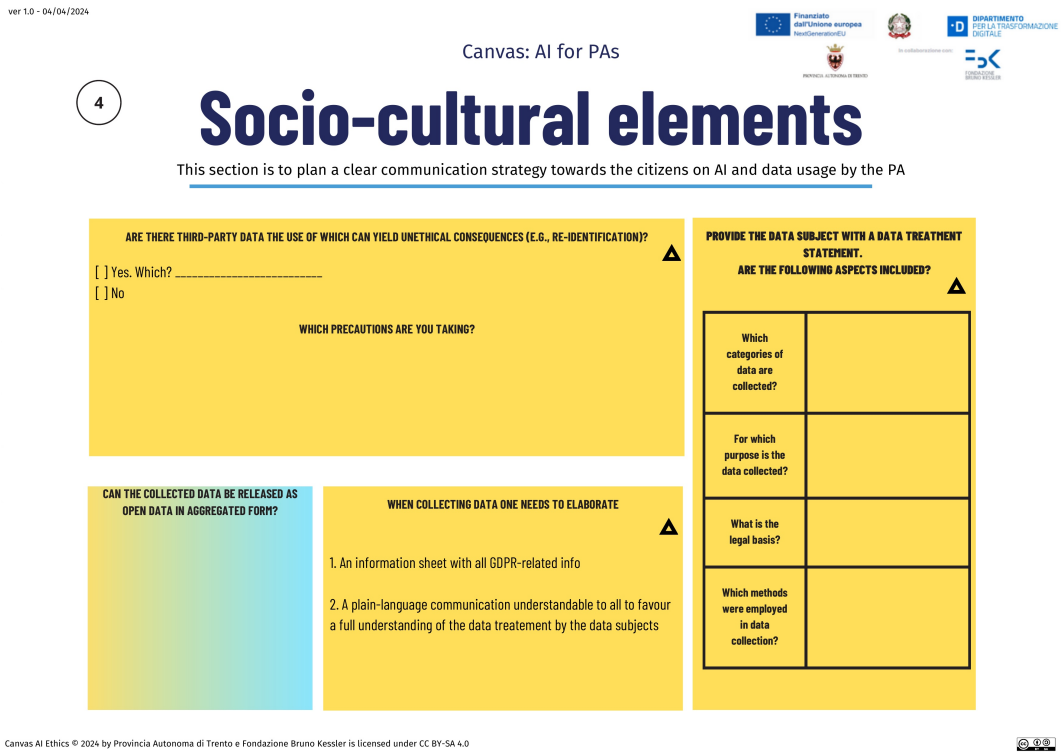


Figure 4.6: AI Ethics Canvas: use-case overview and role legend.

levels [79, 80, 82]. The AI Product Card is particularly relevant to public administrations because it connects descriptive and formal information about an AI product with implementation artifacts. The Canvas precedes and complements such documents: it helps a group decide what must be documented and why, while the cards provide more stable containers for the resulting information.

Requirements must retain links to their reasons. If a system logs overrides, the project record should identify which risk or governance need motivated the log, who may access it, and how long it is retained. Otherwise a technically implemented feature can outlive its purpose or create a new risk.

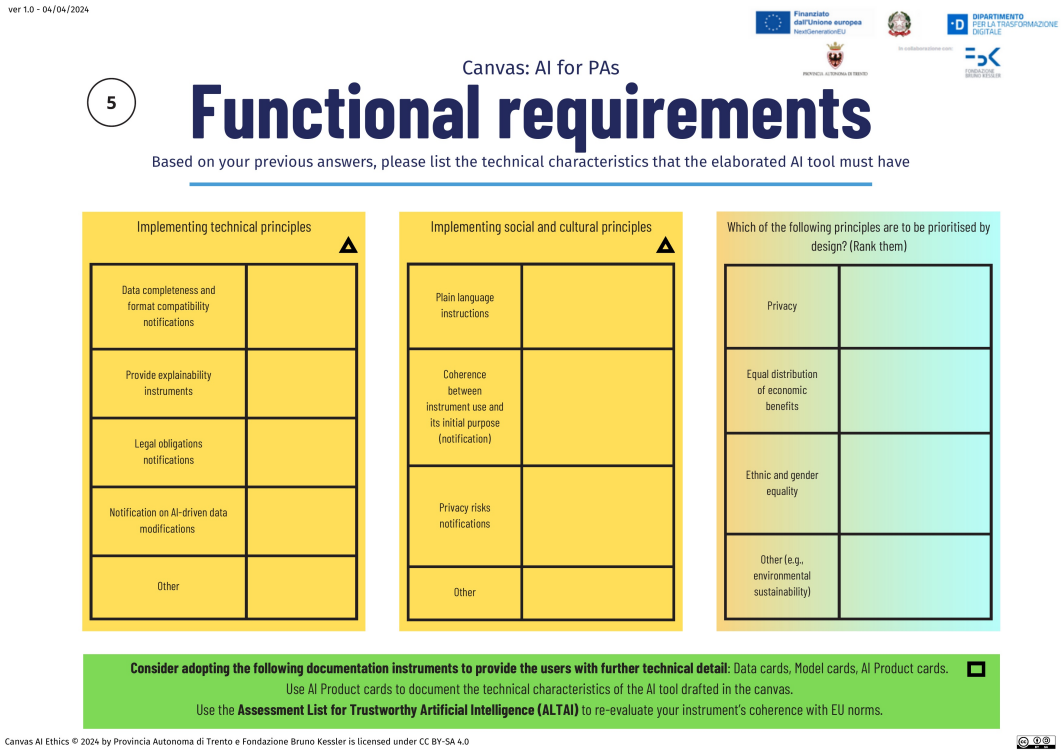


Figure 4.7: AI Ethics Canvas: use-case overview and role legend.

4.5 Formative Evidence

4.5.1 Contextual needs evidence: survey and focus-group activities

The wider AIxPA project included a consultation survey and focus-group activities about AI in the provincial public administration. These activities explored perceptions of benefits and risks, organisational readiness, public values, possible use domains, training needs, human responsibility, regulatory complexity, and concerns about transparency, privacy, quality, and

loss of control. They helped establish the environment in which a practical ethics instrument would be useful.

This chapter treats that material as *contextual needs evidence*, not as direct evaluation of the Canvas. The consultation elicited attitudes and priorities concerning AI and project guidance. The earlier focus-group activities supported use-case and organisational problem identification. Participants were not consistently evaluating the canonical Canvas, and the activities had different questions, units of analysis, timing, and completion rules. Accordingly, no survey response counts, denominators, or quantitative distributions are used here, because those values were not reconciled for this chapter and would create false precision.

The qualitative signal is nevertheless relevant. Officials and project actors repeatedly identified the need for understandable language, AI and data literacy, continued human responsibility, transparent communication, data quality, and guidance through a complex legal setting. These needs informed the action-research context and are compatible with the Canvas design. They do not prove that the Canvas met them. That narrower question was addressed through the interviews and later workshop.

4.5.2 Direct formative evaluation: twelve interviews

Between 3 September 2024 and 28 January 2025, nine semi-structured interview sessions involved twelve participants: six sessions were individual and three were paired. Two participants were categorised as AI scientists, two as AI engineers, and eight as AI users. Their professional backgrounds included software development and research, technical innovation, and provincial administrative work in territorial planning, cartography, cadastre and land registration, research funding, ICT, and digitalisation. These categories retain the study's macro-role classification rather than representing independently assessed technical competence. Interviews lasted approximately thirty to seventy-five minutes. Recruitment used project access and snowballing. The researchers report that recruitment ended when later interviews no longer added significantly new design suggestions; this was a formative stopping rationale rather than a documented formal saturation analysis [133].

The interview guide addressed current AI use, anticipated future use in light of PAT projects, unclear parts of the Canvas and proposed modifications, aspects considered inapplicable to everyday work and proposed modifications, and issues insufficiently covered. These questions reproduce the scope of the guide reported in the associated article; they were asked in Italian, with follow-up questions permitted by the semi-structured format. Feedback concerned both content and presentation. The documented analytical procedure consisted of reviewing suggestions, revising the artifact after interviews, eliciting subsequent participants' responses to earlier changes, and retaining changes judged necessary across the aggregate feedback. This procedure supported formative refinement rather than an independent or causal evaluation of project outcomes [133].

One material finding was that many AI users lacked specialised AI-development knowl-

edge because development and deployment capacity was often supplied outside the administration. Participants could understand the administrative problem while struggling with technical terms. The Italian version therefore gained explanatory material, including a glossary, and the design retained visual cues intended to make role assignment and navigation accessible. The finding also strengthened questions about training and interpretation in the methods-of-analysis area.

A related framing problem emerged because some participants associated AI chiefly with chatbots or generative interfaces. The artifact was clarified to encourage explicit identification of the system type and purpose. This change may appear terminological, but it prevents an important framing error: a user-facing conversational interface can involve retrieval, generation, classification, or access to administrative sources, each with different evidence and control needs.

Technical participants, in turn, pressed for a clearer passage from ethical prompts to implementation and documentation. Feedback contributed to the functional-requirements area and to the relationship with data, model, and product cards. The point was not to expand the Canvas into a complete software specification. It was to ensure that a deliberative conclusion could generate a task or artifact elsewhere in the project rather than remaining on a post-it.

Participants generally found the visual format helpful for entering a complex discussion, while also finding portions technically demanding. That combination is analytically useful. It suggests that accessibility is not achieved merely by simplifying vocabulary; a public-administration tool must preserve difficult questions while providing facilitation and routes to expertise. The evidence supports the claim that the Canvas can structure cross-role reflection and can be improved through participant feedback. It does not show that a team can complete it unaided or that use changes deployed-system performance.

4.5.3 Direct formative application: the 2025 workshop

A two-hour workshop on 14 October 2025 provided a later episode of direct formative application. The final AIxPA project record lists 18 participants in total, encompassing active group members, cross-cutting contributors, auditors or observers, facilitators, and organisation support. The count should not be read as 18 independent evaluators performing the same task. Participants occupied different roles in conducting, observing, and documenting the activity.

The workshop concerned *two broader project cases*, not three unrelated cases. The first addressed the definition and evaluation of policies supporting reconciliation between productive work and family life. It was discussed through the *Family* scenario. The second addressed identification and monitoring of territorial fragilities. It was divided into two application scenarios: *Extreme events*, concerning GeoAI-supported monitoring of territorial effects associated with events such as landslides, floods, windthrow, or fires; and *Over-*

tourism, concerning evidence and scenario analysis for sustainable tourism management. Participants therefore worked at three tables and produced three scenario discussions nested within two project cases.[1]

Each table received an A3 Canvas, individual section sheets, explanatory guidance, a glossary, post-its, and writing materials. Participants moved through all five thematic areas. More time was allocated to data, algorithms, and methods of analysis, while socio-cultural elements and functional requirements received shorter dedicated intervals. Facilitators coordinated timing and capture. The groups then reported to a plenary discussion. Completed canvases, individual notes, post-its, photographs, and discussion records formed the qualitative material for the project report.

The Family table raised questions about representativeness and missing variables in administrative data, including the possibility that recommendations would remain partial where relevant social circumstances were not recorded. Participants also considered turnover among local administrators and the resulting need for recurring training. Proposed responses included exploring how missing information could be handled, ensuring human review, and allowing priorities and system guidance to be revisited as policy objectives changed. The discussion shows the interaction of data, institutional continuity, and functional requirements; it does not establish that any proposed mitigation was effective.

The Extreme-events table worked with territorial rather than person-level data and did not foreground personal-data risk in the same way. Its central concern was the danger of treating a predictive output as definitive. Participants emphasised domain-expert review, continued assessment of output reliability and usefulness, and operational notification mechanisms that fit emergency and territorial decision processes. The scenario demonstrates why the absence of personal data does not end ethical inquiry. Reliability, uncertainty, institutional responsibility, and consequences of action remain material.

The Overtourism table similarly concentrated less on individual identification and more on interpretability for public decision-makers. Participants identified a possible grey area between a technical recommendation and a political or administrative choice. They called for clearer explanations of the criteria and possible consequences of suggested actions, as well as periodic follow-up. The discussion made visible the risk of responsibility being displaced onto an apparently objective system when decision-makers cannot reconstruct how a recommendation was produced.

Across the three scenarios, the project report identified recurring concerns about data quality and representativeness, transparency and explainability, training and interpretive capacity, human supervision, and continued review. The workshop also showed that the same printed questions produce different discussions when the governed object changes. Privacy was central in the Family scenario but less prominent for aggregate territorial and tourism information; interpretability and accountability remained central across them. This

[1] The reconstruction follows the final project deliverable: AIxPA, "Report su Assessment etico relativo ai casi d'uso", AI-M12—Assessment etico, version 1.0 (2025).

contextual variation is one reason the Canvas should not calculate one ethical score.

The workshop supplies evidence that a facilitated group can use the artifact to identify risks, disagreements, evidence needs, and possible mitigations in concrete project discussions. It also generated observations about wording and future use. It does not establish longitudinal effect. The projects were not followed through a controlled comparison of design, procurement, deployment, and service outcomes, and the workshop did not technically test the underlying systems. The term *formative application* therefore remains more accurate than certification or effectiveness validation.

4.5.4 Findings and Evidentiary Boundaries

The documented development process provides an inspectable account of how the Canvas was derived from frameworks and EU legal sources, iteratively revised with twelve interview participants, and applied in the AIxPA setting. It provides the artifact, the role and thematic architecture, methodological limitations, and the interpretation that a visual tool can help public-administration actors organise ethical reflection [133]. The empirical contribution concerns design and formative usability in one institutional setting.

The available evidence does not show that Canvas use causes more lawful systems, fairer outcomes, better accuracy, more effective human oversight, or increased public trust. It does not support a comparison between the Canvas and alternative methods, a measure of inter-group consistency, or an evaluation of adoption across several Member States. This chapter preserves those limitations and adds a more explicit separation of evidence strata. It also narrows claims that project documentation occasionally described as "validation": the activities validate neither the lawfulness of a use case nor the truth of an answer written on the artifact.

4.6 Iterative Project Review

The Canvas can be used at three kinds of checkpoint: ex ante, during development and operation, and at final or periodic review. The same question architecture allows assumptions, requirements, evidence, and responsibilities to be reconsidered as the project changes. These checkpoints support review; they do not constitute certificates.

In AI systems, the case for data stewardship arises from the need to govern data throughout the system lifecycle rather than merely to document it at the point of deployment. This includes allocating and monitoring responsibilities for data collection, quality, access, transformation, use, and review across organisational boundaries [183]. Evidence from high-stakes AI projects shows that weak ownership and coordination of data work can generate cascading problems that emerge only at later stages of system development [184]. Data stewardship can therefore preserve contextual knowledge, coordinate maintenance, and trigger renewed

assessment when data or uses change; it does not, by itself, establish the quality, fairness, or legal compliance of an AI system.

4.6.1 Initial Framing

At the ex-ante checkpoint, the administration defines the problem before committing to a model or supplier. Participants describe the public purpose, legal mandate, intended beneficiaries, affected groups, contemplated decision, data needs, available alternatives, and reasons for using AI. They identify which roles and institutions must participate and which questions require specialist assessment. The output is a structured set of assumptions, unresolved issues, preliminary requirements, evidence requests, and stop conditions.

This checkpoint can alter procurement and project selection. If a required dataset lacks sufficient provenance, if an official cannot lawfully obtain it, if the proposed output cannot be integrated into a reviewable procedure, or if a non-AI alternative better serves the objective, the project can be redesigned before sunk costs narrow the choice. A completed Canvas is not required for permission to proceed; rather, the administration must decide how material uncertainties will be resolved and who owns each action.

For a project using data discovered in a Common European Data Space, the ex-ante checkpoint should distinguish discovery and access from downstream authority. Participants record the source, proposed operation, purpose, access route, transformations, and intended system. They do not infer from catalogue availability that training or deployment is authorised. The information recorded at this stage can support subsequent compliance and assessment activities.

4.6.2 Iterative review during design and operation

During design, answers should be reopened when the facts change. A new dataset, model, supplier, affected group, purpose, interface, threshold, or operational environment can invalidate an earlier conclusion. Iteration is not the repeated completion of an unchanged form. It is comparison between versions: what changed, which earlier assumptions are affected, what new evidence exists, who reviewed it, and whether requirements or responsibilities must change.

Project milestones offer practical triggers. A team might revisit the data and algorithm areas when a training corpus is selected; methods of analysis when an interface and human-review workflow are prototyped; socio-cultural elements when user research or consultation produces new evidence; and functional requirements before procurement acceptance or deployment. Incidents, complaints, drift, and legal or guidance changes create additional triggers. The same process can be used after deployment as part of periodic monitoring, provided that the administration has access to operational evidence.

The iterative proposal responds to a common weakness of ethics tools: a workshop can

produce a thoughtful record that becomes detached from engineering and management. Versioned checkpoints make the record a project dependency. A requirement cannot be marked resolved merely because a feature exists; the group must consider whether the implementation and evidence address the concern and whether any residual risk remains. This is consistent with research that treats reviewability and human–AI interaction as lifecycle properties [67, 134].

4.6.3 Final Review

The final-review checkpoint compares the implemented or operated system with the purpose, assumptions, and requirements recorded earlier. "Final" describes a project gate, not the end of responsibility. Participants ask whether the data, model, user interface, institutional workflow, and assigned roles correspond to what was reviewed; whether promised evidence exists; whether limitations are communicated; whether staff are trained; and whether monitoring and escalation are operational. Unresolved matters can block acceptance, require remediation, or be escalated to the competent authority.

The Canvas cannot supply the technical tests or legal opinions required at this stage. It can help make their absence visible. A data-quality question should link to a quality report; an oversight question to a procedure, training record, and test of reviewer conditions; a classification question to a current legal assessment; and a communication question to the deployed notice or interface. The ex-post session therefore verifies the project record and reopens reasoning. It does not certify legal compliance or technical quality.

Because the projects were not followed through these three checkpoints, the proposed lifecycle use must be tested in future work. A longitudinal evaluation should examine whether issues identified ex ante influence requirements, whether later changes cause reopening, whether evidence remains available, how disagreements are resolved, and what burden the process imposes. It should compare the record with actual project decisions and outcomes without using "number of completed fields" as the principal success measure.

4.7 Information and Evidence Touchpoints

The Canvas creates touchpoints at which its questions and the involvement of different parties allow information to be collected. When recorded, this information may enrich project documentation and metadata and support subsequent compliance and assessment activities. These touchpoints do not prescribe a metadata schema, determine the applicable legal requirements, or establish compliance.

4.8 Relation to Common European Data Spaces

Common European Data Spaces heighten the organisational importance of the Canvas. A federated environment can improve discovery, interoperability, access negotiation, and the exchange of governed data descriptions. A public administration may obtain data from such an environment for analysis or training. It may also contribute administrative, environmental, mobility, tourism, or service data and derived resources. These two directions create a cycle of use and contribution, but the thesis does not claim that the full cycle was empirically observed in AIxPA.

When a public administration consumes data, the Canvas can prompt questions that catalogue discovery alone cannot answer: whether the dataset is fit for the specific public task; whether the proposed operation is within the access conditions; what combinations or inferences are planned; which groups are represented; and what limitations must reach the system user. The resulting evidence needs can inform an access request, a project decision, or a decision not to use the resource.

When an administration contributes data or an AI-related artifact, the Canvas can prompt it to consider provenance, documentation, quality, permitted reuse, affected communities, update responsibility, and the consequences of making the resource discoverable. The functional-requirements area can generate requirements for catalogue descriptions, versioning, access policies, product documentation, and review.

The Canvas is not a data-space connector or usage-control mechanism. It does not negotiate a contract, authenticate a participant, enforce a policy, or prevent copying after data leave a controlled environment. Nor does data-space infrastructure guarantee that a participant has deliberated adequately. The relationship is instead organisational: structured deliberation can help identify information, evidence, and decisions that may need to be maintained or communicated. Technical mechanisms can implement bounded controls where appropriate, while human and institutional actors retain authority over interpretation and exceptions.

This connection also places the Canvas within the thesis's account of human-centred compliance. Human-centred does not mean that a person merely signs off an automated result. It means that relevant professional and affected perspectives can influence design, evidence is available to those who review it, responsibilities are explicit, and a decision can be challenged and reopened. The Canvas supports the conditions for such work; it does not establish that they are present in every session.

4.9 Limits, risks, and conditions for responsible use

Its relevance is necessarily contextual. The artifact was developed with one regional public-administration system and partners in Italy. EU-level legal references support broader relevance, but organisational structures, languages, administrative law, procurement practice,

and available expertise differ across Member States and sectors. Adaptation is therefore expected. The CC BY-SA 4.0 licence permits it, while share-alike terms support continued accessibility of adaptations. A local extension should identify its source version and preserve the distinction between core prompts and jurisdiction-specific additions.

Twelve interviews and one later workshop provide rich formative material on the artifact's development and use, but do not support statistical generalisation. The interview distribution deliberately favoured AI users because they constituted the principal administrative audience. The workshop included several participation functions and three scenario tables rather than eighteen equivalent evaluators. No control group or comparative instrument was used. A later claim that the Canvas improves outcomes would require a research design that observes decisions and systems over time.

The relationship between Canvas use and CEDS described in the previous section is a research proposition requiring empirical evaluation; it was not investigated in AIxPA. Future inquiry should ask whether, compared with existing administrative practices, Canvas use in projects consuming or contributing CEDS data identifies otherwise overlooked evidence needs, responsibility gaps, or affected perspectives, and whether these findings lead to further assessment, revised requirements, or reconsideration of a project decision.

Facilitation and power create a further condition for responsible use. A common visual surface does not remove hierarchy. A supplier may withhold technical detail; a junior official may not challenge a director; an absent community cannot contest how its interests are represented; and a facilitator can compress disagreement into a consensus statement. Responsible use should record dissent, unknowns, and absent perspectives. It should also identify who has authority to decide and who can halt or escalate the project. Completion without these conditions risks becoming checkbox ethics [132].

Its legal boundary is equally clear. The Canvas is not a source of law, legal opinion, conformity assessment, or safe harbour. Its questions may identify the need to assess the GDPR, AI Act, procurement law, intellectual property, equality law, sectoral rules, or other sources. That assessment must use the current authoritative text and facts. The relevant legal roles and competent institutions retain their powers. A positive answer written by a project team neither displaces an obligation nor proves it has been fulfilled.

Technical evaluation remains outside the Canvas's function. The Canvas does not evaluate code, cybersecurity, robustness, accuracy, fairness, privacy protection, or environmental performance. It can request the relevant test and help participants interpret why it matters. Technical results are bounded by methods, data, conditions, and thresholds. A metric must not be turned into a general statement of ethical quality, and a human reviewer must not be presumed effective simply because one is named.

Its usefulness is also time-sensitive. Data drift, model updates, organisational change, legal developments, and new uses can invalidate earlier answers. Release 1.0.0 itself will require maintenance as the legal and technical environment evolves. A project must record

the Canvas version used and define review triggers. Without versioning and reassessment, the artifact may provide a misleading appearance of continuing assurance.

Generality is deliberate, and domain expertise remains indispensable. A social-services project, an environmental prediction system, and a tourism-planning dashboard expose different values, evidence, and consequences. The printed prompts are a floor for discussion, not a ceiling. Participants should add questions when the domain reveals a concern that the artifact does not contain.

4.10 Conclusion

The AI Ethics Canvas addresses the organisational level at which European AI and data regulation must become operational. Its contribution is not a new statement of ethical principles but a structured setting in which AI scientists, AI engineers, and AI users can examine one public-administration use case across data, algorithms, methods of analysis, socio-cultural elements, and functional requirements. The action-research and value-sensitive-design process shows how an artifact can be revised through contact with the organisation expected to use it.

The evidence is deliberately bounded. Contextual survey and focus-group activities identified organisational needs but did not directly evaluate the Canvas. Twelve interviews supplied direct formative feedback on clarity, relevance, roles, terminology, and implementation links. The workshop applied the tool with 18 participants across two broader project cases and three scenarios, revealing both recurrent concerns and domain differences. These activities support the Canvas as a deliberative and requirements-framing instrument. They do not demonstrate longitudinal effectiveness or certify compliance and technical quality.

Used at initial, iterative, and final-review checkpoints, the Canvas allows assumptions, requirements, evidence, and responsibilities to be revisited. The effects of this use remain a question for future longitudinal study.

Chapter 5

From EU Legal Concepts to Metadata Standards

5.1 The Informational Problem

Two developments motivate the informational problem addressed here, yet they remain insufficiently connected. EU law governs the sharing, transfer, access, and reuse of data through a complex body of horizontal and sectoral instruments. These instruments allocate different roles, powers, obligations, conditions, exceptions, and evidentiary requirements according to the actor, governed object, operation, purpose, access route, territory, and time. The legal framework cannot therefore be reduced to a single rule governing whether a dataset may be used.

In parallel, technical communities and European institutions continue to develop vocabularies, application profiles, provenance models, policy languages, quality models, and validation mechanisms for describing datasets and the conditions under which they are exchanged and reused. These specifications can improve discovery, interoperability, traceability, and machine-readable documentation. They do not, by themselves, determine which legal concepts should be represented, how source-specific qualifications should be preserved, or where technical representation must stop [4, 5, 116].

What is missing between these domains is a structured method for moving from an exact legal source to an information need and from that need to the technical elements that may represent it. The contribution of this chapter is a conceptual, source-linked framework for making that transition explicit. The framework is designed to identify possible support, partial representation, representational gaps, and questions that must remain with competent legal, technical, professional, or institutional actors. It does not claim to establish the completeness of the legal corpus, the coverage of existing standards, or the actual distribution of gaps across the European data-space ecosystem.

The running dataset journey makes the distinction concrete. A regional health authority and the relevant data-space actors first make a health dataset discoverable by preparing and maintaining its description, provenance, quality information, and conditions of availability. A public research organisation may subsequently discover the dataset, request access, and propose a specific downstream use, including the development or validation of an AI system. These are related but distinct events. A catalogue entry identifies and describes an offering; it does not make the underlying data unrestricted. An access decision authorises only the operation, actor, purpose, period, and conditions within its scope. A later transformation, purpose, recipient, or system classification may activate other rules. Copyright, data protection, sectoral legislation, contractual terms, and the AI Act can remain relevant even after an upstream access route has been established [1, 2, 32, 56, 57].

The response developed here deliberately confines itself to one analytical chain; comprehensive compliance architectures and all-encompassing ontologies remain outside its scope:

> **Legal source → source-linked legal concept or requirement → metadata need → exact standard element → representational outcome.**

Each arrow is an analytical transformation that must be explained and capable of being reversed. A reader should be able to move from a reported mapping back to the exact element definition, the need for which it was selected, the reviewed legal proposition from which that need arose, and the authoritative source on which the proposition depends. The reverse path should show what technical support was found for a selected source component and where the path ends in a partial representation, an explicit gap, or a matter reserved for human judgment. The chain does not infer one stage automatically from the previous one.

Four principles operate across the whole chain: responsibility, evidence, versioning, and reassessment. They are transversal safeguards, not additional analytical stages. Responsibility identifies who made, reviewed, or maintains an assertion and who has authority to decide the remaining question. Evidence connects an assertion or decision to the material on which it depends. Versioning fixes the legal source, dataset, statement, technical specification, and mapping used at a particular time. Reassessment ensures that a material change can reopen the affected part of the analysis. Treating these principles transversally avoids expanding the conceptual framework while still preventing a mapping from becoming detached from the people, materials, and temporal conditions that give it meaning.

The five stages and their transversal safeguards establish the relations that the method seeks to preserve in operational use. On that basis, data stewardship, a governed corpus, an assessment-supporting knowledge graph, and modular standards can be considered as possible means of preserving those relations in practice, while two bounded micro-applications drawn from Data Act Article 33(1)(a) illustrate how the chain reaches both an exact technical representation and a question reserved for competent judgment; the limits exposed by those applications then frame the discussion of transfer and governance.

Appendix B identifies the legal instruments and versioned technical specifications that informed the framework and states their function in the analysis.

5.2 The Five-Stage Analytical Chain

The chain offers a method for controlled translation; treating law as though it could be compiled directly into metadata would obscure at least three kinds of choice. Legal analysis selects the relevant source and interprets its relation to facts and other rules. Information analysis decides which aspects of a proposition require a retrievable record, evidence, or review. Semantic analysis chooses a technical element whose published meaning is sufficiently close to the stated need. These choices are related, but each has its own evidence and limitations. Separating them prevents a convenient vocabulary term from being treated as the meaning of the legal source and prevents a legal phrase from being reported as though it already specified a technical field.

Figure 5.1 presents the complete conceptual model. The visual contains exactly the five analytical stages used throughout the chapter. The four safeguards shown below the chain apply to every transition; they do not enlarge it.

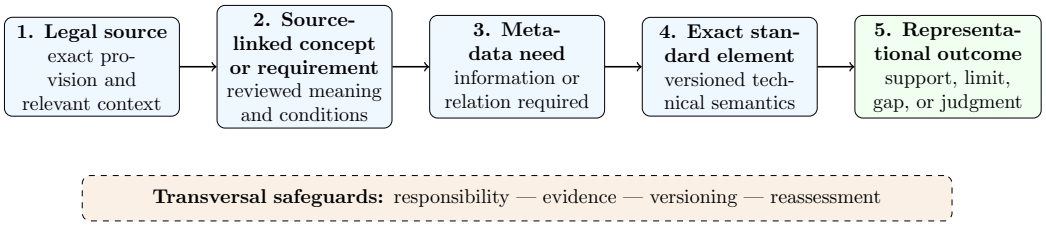


Figure 5.1: The single five-stage conceptual framework

5.2.1 Stage 1: legal source

Analysis begins by fixing the exact legal provision on which it is based. The record points to the official act and the relevant article, paragraph, or point, and notes the language, date, and version examined. It includes only the definitions, exceptions, and cross-references needed to understand that provision. Every later requirement can therefore be traced to the same legal source.

Judgments, official guidance, and scholarship may support the interpretation, but they remain separate from the legal text. ELI can be used to identify the act and its subdivisions in machine-readable form. The output of Stage 1 is a verified reference to the legal source, together with the minimum context needed for the next stage.

5.2.2 Stage 2: source-linked legal concept or requirement

With the source fixed, Stage 2 explains what the selected provision requires, permits, prohibits, or makes conditional. The resulting proposition identifies the relevant actor and legal role, action, governed object, purpose, conditions, exceptions, and timing.

Independently assessable duties are recorded separately, while qualifications that change the meaning of a duty remain attached to it. The proposition stays linked to the exact legal source and records any material uncertainty. It supports the later analysis but does not determine whether the provision applies, or whether it has been complied with, in a concrete case. Requirements-engineering research supports this form of explicit decomposition and traceability while also showing the difficulty of maintaining links across legal and technical work [74, 75, 76, 110, 113, 114].

In this thesis, "reviewed" denotes completion of the author's source-based analytical check. A proposition is treated as reviewed after its formulation has been checked against the exact provision, necessary definitions and cross-references, and relevant conditions and exceptions, with material uncertainty and alternative readings retained. The review concerns the fidelity and defensibility of the proposition for the bounded analytical purpose of the framework. Its authority is scholarly: it does not constitute independent legal validation, institutional approval, or a binding interpretation.

In an operational application, the organisation would need to assign this review to a person or function with relevant legal competence and an identified mandate for the assessment. The review record should identify the reviewer, scope of responsibility, sources and versions examined, factual assumptions, reasoning, date, unresolved issues, and outcome. Legal advice and analytical review remain distinct from a decision adopted by the institution authorised to act. A reviewed proposition may therefore remain contextual and contestable: alternative interpretations should be preserved where material, and changes in the legal sources or factual configuration should trigger reassessment.

Constructing source-linked propositions

Starting from that verified subdivision and the definitions, exceptions, and cross-references necessary to understand it, the operative text is divided into propositions that can be assessed independently, while conditions that alter the meaning of a duty remain attached. Each proposition retains its source pinpoint and identifies the relevant actor and legal role, modality, action, governed object, purpose, conditions, exceptions, timing, dependencies, and material uncertainty.

Atomicity is semantic rather than grammatical. A paragraph may require both publication of a catalogue description and inclusion of specified content. Evidence that an endpoint exists answers the first issue but does not show that the description contains each required category; a prepared description answers the second issue but does not show that it was published through the required channel. Separating those duties makes their information

needs and evidence distinguishable without detaching qualifications from the proposition they govern.

Legal roles and governed objects retain the granularity of their source. A public-sector body under the DGA, a participant offering data under the Data Act, a health data access body under the EHDS, an AI provider under the AI Act, and a research organisation under the DSM Directive are not instances of one interchangeable organisational role. A dataset is similarly distinct from its catalogue record, distribution, quality annotation, licence, policy, permit, evidence item, transformation activity, model, or secure processing environment.

Applicability is configuration-sensitive. A concrete assessment must identify the actor and role, governed object, operation, purpose, access or rights route, recipient, territory, relevant system classification, lifecycle stage, date, and material evidence. If those facts are absent or contested, the proposition records the missing evidence or disputed assumption for the relevant evidence provider and legal assessor, with any necessary determination reserved for the authorised decision-maker, instead of assigning a default legal state.

Some instruments may operate as contextual dependencies rather than as sources of a separate proposition. The GDPR illustrates this distinction: where personal data fall within its scope, it may qualify the interpretation or applicability of a requirement derived from another instrument without independently generating a metadata need or representational outcome in the bounded application below. A mapping derived from another source must not be read as supplying a legal basis or satisfying the GDPR [2, 3].

5.2.3 Stage 3: metadata need

Only then does Stage 3 ask what information must be available to support work on the proposition. A metadata need is expressed in solution-neutral terms before any vocabulary is chosen. It may require the ability to identify a dataset and distinguish its version; connect a catalogue record to the described resource; record an asserted role; describe an intended purpose or a use condition; trace collection and transformation; associate a quality assertion with a method and evidence; link a decision to its source; or make a change discoverable. The need specifies the subject, relation, granularity, and operational purpose of the information.

Not every legally material phrase becomes a metadata property. Four responses are possible before technical selection. A material aspect may require structured information that can be retrieved and exchanged. It may require an addressable evidence item whose content remains outside the graph. It may be suitable for a deterministic structural check after the relevant assumptions and values have been fixed. Or it may depend on legal, factual, professional, or institutional judgment and therefore require a review record rather than a machine-generated result. If an existing need already captures the same operational function, the analysis creates no duplicate field.

This intermediate step is essential because metadata schemas organise information for particular purposes and at particular levels of granularity. A schema may contain a property

called “rights”, “quality”, “purpose”, or “role”, but the label does not show that it represents the legal aspect under examination. Schema and ontology matching research similarly warns that correspondence weakens when labels mask differences in structures, domains, ranges, and conceptual granularity [91, 92, 93]. By fixing the need first, the method can later explain why a candidate element contributes, which part it contributes, and what is lost.

Deriving metadata needs

The move from a reviewed proposition to a solution-neutral metadata need is the conceptual centre of the framework. The analysis asks which information must remain available to identify the governed object, understand the relevant assertion, locate supporting material, perform a closed structural check, or route an open question to review. The resulting need describes an informational function; it neither repeats every word of the provision nor assumes that the function corresponds to one field in an existing vocabulary.

A requirement to describe a dataset’s collection methodology illustrates the distinction. The phrase does not imply one free-text property named “methodology”. An inspectable account may need to identify the collection activity, responsible agent, plan or protocol, inputs, dates, generated dataset version, transformations, and supporting document. Some relations can be represented directly as provenance, while the detailed method remains in a linked artifact.

Five questions guide the derivation:

1. Which object or event must be identified, and at what granularity and version?
2. Which actor, role, action, purpose, condition, or temporal relation must remain connected to it?
3. Which assertion must be directly retrievable, and which supporting material can remain in a linked artifact?
4. Which aspect can be tested through a fixed technical constraint after the relevant assumptions have been supplied?
5. Which aspect depends on authority, truth, sufficiency, proportionality, fitness, or another form of accountable judgment?

The framework groups common metadata needs into the categories shown in Table 5.1. The categories organise the analysis without adding stages to the chain or claiming to exhaust every legal or technical context. One proposition may require several categories, and one technical element may contribute to more than one informational function.

Table 5.1: Metadata-need categories used by the framework

Need	**Required informational function**	**Typical boundary**
Identity and description	Distinguish dataset, distribution, catalogue record, service, policy, decision, evidence item, and version; describe content and access interfaces.	Identifier or description does not establish the legal classification, continued availability, or entitlement to use.
Legal-source trace	Resolve the exact act, expression, subdivision, relevant date, and dependency on which the proposition relies.	Source identity does not decide applicability or make the extracted interpretation authoritative.
Actors and roles	Identify organisations or people and state their legal or organisational role for a defined activity, object, purpose, and period.	A role assertion does not confer statutory competence or prove the asserted status.
Provenance and method	Record collection, generation, use, derivation, transformation, agents, plans, time, and links to method evidence.	A lineage account does not prove completeness, lawful origin, or faithful execution of the stated method.
Purpose and use conditions	Represent asserted permissions, prohibitions, duties, restrictions, licences, purposes, recipients, and temporal or spatial constraints.	Policy syntax does not prove issuer authority, validity, enforceability, priority, or legal effect.
Quality and uncertainty	Link a resource to dimensions, metrics, values, annotations, methods, population or context, limitations, and evidence.	A recorded measurement does not select an adequate metric or establish fitness, representativeness, or truth.
Decision and evidence	Connect an assertion or review outcome to its rationale, evidence, responsible actor, date, and validity period.	Attribution and preservation do not establish evidential completeness, persuasiveness, or decision-maker competence.
Change	Relate successors, revisions, changed evidence, and superseded mappings to the earlier objects they may affect.	A detected dependency does not determine whether the change is legally material.

These needs lead to three forms of technical assistance and one organisational response. A record makes a value or relation retrievable. An evidence link makes supporting material addressable without claiming that all of its content has been represented. A closed check tests a declared syntactic or structural condition. Responsible review preserves the materials and accountable conclusion for a question that cannot be settled by those technical functions.

5.2.4 Stage 4: exact standard element

Once the need has been stated, Stage 4 compares it with an exact element in a versioned technical specification. The unit is not the name of a vocabulary in the abstract. It is a class, property, relation, constraint, or profile requirement whose issuer, version, date, definition, formal status, domain, range, cardinality, and conformance implications have been checked. Selecting DCAT, for example, says too little: a mapping must distinguish among `dcat:Dataset`, `dcat:CatalogRecord`, `dcat:Distribution`, `dcat:DataService`, and the properties used to relate them. Selecting PROV-O similarly requires identification of whether the need concerns an entity, activity, agent, generation, use, derivation, attribution, association, or qualified role.

The comparison asks whether the element denotes the same kind of entity or relation, whether it operates at sufficient granularity, whether its temporal and actor model is compatible, and which additional profile rules or evidence are required. It also preserves the element's formal status. A W3C Recommendation, a Commission draft profile, a Working Group Note, and a community specification have different issuers and claims. Technical use does not make an element a harmonised standard, an adopted common specification, or a source of legal presumption. A later release cannot be silently substituted for the version selected for a particular application.

The method favours reuse and composition. Existing vocabularies already provide strong components for legal-resource identification, catalogue description, provenance, quality assertions, policies, and validation. A bounded application profile can coordinate those elements and add constraints for a specific purpose. A local extension is considered only where the information is genuinely representable, no examined candidate element expresses it with sufficient fidelity, and the proposed term can be linked to the need and gap that justify it. An extension cannot solve a problem of truth, authority, or substantive judgment.

5.2.5 Stage 5: representational outcome

The comparison culminates in Stage 5, which assigns a qualitative representational outcome and records the boundary that remains. The framework defines four possible outcome families. They provide a method for classifying a comparison, not a score of legal coverage or a summary of empirically observed results.

An explicit gap and a referral to judgment are legitimate analytical outcomes. They

Table 5.2: Representational outcomes defined by the framework

Outcome	Meaning within the framework	Boundary that remains
Direct or composite support	One exact element represents the bounded need, or several elements together provide the required descriptive relations.	Support concerns representation within a stated profile; it does not prove truth, authority, completeness, or compliance.
Partial or evidential support	An element represents only part of the need or links an assertion to supporting material without carrying the decisive meaning.	The omitted qualification, interpretation, or evidentiary assessment must remain explicit.
Gap or extension required	No examined standard element adequately represents the identified need, or the available elements are too general.	The gap remains open or justifies a narrowly governed extension; generic free text is not relabelled as semantic coverage.
Non-metadata judgment	The question concerns applicability, authority, legal effect, factual truth, evidential sufficiency, or substantive adequacy.	Metadata may record the question, evidence, reviewer, rationale, and outcome, but it must not generate the determination.

indicate that the chain has reached its proper boundary instead of forcing every legal phrase into a metadata field. This is particularly important for terms such as "sufficient", "appropriate", "representative", "strictly necessary", and "where applicable". A profile may record relevant facts and make missing information visible, but the term's application normally depends on context and accountable reasoning. The recorded outcome should therefore state what may be represented and precisely what cannot be concluded from that representation.

5.2.6 Transversal Safeguards

Responsibility, evidence, versioning, and reassessment accompany every stage. Their role is summarised in Table 5.3. The table is deliberately organised by principle rather than by an additional sequence.

Documentation and metadata support accountability by preserving assertions, evidence, reasoning, and decisions. Their completeness does not, by itself, establish the accuracy of those assertions, the adequacy of the evidence, or substantive compliance. Producing and maintaining documentation may itself be a legal requirement, but satisfying that requirement does not automatically establish compliance with the underlying substantive obligations.

5.3 Operationalising the Analytical Chain

The five stages specify the relations that an accountable analysis must preserve while leaving the choice of database, software product, and organisational workflow open. Their practical

Table 5.3: Transversal safeguards for the five-stage chain

Principle	**Function across the chain**	**What it does not do**
Responsibility	Attribute source selection, legal interpretation, metadata production, semantic mapping, maintenance, review, and decision authority to defined people or institutions.	An internal assignment or metadata role does not create a statutory role or transfer legal authority.
Evidence	Link factual assertions, descriptions, methods, quality claims, policies, decisions, and reviews to addressable supporting material and its provenance.	An evidence link does not authenticate the material, establish completeness, or determine its legal weight.
Versioning	Pin the legal source, proposition, dataset or other governed object, standard element, profile, evidence, and mapping used for the reported outcome.	A version identifier does not show that the selected version is substantively adequate or still controlling.
Reassessment	Use recorded dependencies to identify which parts of the analysis may need review when a source, dataset, purpose, role, standard, evidence item, or decision changes.	The competent person or institution determines whether the change is material and what reassessment it requires.

value depends on whether the legal source, reviewed proposition, governed object, metadata record, supporting evidence, decision, and review state remain connected after the initial mapping has been completed. In a data space those objects may be held by different participants, and the responsibility for maintaining their connections therefore becomes an organisational question as much as a semantic one.

Data stewardship supplies the continuity function used in this chapter. As introduced in Section 1.3 and developed in Chapter 2, it is treated as an organisational capability rather than a statutory office or a universal job description. Responsible-stewardship research emphasises that the work is iterative and systemic, while lifecycle-oriented accounts connect it with continuing attention to metadata, provenance, quality, access, reuse, and coordination [16, 72, 73]. Public-sector guidance similarly distinguishes data governance, which allocates authority and decision rights, from the daily implementation and coordination performed through stewardship; it also recognises that the relevant responsibilities may be centralised, distributed, or combined with other roles according to institutional context [185].

Within the proposed framework, the data-stewardship function is assigned operational responsibility for maintaining the continuity of the record. This allocation means keeping the selected source, proposition, dataset or other governed object, metadata, evidence, responsible actors, decisions, versions, and review state connected over time, and identifying when a change should be returned to the competent actor. It neither creates a legal role nor transfers authority to decide applicability, legal effect, evidential sufficiency, or substantive adequacy. A data owner or authorised body may determine the governing policy, a legal or professional actor may decide the open question, and a technical custodian may operate the platform. Models developed for urban data spaces make this distinction concrete by separating decision bodies and data owners from data stewards and technology stewards, while also showing that the exact allocation is contextual rather than universal [186].

References to a "competent actor" denote competence for a specified function, rather than one person exercising all forms of authority. Four functions must be distinguished. Legal interpretation is undertaken by persons with relevant legal expertise and an identified assessment mandate. Factual evidence is supplied by data holders, custodians, domain specialists, or project participants responsible for the information concerned. Technical assessment is performed by specialists able to examine the relevant data, methods, systems, or validation results. The final institutional decision is taken by the person or body authorised under the applicable legal and institutional arrangements, drawing on those contributions.

These functions may be distributed across organisations or combined within one role, but their contributions and authority should remain separately attributable. The record should identify who supplied the evidence, who interpreted the legal source, who performed the technical assessment, and who adopted the decision, together with their respective scope of responsibility. The data-stewardship function maintains these connections and routes unresolved questions; performing that function does not itself confer legal assessment or decision-making authority.

The allocation must also be legible to participants who receive or reuse the record. An asserted role should indicate the activity, object, purpose, and period for which it is claimed; a decision should identify its issuer and scope; and a technical validation should identify the data and constraints against which it was performed. These distinctions allow a downstream participant to establish what has already been reviewed and which conclusion, if any, can be relied upon in the new configuration. They also expose where the chain ends because no competent decision or supporting material has yet been recorded.

5.3.1 A governed, source-linked corpus

The analytical chain can be maintained as a corpus of connected records. "Corpus" in this sense does not mean a collection of legal texts alone. It includes the reviewed analytical objects that allow a reader or authorised participant to reconstruct how an exact provision led to a proposition, how that proposition generated an information need, which technical element was compared with that need, and why the comparison was classified as support, a limitation, a gap, or a matter for judgment. Each record retains its own subject, provenance, version, responsible actor, and relation to the preceding and following stages.

Source linkage is the first governing constraint. A proposition must resolve to the precise provision from which it was derived; a metadata need must resolve to the proposition that makes the information relevant; and a mapping must resolve both to that need and to the controlling definition of the technical element. Requirements-engineering research describes related qualities through traceability, validity, auditability, and continuity, while warning that a trace cannot guarantee complete or correct legal interpretation [74, 113]. Evidence-based compliance engineering extends the same logic across legislation, formalised requirements, cases, evidence, and change [106]. The corpus proposed here adopts that re-

constructable direction without treating the links as proof of the propositions they connect.

Governance enters because the links require maintenance. Stable identifiers make records addressable, but someone must determine whether two identifiers denote the same object, whether a new dataset version supersedes the one assessed, whether a source or technical definition has changed, and whether the evidence still supports the recorded assertion. The data-stewardship function coordinates that maintenance and routes questions to legal, domain, technical, data-protection, security, or institutional actors. Those actors remain visible in the record because the meaning of an assertion depends not only on its content but also on who was entitled or qualified to make it, for which object, purpose, and period.

The four transversal safeguards therefore become properties of the corpus rather than additional stages. Responsibility attributes each interpretation, mapping, maintenance action, and decision. Evidence makes the supporting material addressable without collapsing it into the assertion. Versioning fixes the state examined. Reassessment uses recorded dependencies to locate the parts of the analysis that may need review. These properties can be implemented through different repositories and exchange arrangements; the conceptual requirement is that their relationships survive organisational and technical boundaries.

Continuity requires record states as well as identifiers. A proposition, mapping, evidence item, or decision may be proposed, reviewed, accepted for a defined use, superseded, disputed, or withdrawn. The framework does not prescribe one universal status vocabulary, but an operational arrangement must prevent a later user from confusing a draft interpretation with an approved decision or a superseded assessment with the current one. Revision should create an attributable successor or change record rather than silently replacing the basis on which an earlier conclusion was reached.

The quality of this corpus can be assessed at a technical and organisational level without turning that assessment into a legal conclusion. A profile may test whether required identifiers, provenance links, issuers, dates, and review states are present; stewardship may identify stale evidence, unresolved dependencies, orphaned decisions, or records whose responsible actor is no longer available. Those checks indicate whether the analytical account is maintainable and ready for review. They do not show that the underlying facts are true or that the recorded decision is substantively correct.

A governed corpus also preserves negative and unresolved outcomes. If no examined standard element adequately represents an identified need, the gap remains linked to the source and the comparison that exposed it. If a question depends on authority, factual truth, proportionality, sufficiency, or another form of accountable judgment, the record may retain its evidence, reviewer, reasoning, date, and outcome while leaving the determination outside metadata. This prevents later users from treating an empty field as proof that the issue was considered or a populated field as proof that it was resolved.

5.3.2 From a bounded exploration to operational use

The framework can be explored through a bounded selection of legal sources and technical specifications. In this thesis, the selected material makes the analytical stages and the two micro-applications concrete. It is not presented as a complete account of the applicable legal corpus, a representative evaluation of existing standards, or evidence of how representational gaps are distributed across the legal and technical landscape.

An operational corpus would begin from a defined activity rather than from the thesis's research question. Alongside the applicable legal sources and reviewed requirements, it would need to identify the actual dataset and version, intended operation and purpose, participants and roles, access route, policies and permits, quality and provenance evidence, responsible decision makers, decisions already taken, and conditions for later review. Its contents would change as data are updated, transformed, shared, retained, deleted, or assigned to a new purpose. The corpus would thus support a continuing assessment record rather than a completed research release.

This distinction prevents the exploratory exercise from being mistaken for an operational implementation. The bounded application illustrates how the five-stage chain can be followed for selected aspects of one identified provision. It does not demonstrate that an administration can maintain the same relations at reasonable cost, that the proposed allocation fits every institution, or that the availability of the record improves the outcome of a disputed assessment. Those questions require implementation with real organisational actors and evaluation over time.

The exploratory material can inform a later organisational application by showing the kinds of relationships that the framework asks an assessor to examine: identity, legal-source trace, roles, provenance, purposes and use conditions, quality, evidence, decisions, and change. These are design categories defined by the framework, not recurring results established across the legal and technical landscape. Transfer remains conditional: each organisation must repeat the source and applicability analysis for its facts, select the relationships relevant to its tasks, and allocate decision and maintenance responsibilities within its own legal and institutional setting.

Data stewardship connects that design to continuing practice. The steward does not inherit the thesis's interpretations or convert the exploratory material into an organisational compliance file. The function instead maintains the operational relationships accepted by the organisation, records their provenance and review state, coordinates corrections, and makes a material change visible to the competent actor. In a federated data space the function may itself be distributed: one participant can maintain the catalogue record, another the source dataset and provenance, and another the access decision or quality assessment. Continuity then depends on shared identifiers and exchangeable relations as well as an agreed route for correction and reassessment.

An operational record would normally be assembled around an assessment event. A pro-

posed access, transformation, sharing operation, or training use supplies the configuration against which relevant sources and prior decisions are retrieved. The resulting record can collect the facts and evidence required for review, preserve the conclusion reached for that configuration, and state which changes would reopen it. A later purpose or dataset version may reuse part of that material, but reuse remains explicit: the new assessment identifies which earlier propositions and evidence are carried forward and which require renewed consideration.

At this transition, stewardship contributes most directly to assessment and compliance work. The function can prepare an attributable assessment package, identify missing contributors, maintain the decision after it has been made, and return it for review when one of its dependencies changes. It cannot guarantee that every applicable source has been found or that the competent actor will reach the correct conclusion. Its contribution is to prevent a decision from becoming detached from the configuration, evidence, and reasoning that made it intelligible.

5.3.3 An assessment-supporting knowledge graph

A knowledge graph is one possible technical realisation of these recorded relations. Graph-based models represent identifiable entities and the relations among them and can support query and validation over information drawn from heterogeneous sources [187]. In the present framework, possible nodes include the legal source and subdivision, proposition, dataset and version, activity, actor and role, policy, evidence item, decision, standard element, mapping, and review. Edges express relations such as derivation, description, attribution, use, dependence, support, supersession, or responsibility.

Graph representation is not a prerequisite for applying the conceptual framework. An organisation may preserve the same relations through linked registries, relational databases, document-management systems, or an exchange package assembled for a particular review. The relevant criterion is whether the route from source to outcome remains addressable and reconstructable across those components. A knowledge graph becomes especially useful where heterogeneous objects and many-to-many dependencies must be queried across organisational boundaries, but selecting that implementation remains a separate architectural decision.

This representation would allow an assessor to move through the chain rather than search separate documents independently. A query could retrieve the source and interpretation on which a metadata requirement depends, identify the evidence associated with a quality assertion, distinguish the dataset from its catalogue record and distributions, or locate mappings affected by a revised standard element. Regulation-to-standards work has already demonstrated that normative concepts, requirements, standards, and implementation or tracking responsibilities can be published as navigable and queryable relations [107]. Evidence-based compliance models and recent compliance-as-code implementations similarly show how legis-

lation, requirements, controls, cases, versions, and evidence can be connected for inspection and change analysis [106, 108].

The graph supports assessment through retrieval, comparison, and routing. Some questions are demonstrative: the requested record or evidence is present and can be returned. Others permit a closed structural check, such as testing whether a required identifier, relation, controlled value, or cardinality has been supplied. Evaluative questions remain open even when the related information is complete, because the conclusion depends on further criteria or competent judgment. Research on queryable provenance for GDPR compliance makes a comparable distinction and observes that, where compliance cannot be evaluated automatically, the presence of the required metadata can still be checked and the related information retrieved for assessment [188].

Dependency queries can assist reassessment in the same limited sense. When a legal source, dataset version, purpose, role, evidence item, decision, or technical definition changes, recorded links can identify propositions, needs, mappings, and reviews that depend on the changed object. The resulting set is a review queue, not a legal conclusion. The competent person or institution determines whether the change is material, whether previous reasoning remains applicable, and which part of the analysis must be repeated.

Graph representation also makes omissions and disagreements more visible. A proposition without a source, a quality claim without a method, a policy without an issuer, or a review without an identified object can be queried as an incomplete relation. Conflicting assertions can remain separately attributed instead of being overwritten by one unqualified value. The graph thereby improves inspectability of the recorded account. Its epistemic boundary remains unchanged: W3C guidance on provenance expressly separates access to a provenance record from trust in its authority or correctness [103]. A complete-looking graph may still contain an inaccurate assertion, incomplete evidence, an unauthorised decision, or an incorrect legal interpretation.

The knowledge graph is an implementation possibility rather than an output of this thesis. The framework identifies the kinds of entities and relations that such a graph could contain, but the study has not deployed a shared graph, connected it to administrative systems, tested access controls, measured maintenance costs, or evaluated its use by legal and operational staff. Describing the possible representation clarifies how the framework could support organisational work; it supplies no evidence of production readiness or practical effectiveness.

Useful queries must remain configuration-sensitive. A request for the policy attached to a dataset differs from a request for the decision authorising a particular actor, purpose, operation, and period; retrieving every node labelled as a licence or legal basis would obscure that distinction. The graph should therefore support navigation through qualified relations and versions rather than offer a global “compliant” state. Its outputs may include a source trail, an evidence package, a list of unresolved relations, or a queue of records affected by

change, each of which is an input to assessment rather than its legal result.

Graph governance is equally necessary in a federated setting. Participants may expose only selected metadata while retaining confidential evidence or decisions in controlled repositories. A relation can identify the existence, custodian, access route, and review status of an evidence item without publishing its protected content. Access control, identity management, retention, and security remain implementation requirements beyond the semantic model developed here, but the conceptual chain requires that restricting an artifact should not erase its place in the assessment or the responsibility for making it available to an authorised reviewer.

5.3.4 Standards as modular building blocks

No single specification covers the full chain. The standards examined in this study instead provide building blocks for different informational functions, and their use begins only after a need has been stated without reference to a preferred vocabulary. The unit of comparison is an exact class, property, relation, constraint, or profile requirement whose issuer, version, definition, formal status, subject, value model, and conformance implications have been checked. This element-level discipline keeps a broad vocabulary label from standing in for the particular semantics required by the analysis.

Core ELI supports identification of legal resources and subdivisions, while DCAT and DCAT-AP distinguish catalogues, catalogue records, datasets, distributions, and data services and provide the principal descriptive relations used for discovery [26, 89, 98]. These vocabularies can identify the source and governed resource at the required granularity. They leave the extracted legal proposition, concrete applicability, truth of the description, and authority to access or use the data to the surrounding assessment.

PROV-O, DQV, and ODRL address different relations around that resource. PROV-O connects entities, activities, agents, derivations, uses, associations, roles, and plans; DQV represents quality dimensions, metrics, measurements, annotations, and policies; ODRL expresses asserted permissions, prohibitions, duties, parties, actions, assets, and constraints [27, 28, 29]. Their combination can make a collection method, quality assessment, or stated use condition attributable and retrievable. The method, metric, issuer, scope, and supporting evidence remain material because representing an assertion neither authenticates it nor establishes its legal effect.

SHACL adds closed structural validation. A shape can require a property, datatype, class, controlled value, pattern, or cardinality and can report where supplied RDF data do not conform to those declared constraints [189, 190]. This is useful for questions whose assumptions and expected structure have already been fixed. A successful validation result remains specific to the data graph, shapes graph, processor, and entailment assumptions; it does not answer open questions of accuracy, adequacy, authority, or compliance.

HealthDCAT-AP, MLDCAT-AP, Croissant, and Croissant Responsible AI add more spe-

cialised detail for health catalogues, machine-learning datasets, files, record sets, features, collection, transformations, limitations, bias, and related documentation [31, 94, 95, 191]. Their greater granularity is useful only where the subject and release status fit the need. A health-specific access-body relation or an ML documentation field cannot be generalised to every data space merely because its label is convenient.

Application profiles provide the principal means of composition. A profile can select elements from several namespaces, fix their versions, define required and optional values, constrain ranges and cardinalities, use controlled vocabularies, and add validation shapes for a stated community and purpose [90]. Composition must preserve the meaning of each reused element. A generic provenance role cannot be renamed as a statutory status, and an asserted ODRL permission cannot be treated as proof that the represented use is lawful.

Formal status remains visible throughout this technical stage. A W3C Recommendation, a Working Group Note, a European application profile, a draft profile, and a community specification make different claims. Use within the framework does not convert any of them into a harmonised standard, adopted common specification, or source of legal presumption. The modular outcome is therefore a conceptual profile linking legal-source trace, catalogue identity and description, provenance, quality, asserted policy, supporting evidence, and structural checks. It is neither an adopted EU profile nor a production compliance architecture.

Modularity is particularly important in a data space because participants need not operate the same internal systems or maintain every part of the assessment. A shared profile can state which identifiers and relations must cross the organisational boundary, while local records preserve more detailed evidence, decisions, or sector-specific concepts. Instrument-specific extensions can then add the actors, procedures, conditions, and legal effects required by one regime without redefining the reusable core. Their governance must identify who maintains the extension, how changes are communicated, and which version was used for an assessment; otherwise, formal interoperability at the exchange layer would conceal divergent meanings underneath it.

5.4 Applying the Conceptual Framework to Data Act Article 33(1)(a)

Data Act Article 33(1)(a) provides a bounded application of the five-stage chain. The application does not attempt to map the whole provision or evaluate the coverage of the standards landscape. It selects two aspects that expose different functions of the framework: the technical representation of a relation relevant to the collection methodology, and the qualification that the listed information must be "sufficiently described".[1]

[1] Article 33(1)(a) provides that "the dataset content, use restrictions, licences, data collection methodology, data quality and uncertainty shall be sufficiently described, where applicable, in a machine-readable

The two micro-applications are illustrative demonstrations of the framework's internal coherence and intended operation. They show how the five stages preserve source linkage and representational boundaries.

Article 33(1)(a) forms part of a wider interoperability framework. Article 33(3) establishes a presumption of conformity for participants meeting harmonised standards whose references are published in the Official Journal, limited to the essential requirements covered. Paragraph 4 provides for standardisation requests, while paragraph 5 allows the Commission to adopt common specifications under specified conditions; paragraph 8 attaches a similarly limited presumption to conformity with those specifications. The use of PROV-O, DCAT-AP, or SHACL in the following examples does not by itself establish either presumption or overall legal compliance [32, art. 33(3)–(5), (8)].

Stage 1 fixes the legal source as Regulation (EU) 2023/2854, Article 33(1)(a). The record identifies the official act and exact subdivision, together with the language, date, version, and minimum contextual provisions needed to understand the relevant actors and relationship. The output is therefore a verified and addressable source, not an interpretation of whether the provision applies to a particular organisation.

At Stage 2, the provision is expressed as a source-linked proposition. It concerns a participant in a data space that offers data or data services to another participant. The dataset content, use restrictions, licences, data-collection methodology, data quality, and uncertainty must be sufficiently described, where applicable, in a machine-readable format, so that the recipient can find, access, and use the data. The proposition preserves the offering participant, recipient, governed object, required description, informational categories, functional purpose, and qualifying expressions. It does not determine whether particular parties occupy those legal roles or whether a particular description is sufficient.

The proposition remains intact at Stage 2. Selecting two aspects for closer examination does not convert the reference to collection methodology or the qualification of sufficiency into autonomous obligations. It allows the later stages to test bounded information needs while retaining the conditions and purpose of the provision.

5.4.1 Collection methodology

For this aspect, Stage 3 derives a deliberately narrow metadata need: connect an identified version of the offered dataset to the asserted collection activity that generated it, the agent associated with that activity, and the plan or method intended for the activity. The need does not prescribe PROV-O in advance. It states the entities and relations that an assessor would need to retrieve before candidate elements are examined.

At Stage 4, PROV-O supplies exact elements for this bounded need. The dataset version can be represented as a `prov:Entity` and the collection event as a `prov:Activity`; `prov:wasGeneratedBy` relates the entity to that activity. `prov:wasAssociatedWith`

format". Its stated function is to support discovery, access, and use by the recipient [32, art. 33(1)(a)].

can identify an associated agent. Where the intended plan must be stated explicitly, `prov:qualifiedAssociation` connects the activity to a qualified association and `prov:hadPlan` connects that association to a `prov:Plan` [28]. These elements retain their published provenance semantics; their use here does not redefine them as legal concepts.

Stage 5 records *composite support* for the bounded metadata need. The elements together can represent the asserted relation among a dataset version, collection activity, agent, and plan. They do not necessarily contain the substantive account of the whole collection methodology, prove that the activity occurred as recorded, demonstrate fidelity to the plan, or establish that the methodology is appropriate. The outcome is therefore a representational contribution with an explicit boundary, not a conclusion that Article 33(1)(a) has been implemented.

5.4.2 Sufficiency of the description

The second micro-application reaches a different outcome. At Stage 3, the need is not a Boolean metadata value stating that the description is sufficient. The record must instead preserve the description supplied, the dataset or service and version to which it refers, the recipient task and context against which it was reviewed, the supporting evidence, the reviewer, the rationale, and the resulting decision.

At Stage 4, the DCAT-AP 3.0.1 Dataset description requirement, which reuses `dct:description`, is examined as an exact profile element capable of carrying the supplied dataset description, while a SHACL property constraint using `sh:minCount` can test whether a declared description is present under a specified shape [26, 189]. Neither element denotes or decides the legal qualification "sufficiently described". Sufficiency depends on the provision's purpose, the information supplied, the recipient's ability to find, access, and use the data, and the facts of the particular offering.

Stage 5 therefore classifies sufficiency as a *non-metadata judgment*. Metadata can make the relevant description, evidence, attribution, and review state retrievable; structural checks can identify missing required relations under a declared profile. The competent legal or institutional actor must still determine whether the description is sufficient and whether Article 33(1)(a) applies to the configuration.

Taken together, the two micro-applications complete the five-stage sequence without presenting the entire provision as mapped. The first reaches exact technical elements and a bounded representational outcome; the second shows why the same framework must preserve a question for competent judgment. An operational application would additionally require the actual dataset and version, the participants and their roles, the intended activity, available evidence, the responsible decision maker, and the conclusion reached for that configuration.

The data-stewardship function would maintain the continuity of this operational record. It would keep the source, dataset, description, policies, evidence, decisions, and versions

connected; identify missing or superseded records; and route a material change to the competent actor. A knowledge graph could support this work by making dependencies and affected records queryable, but it would not decide whether the provision applies or whether the description is sufficient. Those determinations remain with the person or institution competent to assess the relevant facts and legal context.

The example therefore illustrates the division of labour established by the framework. Technical representations can make information retrievable, related, and structurally testable. Data stewardship can preserve that information and maintain the route to review. Neither function determines legal applicability, descriptive sufficiency, access authority, or the lawfulness of a later use. The contribution of the framework is to keep those questions, their supporting materials, and the actors responsible for answering them connected rather than to replace the assessment itself.

5.5 Sectoral Transfer: The EHDS

The European Health Data Space provides a concise illustration of how the framework may be transferred to a sector-specific setting. Article 77 requires health data access bodies to provide descriptions of available datasets and their characteristics through standardised, machine-readable catalogues [1]. This raises an informational problem comparable to Data Act Article 33(1)(a), but the relevant actors, objects, required content, procedures, timetable, and legal effects are specific to the EHDS. Those differences require the analysis to begin again from the EHDS source and its concrete configuration, and prevent the Data Act micro-applications or their outcomes from being transferred unchanged.

Article 77(3) requires the dataset catalogue, rather than the electronic health data described in it, to be made available from 26 March 2029 to the single information points established or designated under Article 8 of the Data Governance Act [1, 56]. Applying the framework would begin again with the exact provision and its context, identify the catalogue and responsible actors, derive the information needed to describe and review the availability relation, and then examine candidate standard elements. The provision does not itself prescribe a technical connection, and recording an endpoint, exchange, or validation result would not establish that the legal duty had been fulfilled.

What transfers across the two settings is therefore the analytical method rather than a completed mapping or representational outcome. Standards used for catalogue description, legal-source identification, provenance, or structural validation may reappear as candidates, but their suitability must be reassessed against the EHDS provision and the concrete configuration. The catalogue must also remain distinct from the health data it describes: its availability neither grants access to those data nor authorises their subsequent use. Other Common European Data Spaces may define different actors, procedures, and informational requirements, so each application must repeat the five stages from its own authoritative

sources.

5.6 A Prospective Minimal Profile

The framework could inform the development of a minimal, source-linked application profile for a defined organisational and legal setting. Such a profile would select only the relations needed to connect the authoritative source and reviewed proposition to the governed object, its description or asserted policy, supporting evidence, responsible actors, relevant versions, decision, and review state. It would not itself constitute a new vocabulary, a harmonised standard, or a compliance architecture.

The exact elements included in such a profile would have to be selected case by case from the modular specifications discussed above and assessed against the relevant information need, context, and specification version. A controlled extension should be considered only where the identified need cannot be represented adequately by the elements examined. Any conclusion of this kind would remain bounded by that examination and would not establish the absence of a suitable element from every existing or future standard.

The governance of the profile would be organisational as well as semantic. Data stewardship could preserve identities, provenance, versions, evidence links, and routes to review, while legal, domain, technical, and institutional actors retained their respective judgments. A shared catalogue or exchange record would not need to expose the complete operational record: it could identify the existence, custodian, version, and access route of an assessment, permit, or evidence item while protected content remained in a controlled repository. These arrangements could keep decisions connected to the dataset, evidence, and configuration to which they referred without turning technical representation into a compliance determination.

5.7 Validity and Limitations

The contribution rests on making the source, interpretive step, information need, candidate technical representation, and remaining judgment distinguishable and reviewable. The Data Act micro-applications illustrate the framework's internal coherence and intended operation, without independently validating its effectiveness, operational feasibility, or standards coverage.

The legal and technical materials considered are selective and time-bound. Other provisions, interpretations, implementing measures, national rules, cases, or later specification versions may introduce different information needs or candidate elements. A concrete application should therefore repeat the five stages for its own sources, actors, objects, purposes, and technical context rather than transfer an earlier outcome unchanged.

The framework has not yet been evaluated through a longitudinal public-administration

project or a live data-space transaction. A future pilot could examine its usability and maintenance burden and whether participants can locate sources and evidence, distinguish recorded assertions from legal effects, allocate responsibilities, and recognise reassessment triggers [192, 193]. This is the next empirical step rather than a condition for presenting the framework as a conceptual method.

5.8 Conclusion

The analysis connects two complex bodies of work: EU legal rules governing data sharing, access, transfer, and reuse, and the technical standards used to describe datasets, provenance, quality, policies, and structural constraints. The central difficulty lies in the absence of a structured method for examining how a selected legal concept may be represented and where technical representation remains incomplete or dependent on judgment.

The proposed framework responds through a five-stage chain: authoritative legal source, source-linked legal concept or requirement, metadata need, exact standard element, and representational outcome. The two Data Act micro-applications make this sequence concrete by reaching, respectively, composite technical support and a non-metadata judgment. The EHDS illustration shows why the analysis must be repeated when the sectoral actors, objects, procedures, and legal effects change. The framework preserves possible technical support, representational gaps, and questions reserved for competent judgment without treating metadata as a substitute for legal assessment.

In operational use, the framework could guide the construction of a governed corpus and, where appropriate, an application profile or assessment-supporting knowledge graph. Data stewardship could maintain the continuity of sources, governed objects, evidence, versions, decisions, and review states, while legal, domain, technical, and institutional actors retained their respective authority. The contribution of the chapter is therefore the conceptual and methodological connection between legal analysis and metadata representation, not a completed mapping or an automated determination of compliance.

Chapter 6

Synthesis and Conclusions

6.1 The Dataset Journey Revisited

The thesis began with public actors preparing a health dataset to be made discoverable through a European data space. Before any downstream use, the resource had to be identified and described, its provenance, quality, and applicable conditions documented, and responsibility for maintaining that information allocated. A public research organisation could subsequently discover the dataset and propose to use it for a specific purpose, including the development or validation of an AI system for a hospital. Discovery and access were necessary, but neither determined whether that operation could proceed. The example therefore exposed two connected challenges: public administrations must first be able to supply and govern data within CEDS, and must then assess the authority and conditions governing each proposed use.

A minimal hypothetical example makes this interaction concrete. Suppose that the public research organisation considers a dataset version for validating an AI system for a hospital, but the available description does not clearly document its collection methodology. Table 6.1 connects possible deliberative outputs with the collection-methodology micro-application already presented in Section 5.4.

Table 6.1: A bounded illustration of the health-data journey

Component	Illustrative output
Canvas discussion	Two questions from the Data area: "Who collects the data?" and "Is data well documented?". A model researcher, integration engineer, and hospital domain professional contribute through the three macro-roles; the data steward is consulted for supporting documentation.
Recorded response	The collection methodology remains unclear for the selected version. The group keeps the suitability assessment open, requests the relevant documentation, and assigns responsibility for obtaining it.
Legal analysis	Assuming that Article 33(1)(a) applies to this offering, the participant offering the dataset must describe how the data were collected. To assess this requirement, the description and its supporting evidence must be retrievable and linked to the dataset version being considered.
Technical representation	Stage 4 reuses `prov:Entity`, `prov:Activity`, `prov:Agent`, `prov:wasGeneratedBy`, and `prov:wasAssociatedWith`. These elements represent part of the asserted provenance examined in the existing micro-application.
Outcome	The example produces a documentation request and a proposed representation of the dataset's provenance. The accuracy of that provenance, the sufficiency of the description, the dataset's suitability, and the authority for the intended use remain to be assessed.

A small conceptual graph could preserve the following links:

Record category	Relation	Record category
Article 33(1)(a)	grounds	Legal proposition
Legal proposition	motivates	Metadata need
Metadata need	concerns	Dataset version
Dataset version	`prov:wasGeneratedBy`	Collection activity
Collection activity	`prov:wasAssociatedWith`	Associated agent
Assessment	assesses	Dataset version

Only the relations prefixed with `prov:` reproduce standard properties; the other labels indicate conceptual links, not a proposed vocabulary. "Review" records an open assessment and the outstanding documentation request. The Canvas makes an evidence need visible, while the five-stage chain independently examines a source-linked legal proposition. Their records can refer to the same dataset without forming a necessary or automatic pipeline.

Such a study would test the proposed organisational relationship between the Canvas and CEDS, rather than assume its effectiveness. It could address two questions: does Canvas use identify material evidence needs or responsibility gaps that existing procedures overlook; and are the concerns identified subsequently addressed through specialist assessment, revised requirements, or an authorised decision? Evaluation should trace these consequences over time and record the organisational effort involved.

The journey can now be reconstructed across three connected levels. Strategic and infrastructural conditions make the data available through institutions, standards, computing

resources, and organisational arrangements. Once that possibility exists, the public administration must determine what it seeks to achieve, who may be affected, which expertise is required, which risks and values matter, and what evidence should support a decision to continue or change course. Selected legal concepts and those decisions must in turn remain connected to the dataset, the applicable source, and the exact technical representations through which they are communicated or checked.

Each level performs a distinct function. Regulation does not create industrial capacity merely by declaring a right. Infrastructure may enable access, while the desirability and lawfulness of a public-sector AI project remain for determination through the competent legal and institutional processes. Participatory workshops structure deliberation without establishing technical quality or conferring legal authority, and metadata represents assertions whose truth and sufficiency still require assessment. The contribution of the thesis is consequently not a claim that these components form a completed compliance machine. It is a reasoned account of why public administrations engaging with Common European Data Spaces (CEDS) need institutional capacity, structured deliberation, and technical traceability to interact without being confused.

This position also explains the term "human-centred compliance", which denotes a continuing process extending beyond the completion of a checklist. People with relevant knowledge and authority use evidence and technical tools to formulate, document, challenge, and revisit decisions. Automation can retrieve information, validate structure, enforce a bounded policy, or signal change. People and institutions remain responsible for interpreting sources, evaluating evidence, hearing affected perspectives, resolving disagreement, and deciding whether a project may proceed. Human-centredness lies in the allocation of knowledge, participation, and authority, not in the mere presence of a person beside an automated system.

6.2 Operationalising EU Data and AI Regulation

The main research question asked:

> How can the complexity of EU data and AI regulation be translated into organisational and metadata-based tools that support public administrations engaging with Common European Data Spaces?

The answer is that translation should be treated as a controlled sequence of partial transformations within an adequately resourced institution. At the organisational level, a deliberative artifact can help technical and non-technical participants formulate a shared account of the public purpose, affected interests, data, methods, functional expectations, responsibilities, and evidence. At the informational level, a selected legal source can be reconstructed as a source-linked concept or requirement, translated into a metadata need,

compared with an exact standard element, and assigned a bounded representational outcome. At every point, provenance, responsibility, versioning, evidence, uncertainty, and reassessment must remain visible.

The qualification "partial" is decisive. A legal concept often contains evaluative language, institutional competence, exceptions, and relations to other sources. An organisational discussion may identify competing values without resolving them. A metadata standard operates with published semantics that may represent a description, policy, activity, agent, or quality assertion without representing the legal conclusion that motivated the record. Translation is therefore successful when it preserves what can be communicated and exposes what cannot, rather than when it forces every question into a machine-readable field.

For a public administration, the practical result is a division of labour that begins before data reuse. As a data holder or provider, it must identify and prepare relevant resources, maintain their descriptions, provenance, quality information, and conditions, and assign responsibility for keeping those records current. As a data user, it must assess whether access covers the proposed operation and what evidence and review remain necessary. CEDS and their associated services can support discovery, identity, negotiation, controlled access, exchange, and shared description. The AI Ethics Canvas can create a recurring space in which the administration examines the purpose of an AI project, the available evidence, the actors who should participate, and the effects of design choices. The conceptual framework can connect selected requirements and decisions to data descriptions, policies, provenance, quality information, and review triggers. Competent officials and professional functions determine the legal and institutional significance of those records, while data stewardship maintains their continuity across organisations, datasets, and project stages [17, 18].

These tools support an administration by helping it recognise dependencies, structure conversations, preserve reasons, identify missing evidence, and support technical actions within a defined scope, without displacing administrative procedure, procurement law, data-protection assessment, cybersecurity work, sectoral obligations, judicial interpretation, political responsibility, or professional judgment. The same metadata can support a sound or unsound decision depending on its source and maintenance. The same Canvas can support serious deliberation or ritual completion depending on facilitation, participation, authority, and follow-up. The appropriate claim is that the tools make important parts of regulatory operationalisation more explicit and reviewable.

6.3 Digital sovereignty, material capacity, and CEDS

Meaningful agency depends on more than the formal recognition of rights. Chapter 3 answers the first subquestion by preserving the distinction established in the EU–China comparison and then extending the analysis to the European data-space programme. A person may possess rights concerning data and automated decisions while still lacking meaning-

ful alternatives to dominant infrastructures and services. Regulation structures power and creates enforceable claims, but the exercise of those claims depends on material, industrial, organisational, and institutional capacity [9].

This thesis does not turn that comparative result into a general score of European or Chinese sovereignty. The two legal orders allocate state, corporate, and individual power differently, and the comparison exposes tensions within the European promise of individual control. The subsequent treatment of CEDS and the policy landscape through 2026 extends that analysis for this thesis.

The extension identifies CEDS as a material and institutional response to part of the capacity problem. A data space can provide common governance, catalogues, interoperability arrangements, trust services, and technical means for making data available across organisational boundaries. It can help European public and private actors discover and reuse resources under conditions that remain connected to European law and sectoral institutions [7]. In that limited sense, CEDS do more than announce an entitlement: they invest in an environment through which access and reuse may become practically possible.

They remain a partial response. Data spaces do not remove dependence on cloud providers, accelerators, advanced manufacturing, foundation models, platforms, connectivity, energy, or specialised labour. Nor do they ensure that a small municipality, health authority, or public research organisation can evaluate a provider, negotiate access, curate a dataset, run a model, or maintain evidence over time. Europe's AI Continent Action Plan and Data Union Strategy recognise the interdependence of data, computing, skills, adoption, and infrastructure [23, 24]. These policy programmes demonstrate institutional direction and announced measures; they do not, by publication alone, demonstrate that dependencies have been overcome.

CEDS can therefore strengthen European agency by making governed data resources and interoperable services more practically available. Their contribution to digital sovereignty remains conditional on the actors able to use them, the infrastructures on which they run, the alternatives available in the market, and the governance through which control is exercised. For public administrations, data availability becomes useful capacity only when it is matched by procurement competence, domain knowledge, legal and ethical review, security, technical expertise, stewardship, and authority to act on identified problems [17, 183, 194]. At that level, the capacity problem becomes an organisational problem of competence, deliberation, and authority.

6.4 The Canvas in public-sector AI design and review

For public-sector AI design and review, the formative evidence supports a specific role for the Canvas. Chapter 4 shows that it is best understood as a boundary object: a shared, visual artifact through which participants with different professional vocabularies can formulate

questions around data, algorithms, methods of analysis, socio-cultural considerations, and functional requirements. Its three roles distribute attention across relevant perspectives, while its five areas prevent the project from being reduced to model performance.

The direct formative evidence consists of twelve qualitative interviews and a workshop involving 18 participants. The workshop used two broader project cases and three scenarios. Those activities informed the wording, organisation, and practical use of the artifact. Surveys and focus groups occupy a different evidentiary level. They helped identify the public-sector context and the need for a multidisciplinary tool; they are not combined with the interviews and workshop as a single measure of effectiveness. Maintaining this separation avoids inflating the empirical basis of the study.

On that evidence, the Canvas can support question formulation, shared understanding, the identification of responsibilities and evidence needs, and the reconsideration of design choices. It can make disagreement visible before it becomes embedded in procurement or code. It can also help a team recognise that it lacks an affected perspective, a data-quality account, a public-purpose justification, a competent legal assessment, or a plan for monitoring. These are meaningful organisational functions even when the Canvas does not resolve the question it raises.

The thesis extends the artifact's role across three moments. An initial session can test whether the problem has been framed adequately before a solution is fixed. Iterative sessions can be triggered when the data, intended purpose, provider, model, deployment setting, or affected community changes. A final review can revisit the recorded questions and evidence before an authorised lifecycle transition. This use pattern is conceptually consistent with the Canvas's iterative design, but it has not been evaluated through a complete longitudinal project. It should therefore be described as a proposal for deployment rather than as an empirically demonstrated effect.

A final review may direct the administration to a data-protection impact assessment, procurement condition, security test, equality analysis, consultation, model evaluation, legal opinion, or decision by an authorised body; it does not certify the system. Completion of every area, the presence of a lawyer, or discussion of risks does not by itself establish ethical acceptability, legal compliance, or technical quality. The Canvas supports the organisation in identifying and connecting these activities without inheriting their authority.

The canonical reusable artifact is release 1.0.0 in the repository, licensed under CC BY-SA 4.0 [33]. The 2024 version documents an earlier point in the design process and is not treated as the operative reference for this thesis. Its availability supports inspection, adaptation, and further evaluation. It does not prove that an adaptation preserves the intended facilitation method or works equally well in every administration.

Within those evidentiary limits, the Canvas can provide a recurring organisational setting for multidisciplinary inquiry and review, provided that the administration allocates time, facilitation, evidence, participation, and decision authority. Its value lies in the quality and

continuity of the deliberative process, not in the completion of the graphic. The current evidence supports formative development and plausibility; longitudinal effects, comparative performance, institutional fit, and impact on actual decisions remain open.

6.5 The Conceptual Framework

The answer at the informational level is methodological. Chapter 5 relates selected EU legal concepts and requirements to metadata needs and exact standard elements through one five-stage chain:

$$\begin{aligned}\text{legal source} &\longrightarrow \text{source-linked legal concept or requirement} \\ &\longrightarrow \text{metadata need} \longrightarrow \text{exact standard element} \longrightarrow \text{representational outcome.}\end{aligned} \tag{6.1}$$

The order protects the analysis from vocabulary-driven reasoning. Starting with an available field and searching for a legal sentence that resembles it would encourage false equivalence. Starting with the authoritative source requires the analyst to preserve the addressee, action, object, purpose, condition, exception, temporal context, and uncertainty that determine meaning. A metadata need then states what information would have to be carried or linked without assuming that a particular standard supplies it. Only after that step is an exact, versioned standard element examined within its published semantics.

Data Act Article 33(1)(a) supplies the bounded application. It requires the dataset content, use restrictions, licences, data-collection methodology, data quality, and uncertainty to be sufficiently described, where applicable in a machine-readable format [32]. The provision remains one source-linked proposition, while its informational categories are treated in a structured form so that selected needs can be examined separately without converting them into autonomous obligations. The first micro-application derives the bounded need to connect a dataset version to its asserted collection activity, associated agent, and intended plan. The exact PROV-O elements `prov:wasGeneratedBy`, `prov:wasAssociatedWith`, `prov:qualifiedAssociation`, and `prov:hadPlan` provide composite support for those relations [28]. They do not prove that the activity occurred as recorded, that the plan was followed, or that the methodology was appropriate.

The second micro-application concerns the qualification "sufficiently described". The DCAT-AP 3.0.1 Dataset description requirement, which reuses `dct:description`, can carry a dataset description, and a SHACL constraint using `sh:minCount` can test whether a description is present under a declared shape [26, 189]. Neither determines whether the supplied information is sufficient for the recipient to find, access, and use the data. Stage 5 therefore records sufficiency as a non-metadata judgment for the competent actor. The two outcomes illustrate the framework's operation without claiming a complete mapping of

Article 33(1)(a) or measuring the coverage of existing standards.

The EHDS provides a shorter illustration of sectoral transfer. Article 77(3) concerns the availability of a dataset catalogue within a specific institutional arrangement involving health data access bodies and Data Governance Act single information points [1, 56]. It illustrates why the five stages must be repeated when the actors, objects, procedures, timetable, and legal effects change. It does not confirm a completed cross-sectoral mapping or make the EHDS the universal architecture of every Common European Data Space.

Responsibility, evidence, versioning, and reassessment traverse every stage of the chain. The source record identifies the exact provision, language, date, version, and minimum interpretive context; the legal proposition is an attributable interpretation; the metadata need is derived for a stated purpose; the standard element belongs to a particular specification; and the outcome records both possible support and limitation. Standards structure information and bounded checks, while evidence supports the assertions made in the concrete configuration. Data stewardship maintains these relations, identifies missing or stale records, and makes a material change visible. The competent legal, technical, or institutional actor determines the significance of the evidence and whether the change requires a new decision.

The five-stage method accordingly provides a disciplined, source-linked, and revisable way to examine how selected legal requirements may be related to standard-based technical representations and supporting evidence, while preserving the questions reserved for legal and institutional judgment. Existing standards form a modular toolkit rather than a legal model of the whole regulatory environment. The framework supports traceability and reviewability; it does not establish applicability, authority, factual accuracy, evidential sufficiency, substantive adequacy, or compliance.

6.6 Evidence and Proposals

Because the thesis connects contributions with different methods and evidentiary bases, its final claims must preserve their respective status. Table 6.2 distinguishes established bases and conceptual contributions from the developments proposed by the monographic synthesis.

Table 6.2: Evidentiary bases, conceptual proposals, and retained boundaries

Level	**Established basis or conceptual contribution**	**Thesis-level proposal**	**Boundary**
Strategic and infrastructural	Comparative legal analysis identifies a gap between formal individual control and the material capacity to exercise it.	Treat CEDS as one component of a wider European capacity strategy for data and AI.	No empirical proof that CEDS remove technological dependence or deliver sovereignty.
Organisational and deliberative	Interviews and a workshop support formative development of the AI Ethics Canvas.	Use the Canvas initially, iteratively, and at a final-review checkpoint in public-sector projects.	No longitudinal proof of effectiveness and no certification of ethics, legality, or technical quality.
Informational and technical	A five-stage conceptual framework specifies how legal sources, metadata needs, candidate standard elements, and representational boundaries should be related; Article 33(1)(a) illustrates its intended operation.	Use the framework to connect reviewed requirements, standard-based representations, supporting evidence, and review states in a concrete assessment.	No comprehensive mapping of Article 33(1)(a) or the standards landscape, production implementation, aggregate coverage claim, or automated legal conclusion.
Cross-cutting organisational continuity	Federated data arrangements distribute data, metadata, evidence, and decision authority, creating a continuing coordination need.	Treat data stewardship as the central function maintaining connections, versions, and routes to review across CEDS participants.	No universal statutory office, organisational model, or demonstrated effectiveness in an operational CEDS.
Integrated trajectory	Two bodies of work and the conceptual framework identify complementary capacity, deliberation, continuity, and representation problems.	Examine, where useful, how selected contributions interact in a concrete public-administration setting.	No necessary Canvas-to-metadata pipeline; their combined use is one possible object of empirical study.

Each contribution is supported by a distinct evidentiary base. The digital-sovereignty analysis diagnoses the relationship between regulatory ambition and material capacity; it supplies no evidence of the operational maturity of CEDS or the quality of the Canvas. The formative Canvas evidence documents artifact development and use in one setting, without establishing the correctness of a legal interpretation or the semantic adequacy of a metadata standard. The conceptual framework specifies relations for later testing, but does not show that a public administration can maintain them at reasonable cost or that their presence will improve a contested decision. Respecting these boundaries is not a weakness in the synthesis. It is the condition under which the contributions can be connected honestly.

The thesis nevertheless establishes more than juxtaposition. The macro-level analysis explains why regulatory ambition requires institutions and infrastructures capable of acting. The Canvas shows one way to structure the organisation's deliberative work once a concrete public-sector project is contemplated. Separately, the conceptual framework specifies how selected legal concepts and requirements may be related to technical information in a source-linked and semantically bounded form. Information collected during deliberation may, where relevant, enrich project documentation or metadata used in subsequent assessment, but the two methods do not form a necessary or automatic pipeline. Each addresses a different failure mode: rights without capacity, technical development without structured multidisciplinary review, and decisions without persistent source-linked representation.

The resulting account is deliberately asymmetric. The three elements need not always appear in the same form or at the same time. A public administration may use another participatory method instead of the Canvas, a data space may adopt different technical components, and a legal assessment may conclude that a proposed use cannot proceed. What must be preserved is the functional relationship: material access and authority should be distinguished; organisational questions should be formulated by people with relevant knowledge and affected perspectives; and machine-readable records should state their sources, evidence, limits, assigned maintenance responsibilities, and review conditions. Data stewardship maintains these connections, while competent actors retain authority over the decisions they support.

6.7 Original Contribution

The thesis contributes a human-centred account of regulatory operationalisation across institutional, organisational, and informational boundaries. European digital policy often presents law, infrastructure, innovation, and public values in a common programme, while implementation distributes them across strategies, authorities, platforms, standards, procurement arrangements, and professional practices. The research reconnects those domains around the concrete journey of data into a public-sector AI project.

Across the three boundaries, the existence of a European rule remains distinct from the

exercise of practical agency; the Canvas provides a setting in which a public administration can expose assumptions and responsibilities before and during design; and the five-stage method relates selected legal language to information needs, candidate elements of technical standards, supporting evidence, and questions reserved for competent judgment. The originality lies in the disciplined connection among these boundaries, not in claiming that digital sovereignty, canvases, metadata vocabularies, or legal requirements engineering were created by this thesis.

This synthesis also yields a general principle: operationalisation should preserve the type of every claim. Policy announcements must be distinguished from outcome evaluations, legal entitlements from material capacity, and access from authority for downstream use. The same discipline applies at the organisational and technical levels: a workshop observation supplies no proof of long-term effect, an assertion remains distinct from its evidence, and structural conformance does not establish factual truth. A metadata match does not amount to legal coverage, and a final review does not amount to certification. Keeping these types separate prevents tools from acquiring authority by rhetorical association.

The contribution is human-centred in a second sense. It recognises that institutional knowledge is distributed and that affected communities may hold knowledge unavailable to the project team. Traceability should therefore serve challenge and reassessment, not only audit and enforcement. Data stewardship gives this requirement an organisational form by maintaining the relations among the governed data, sources, evidence, responsible contributions, decisions, versions, and review states. A source-linked record is useful when it lets a later reviewer ask who made a decision, on which evidence, under which source version, for which dataset and purpose, and what has changed. The stewardship function keeps that record current and routes unresolved or changed matters to the competent actor; it does not decide their legal or institutional significance. Traceability remains insufficient if the organisation has no path for correcting the record or no actor empowered to respond.

Finally, the thesis contributes a reusable Canvas and a conceptual framework whose stages, safeguards, source inventory, and bounded Data Act application can be examined and challenged. The availability of the Canvas supports adaptation without supplying the institutional authority, resources, tacit knowledge, or affected perspectives needed for a concrete decision. Reuse must therefore include local governance and evaluation rather than treating availability as readiness.

6.8 Implications for policy, administration, and standardisation

The findings suggest that European policy should avoid measuring digital capacity solely through the quantity of enacted rules, funded infrastructure, or available datasets. These are important inputs, but a public body experiences capacity through the concrete possibility

of finding suitable data, understanding their conditions, obtaining compute and expertise, evaluating suppliers, and retaining authority over the service it delivers. Monitoring of CEDS should therefore include organisational capability, downstream use, and the presence of stewardship arrangements able to maintain metadata, evidence, versions, and routes for correction, while keeping such indicators separate from the legal question of whether an individual project complies with applicable law.

Public administrations, in turn, need to assign operational responsibility for maintaining the connections through data stewardship, not only responsibility for the component documents. A data-protection assessment, procurement file, security test, catalogue record, model card, permit, and ethics discussion may each be well prepared while referring to different dataset versions or assumptions. The administration needs a proportionate way to show which project decision each artifact supports, who is responsible for maintaining it, and what event makes it stale. The stewardship function may be distributed across existing functions; creating a new title is less important than providing time, resources, escalation routes, and access to the actor competent to decide the unresolved question.

For designers of ethics tools, the study suggests evaluating what happens after a session. A well-facilitated discussion may surface a material concern, but organisational value depends on whether a responsible actor is assigned, supporting evidence is identified, a deadline is set, and a path into design or authorised decision-making exists. Evaluation should include unresolved disagreements and abandoned projects, rather than treating positive participant reactions as the principal outcome. The information and evidence touchpoints identified through the Canvas should therefore be evaluated by observing whether they receive an assigned responsible actor, supporting evidence, and follow-up.

For standards communities and data-space implementers, the framework favours explicit composition and limitation over a claim that one vocabulary covers regulation. A profile can identify how catalogue, policy, provenance, and quality elements work together, require source and evidence links, and reserve fields for uncertainty or human review. Standards structure information and bounded checks; the evidence supporting a concrete assertion must still be produced, maintained, and assessed separately. Conformance tests should state whether they verify syntax, graph shape, controlled values, credentials, or a declared policy condition. They should not describe those bounded checks as validation of legal applicability or substantive adequacy. This vocabulary of limits can improve interoperability by making participant expectations more precise.

6.9 Limitations

The contributions use distinct methods and evidentiary bases. Their relationship is analytical rather than causal and has not yet been examined longitudinally. A future study may show that some links require different tools, responsibilities, or timings.

The digital-sovereignty analysis compares legal and policy approaches and cannot measure all dimensions of industrial capacity or user agency. The update to 2026 relies substantially on official European sources, which are authoritative evidence of institutional policy but interested accounts of progress. Independent empirical measures of cloud dependence, compute access, adoption, and administrative capability would be necessary to evaluate outcomes.

The Canvas evidence is formative. Twelve interviews and one workshop can reveal design needs and practical reactions, but they do not establish generalisability across Member States, sectors, administrative sizes, or project types. The workshop's 18 participants, two project cases, and three scenarios do not substitute for observation of decisions across a full lifecycle. Facilitation effects, organisational hierarchy, selection of participants, time constraints, and local culture may influence what is surfaced or silenced.

The conceptual framework is informed by selected legal and technical materials, authorial interpretation, and particular source versions. It does not establish exhaustive legal or technical coverage or provide independent legal validation. Additional rules, interpretations, factual configurations, or later standards may require the five-stage analysis to be repeated.

The framework was not tested through a live catalogue, data-space transaction, secure processing environment, or production AI system. It therefore provides no direct estimate of integration effort, maintenance cost, staff burden, or impact on project outcomes. These matters cannot be inferred from the conceptual coherence of the five-stage chain.

6.10 Possible Directions for Empirical Evaluation

Future empirical work could examine the contributions in a concrete public-administration project involving data made available through a CEDS. This is one possible research setting, not an operational model prescribed by the thesis. The Canvas, the conceptual framework, and data-stewardship arrangements could be evaluated separately or in different combinations, and an administration may use other deliberative instruments, technical architectures, or allocations of responsibility.

A longitudinal study would be useful where the research question concerns change, maintenance, and transfers between organisational actors. It could observe whether a deliberative instrument affects project framing, whether source-linked records help later reviewers reconstruct the basis of an assessment, and whether stewardship arrangements keep relevant evidence and review responsibilities connected over time. Evaluation should also consider implementation effort, institutional fit, disagreements, missing information, and cases in which a proposed connection proves unnecessary or impractical. Comparing these observations with the administration's existing practices would provide a stronger basis than participant approval alone [192, 193].

Such research would test compatibility and practical value without assuming that the

contributions must form a single sequence. Decisions concerning access, project approval, legal applicability, technical adequacy, or lifecycle transitions would continue to be taken through the competent administration's ordinary procedures. The relevant empirical question is whether one or more of the proposed contributions improve deliberation, traceability, maintenance, or review in a defined setting, and at what organisational cost.

6.11 Human-Centred Compliance in Practice

Europe's data and AI regulation cannot become operational through law, infrastructure, organisational ethics, or metadata in isolation. CEDS depend on public administrations and other participants being able to prepare, describe, govern, and maintain the data they contribute; they can also expand governed discovery, access, and reuse. Neither availability nor access, however, establishes authority for a particular downstream operation. Public administrations can deliberate through the AI Ethics Canvas, but deliberation remains distinct from an authorised and legally adequate decision. Standards can represent descriptions, policies, provenance, and quality, but representation remains distinct from truth and compliance.

Human-centred compliance consists in maintaining these distinctions while building reliable connections among them. It requires material capacity to act, institutional arrangements that give people time and authority to question a project, evidence that can be inspected, technical records that preserve source and limitation, and mechanisms for reopening decisions when circumstances change. Data stewardship supplies the organisational continuity needed to keep legal sources, governed data, supporting evidence, standard-based representations, decisions, and review states connected over time. The thesis brings together comparative legal analysis, formative artifact development, and a conceptual framework illustrated through two bounded Data Act micro-applications. Their relationship within the thesis is conceptual and methodological rather than an implemented workflow.

The final lesson is therefore methodological as well as substantive. Regulatory complexity should not be hidden behind a green status, an ethics label, or a populated metadata field. It should be decomposed carefully, discussed by the people whose knowledge and interests matter, represented where representation is faithful, maintained through data stewardship, and left visibly open where judgment is still required. Standards can structure information and bounded checks, but they do not themselves prove that an assertion is true or that a legal requirement has been satisfied. Organisational and metadata-based tools can properly support regulatory operationalisation by structuring deliberation, preserving traceable information, and making visible the questions that still require competent judgment.

Bibliography

[1] European Parliament and Council of the European Union. *Regulation (EU) 2025/327 on the European Health Data Space*. Official Journal of the European Union. Accessed 2026-07-17. 2025. URL: https://eur-lex.europa.eu/eli/reg/2025/327/oj.

[2] European Parliament and Council of the European Union. *Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence*. Official Journal of the European Union. Accessed 2026-07-17. 2024. URL: https://eur-lex.europa.eu/eli/reg/2024/1689/oj.

[3] European Parliament and Council of the European Union. *Regulation (EU) 2016/679 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data*. Official Journal of the European Union. Accessed 2026-07-17. 2016. URL: https://eur-lex.europa.eu/eli/reg/2016/679/oj.

[4] Simon Scerri et al. "Common European Data Spaces: Challenges and Opportunities". In: *Data Spaces*. Ed. by Edward Curry, Simon Scerri, and Tuomo Tuikka. Springer International Publishing, 2022, pp. 337–357. DOI: 10.1007/978-3-030-98636-0_16.

[5] Manlio Bacco et al. "What Are Data Spaces? Systematic Survey and Future Outlook". In: *Data in Brief* 57 (2024), p. 110969. DOI: 10.1016/j.dib.2024.110969.

[6] European Commission. *A European Strategy for Data*. Tech. rep. COM(2020) 66 final. Accessed 2026-07-17. European Commission, 2020. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52020DC0066.

[7] European Commission. *Common European Data Spaces*. Accessed 2026-07-17. 2026. URL: https://digital-strategy.ec.europa.eu/en/policies/data-spaces.

[8] Huw Roberts et al. "Safeguarding European Values with Digital Sovereignty: An Analysis of Statements and Policies". In: *Internet Policy Review* 10.3 (2021). DOI: 10.14763/2021.3.1575.

[9] Riccardo Nanni, Pietro G. Bizzaro, and Maurizio Napolitano. "The False Promise of Individual Digital Sovereignty in Europe: Comparing Artificial Intelligence and Data Regulations in China and the European Union". In: *Policy & Internet* 16.4 (2024), pp. 711–726. DOI: 10.1002/poi3.424.

[10] Mélanie Dulong de Rosnay and Felix Stalder. "Digital Commons". In: *Internet Policy Review* 9.4 (2020). DOI: 10.14763/2020.4.1530.

[11] Saffron Huang and Divya Siddarth. *Generative AI and the Digital Commons*. 2023. DOI: 10.48550/arXiv.2303.11074. arXiv: 2303.11074 [cs.CY]. URL: https://arxiv.org/abs/2303.11074.

[12] Gordon Burtch, Dokyun Lee, and Zhichen Chen. "The Consequences of Generative AI for Online Knowledge Communities". In: *Scientific Reports* 14.1 (2024), p. 10413. DOI: 10.1038/s41598-024-61221-0.

[13] Shayne Longpre et al. "Consent in Crisis: The Rapid Decline of the AI Data Commons". In: *Advances in Neural Information Processing Systems 37*. NeurIPS 2024 Datasets and Benchmarks Track. 2024. DOI: 10.52202/079017-3431. URL: https://proceedings.neurips.cc/paper_files/paper/2024/hash/c3738949a80306cc48a8ea8ba0560f9d-Abstract-Datasets_and_Benchmarks_Track.html.

[14] Arman Noroozian et al. *Generative AI and the Future of the Digital Commons: Five Open Questions and Knowledge Gaps*. Preprint, version 1; accessed 20 July 2026. Aug. 2025. DOI: 10.48550/arXiv.2508.06470. arXiv: 2508.06470 [cs.CY]. URL: https://arxiv.org/abs/2508.06470.

[15] Ben Green. "The Flaws of Policies Requiring Human Oversight of Government Algorithms". In: *Computer Law & Security Review* 45 (2022), p. 105681. DOI: 10.1016/j.clsr.2022.105681.

[16] Joe Massey et al. *Defining Responsible Data Stewardship*. Tech. rep. Accessed 2026-07-17. Open Data Institute, 2023. DOI: 10.61557/WPQE9992. URL: https://theodi.org/insights/reports/defining-responsible-data-stewardship/.

[17] OECD. *The Path to Becoming a Data-Driven Public Sector*. OECD Digital Government Studies. Paris: OECD Publishing, 2019. DOI: 10.1787/059814a7-en.

[18] UK Government Digital Service. *Data Ownership Model*. Guidance. Published 16 April 2026; accessed 2026-07-20. Apr. 2026. URL: https://www.gov.uk/government/publications/data-ownership-in-government/data-ownership-model.

[19] Adam Faifr and Martin Januška. "Factors Determining the Extent of GDPR Implementation within Organizations: Empirical Evidence from the Czech Republic". In: *Journal of Business Economics and Management* 22.5 (2021), pp. 1124–1141. DOI: 10.3846/jbem.2021.15095.

[20] Ze Shi Li et al. “Towards Privacy Compliance: A Design Science Study in a Small Organization”. In: *Information and Software Technology* 146 (2022), p. 106868. DOI: 10.1016/j.infsof.2022.106868.

[21] Patrik Hummel et al. “Data Sovereignty: A Review”. In: *Big Data & Society* 8.1 (2021). DOI: 10.1177/2053951720982012.

[22] Mark Ryan, Paula Gürtler, and Artur Bogucki. “Will the Real Data Sovereign Please Stand Up? An EU Policy Response to Sovereignty in Data Spaces”. In: *International Journal of Law and Information Technology* 32 (2024). DOI: 10.1093/ijlit/eaae006.

[23] European Commission. *AI Continent Action Plan*. Tech. rep. COM(2025) 165 final. European Commission communication, 9 April 2025; accessed 20 July 2026. European Commission, 2025. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025DC0165.

[24] European Commission. *Data Union Strategy: Unlocking Data for AI*. Tech. rep. COM(2025) 835 final. European Commission communication, 19 November 2025; accessed 20 July 2026. European Commission, 2025. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025DC0835.

[25] Mark D. Wilkinson et al. “The FAIR Guiding Principles for Scientific Data Management and Stewardship”. In: *Scientific Data* 3 (2016), p. 160018. DOI: 10.1038/sdata.2016.18.

[26] European Commission, Interoperable Europe, SEMIC. *DCAT-AP 3.0.1*. SEMIC Recommendation. Version 3.0.1. Accessed 2026-07-17. 2025. URL: https://semiceu.github.io/DCAT-AP/releases/3.0.1/.

[27] World Wide Web Consortium. *ODRL Information Model 2.2*. W3C Recommendation. Accessed 2026-07-17. 2018. URL: https://www.w3.org/TR/odrl-model/.

[28] World Wide Web Consortium. *PROV-O: The PROV Ontology*. W3C Recommendation. Accessed 2026-07-17. 2013. URL: https://www.w3.org/TR/prov-o/.

[29] World Wide Web Consortium. *Data on the Web Best Practices: Data Quality Vocabulary*. W3C Working Group Note. Accessed 2026-07-17. 2016. URL: https://www.w3.org/TR/vocab-dqv/.

[30] Harshvardhan J. Pandit et al. “Data Privacy Vocabulary (DPV)—Version 2.0”. In: *The Semantic Web—ISWC 2024*. Ed. by Gianluca Demartini et al. Vol. 15233. Lecture Notes in Computer Science. Cham: Springer, 2025, pp. 171–193. DOI: 10.1007/978-3-031-77847-6_10.

[31] MLCommons. *Croissant Format Specification 1.1*. Community Specification. Accessed 2026-07-17. 2026. URL: https://docs.mlcommons.org/croissant/docs/croissant-spec-1.1.html.

[32] European Parliament and Council of the European Union. *Regulation (EU) 2023/2854 on Harmonised Rules on Fair Access to and Use of Data*. Official Journal of the European Union. Accessed 2026-07-17. 2023. URL: https://eur-lex.europa.eu/eli/reg/2023/2854/oj.

[33] Riccardo Nanni et al. *AI Ethics Canvas, Version 1.0.0*. Public software and documentation repository. Released under CC BY-SA 4.0; accessed 20 July 2026. 2026. URL: https://github.com/dclfbk/canvasaiethics.

[34] Michael Franklin, Alon Halevy, and David Maier. "From Databases to Dataspaces: A New Abstraction for Information Management". In: *ACM SIGMOD Record* 34.4 (2005), pp. 27–33. DOI: 10.1145/1107499.1107502.

[35] Alon Y. Halevy, Michael J. Franklin, and David Maier. "Principles of Dataspace Systems". In: *Proceedings of the Twenty-Fifth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. Association for Computing Machinery, 2006, pp. 1–9. DOI: 10.1145/1142351.1142352.

[36] Doris Marquardt. *Blueprint Proposal for the Common European Agricultural Data Space*. Updated 27 August 2026; account of the completed AgriDataSpace preparatory action. European Commission. May 2, 2024. URL: https://digital-strategy.ec.europa.eu/en/library/blueprint-proposal-common-european-agricultural-data-space (visited on 10/03/2026).

[37] *DeployEMDS, the Common European Mobility Data Space Starts to Become Reality*. Digital Europe deployment action; programme objectives and local implementation projects. European Commission. Mar. 19, 2024. URL: https://digital-strategy.ec.europa.eu/en/news/deployemds-common-european-mobility-data-space-starts-become-reality (visited on 10/03/2026).

[38] Fraunhofer IOSB. *SM4RTENANCE – Trusted Data Sharing for Manufacturing Equipment Industry*. Project-participant account of objectives and integration work; no publication date stated. URL: https://www.iosb.fraunhofer.de/en/projects-and-products/sm4rtenance.html (visited on 10/03/2026).

[39] Lise S. Svingel et al. "Shaping the Future EHDS: Recommendations for Implementation of Health Data Access Bodies in the HealthData@EU Infrastructure for Secondary Use of Electronic Health Data". In: *European Journal of Public Health* 35.Supplement 3 (2025), pp. iii32–iii38. DOI: 10.1093/eurpub/ckaf033.

[40] Europeana Foundation. *Common European Data Space for Cultural Heritage – Annual Report 2024–2025*. Implementation-progress summary accompanying the annual report; updated 13 April 2026. Nov. 17, 2025. URL: https://pro.europeana.eu/post/common-european-data-space-for-cultural-heritage-annual-report-2024-2025 (visited on 10/03/2026).

[41] European Commission. *The Public Procurement Data Space (PPDS)*. Architecture and implementation overview; no publication date stated. URL: https://single-market-economy.ec.europa.eu/single-market/public-procurement/digital-procurement/public-procurement-data-space-ppds_en (visited on 10/03/2026).

[42] *Periodic Reporting for Period 2 – ENERSHARE*. CORDIS project-reported summary; grant agreement 101069831; reporting period: 2024-01-01 to 2025-06-30. Date denotes page update. European Commission. Dec. 8, 2025. URL: https://cordis.europa.eu/project/id/101069831/reporting (visited on 10/03/2026).

[43] *Periodic Reporting for Period 2 – OMEGA-X*. CORDIS project-reported summary; grant agreement 101069287; reporting period: 2023-11-01 to 2025-04-30. Date denotes page update. European Commission. Oct. 27, 2025. URL: https://cordis.europa.eu/project/id/101069287/reporting (visited on 10/03/2026).

[44] *Periodic Reporting for Period 2 – DATA CELLAR*. CORDIS project-reported summary; grant agreement 101069694; reporting period: 2023-06-01 to 2024-11-30. Date denotes page update. European Commission. May 23, 2025. URL: https://cordis.europa.eu/project/id/101069694/reporting (visited on 10/03/2026).

[45] *Periodic Reporting for Period 1 – EDDIE*. CORDIS project-reported summary; grant agreement 101069510; reporting period: 2023-01-01 to 2024-06-30. Date denotes page update. European Commission. Dec. 18, 2024. URL: https://cordis.europa.eu/project/id/101069510/reporting (visited on 10/03/2026).

[46] *Periodic Reporting for Period 2 – SYNERGIES*. CORDIS project-reported summary; grant agreement 101069839; reporting period: 2023-09-01 to 2025-02-28. Date denotes page update. European Commission. Oct. 30, 2025. URL: https://cordis.europa.eu/project/id/101069839/reporting (visited on 10/03/2026).

[47] *Establish the Grounds for a Common European Energy Data Space*. CORDIS topic record; HORIZON-CL5-2021-D3-01-01. Date denotes page update. European Commission. June 21, 2022. URL: https://cordis.europa.eu/programme/id/HORIZON_HORIZON-CL5-2021-D3-01-01 (visited on 10/03/2026).

[48] INSIEME Project. *The INSIEME Vision: Towards a Common European Data Space for Energy*. Vision paper. Digital Europe Programme, grant agreement 101194952. INSIEME, Oct. 2025. URL: https://insieme.energy/library/vision/20251105_insieme_visionpaper.pdf (visited on 10/03/2026).

[49] International Data Spaces Association. *Processing Services*. IDSA Rulebook, online chapter; no publication date stated. URL: https://kb.internationaldataspaces.org/external/rulebook/124_ProcessingServices/ (visited on 10/03/2026).

[50] Dataspace Protocol Working Group. *Dataspace Protocol 2025-1*. Technical Specification. Accessed 2026-07-17. 2025. URL: https://eclipse-dataspace-protocol-base.github.io/DataspaceProtocol/2025-1/.

[51] Wout Von Voornveld et al. *Data Exchange Reference Architecture (DERA): Updated Report 2025*. BRIDGE Data Management Working Group report. Publications Office of the European Union, Oct. 2025. DOI: 10.2926/4460045. URL: https://bridge-smart-grid-storage-systems-digital-projects.ec.europa.eu/sites/default/files/bridge-reports/data%20exchange%20reference%20architecture%20dera-HZ0125151ENN.pdf (visited on 10/03/2026).

[52] European High-Performance Computing Joint Undertaking. *The EuroHPC Joint Undertaking Launches AI Factories Calls to Boost European Leadership in Trustworthy AI*. Press release. Sept. 10, 2024. URL: https://www.eurohpc-ju.europa.eu/eurohpc-joint-undertaking-launches-ai-factories-calls-boost-european-leadership-trustworthy-ai-2024-09-10_en (visited on 10/03/2026).

[53] Marcello Ienca et al. *How the General Data Protection Regulation Changes the Rules for Scientific Research*. Tech. rep. PE 634.447. ISBN 978-92-846-5045-3. Brussels: European Parliamentary Research Service, Scientific Foresight Unit (STOA), July 2019. DOI: 10.2861/17421. URL: https://www.europarl.europa.eu/thinktank/en/document/EPRS_STU(2019)634447.

[54] European Data Protection Board. *Opinion 28/2024 on Certain Data Protection Aspects Related to the Processing of Personal Data in the Context of AI Models*. Tech. rep. 28/2024. Accessed 2026-07-17. European Data Protection Board, 2024. URL: https://www.edpb.europa.eu/system/files/documents/2024-12/edpb_opinion_202428_ai-models_en.pdf.

[55] Pablo Trigo Kramcsák. "Can Legitimate Interest Be an Appropriate Lawful Basis for Processing Artificial Intelligence Training Datasets?" In: *Computer Law & Security Review* 48 (2023), p. 105765. DOI: 10.1016/j.clsr.2022.105765.

[56] European Parliament and Council of the European Union. *Regulation (EU) 2022/868 on European Data Governance*. Official Journal of the European Union. Accessed 2026-07-17. 2022. URL: https://eur-lex.europa.eu/eli/reg/2022/868/oj.

[57] European Parliament and Council of the European Union. *Directive (EU) 2019/790 on Copyright and Related Rights in the Digital Single Market*. Official Journal of the European Union. Accessed 2026-07-17. 2019. URL: https://eur-lex.europa.eu/eli/dir/2019/790/oj.

[58] João Pedro Quintais. "What Is a 'Research Organisation' and Why It Matters: From Text and Data Mining to AI Research". In: *GRUR International* 74.5 (2025), pp. 397–398. DOI: 10.1093/grurint/ikaf030.

[59] Thomas Margoni and Martin Kretschmer. "A Deeper Look into the EU Text and Data Mining Exceptions: Harmonisation, Data Ownership, and the Future of Technology". In: *GRUR International* 71.8 (2022), pp. 685–701. DOI: 10.1093/grurint/ikac054.

[60] Pietro Giovanni Bizzaro and Maurizio Napolitano. "Balancing Innovation and Regulation: A Critical Look at the AI Act's Research Exemptions". In: *BioLaw Journal – Rivista di BioDiritto* 3 (Oct. 2025), pp. 289–292. DOI: 10.15168/2284-4503-3741. URL: https://teseo.unitn.it/biolaw/article/view/3741.

[61] Michèle Finck. "In Search of the Lost Research Exemption: Reflections on the AI Act". In: *GRUR International* 74.10 (2025), pp. 903–904. DOI: 10.1093/grurint/ikaf100.

[62] Philipp Hacker. "A Legal Framework for AI Training Data—From First Principles to the Artificial Intelligence Act". In: *Law, Innovation and Technology* 13.2 (2021), pp. 257–301. DOI: 10.1080/17579961.2021.1977219.

[63] Hajo Michael Holtz and Jonas Ledendal. "AI Data Governance—Overlaps between the AI Act and the GDPR". In: *Law, Innovation and Technology* 18.1 (2026), pp. 380–409. DOI: 10.1080/17579961.2026.2633677.

[64] European Commission. *Template for the Public Summary of Training Content of General-Purpose AI Models*. Page updated 26 March 2026; accessed 17 July 2026. 2025. URL: https://digital-strategy.ec.europa.eu/en/library/explanatory-notice-and-template-public-summary-training-content-general-purpose-ai-models.

[65] Michael A. Madaio et al. "Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI". In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, 2020, pp. 1–14. DOI: 10.1145/3313831.3376445.

[66] Bogdana Rakova et al. "Where Responsible AI Meets Reality: Practitioner Perspectives on Enablers for Shifting Organizational Practices". In: *Proceedings of the ACM on Human-Computer Interaction* 5.CSCW1 (2021), pp. 1–23. DOI: 10.1145/3449081.

[67] Jennifer Cobbe, Michelle Seng Ah Lee, and Jatinder Singh. "Reviewable Automated Decision-Making: A Framework for Accountable Algorithmic Systems". In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, 2021, pp. 598–609. DOI: 10.1145/3442188.3445921.

[68] Kieron O'Hara. *Data Trusts: Ethics, Architecture and Governance for Trustworthy Data Stewardship*. Tech. rep. WSI White Paper 1. Web Science Institute, University of Southampton, 2019. DOI: 10.5258/SOTON/WSI-WP001. URL: https://eprints.soton.ac.uk/428276.

[69] World Economic Forum. *Empowered Data Societies: A Human-Centric Approach to Data Relationships*. Tech. rep. Prepared through collaboration with the Government of Finland and the City of Helsinki; accessed 17 July 2026. World Economic Forum, Sept. 2021. URL: https://www.weforum.org/publications/empowered-data-societies-a-human-centric-approach-to-data-relationships/.

[70] High-Level Expert Group on Business-to-Government Data Sharing. *Towards a European Strategy on Business-to-Government Data Sharing for the Public Interest: Final Report*. Tech. rep. Independent expert group set up by the European Commission; accessed 2026-07-20. Luxembourg: Publications Office of the European Union, 2020. DOI: 10.2759/731415. URL: https://op.europa.eu/en/publication-detail/-/publication/d96edc29-70fd-11eb-9ac9-01aa75ed71a1.

[71] Rene Abraham, Johannes Schneider, and Jan vom Brocke. "Data Governance: A Conceptual Framework, Structured Review, and Research Agenda". In: *International Journal of Information Management* 49 (2019), pp. 424–438. DOI: 10.1016/j.ijinfomgt.2019.07.008.

[72] Barend Mons. *Data Stewardship for Open Science: Implementing FAIR Principles*. CRC Press, 2018. DOI: 10.1201/9781315380711.

[73] Anna Wałek. "Data Librarian and Data Steward: New Tasks and Responsibilities of Academic Libraries in the Context of Open Research Data Implementation in Poland". In: *Przeglad Biblioteczny* 4 (2019), pp. 497–512. DOI: 10.36702/pb.634.

[74] Travis D. Breaux and Thomas Norton. "Legal Accountability as Software Quality: A U.S. Data Processing Perspective". In: *2022 IEEE 30th International Requirements Engineering Conference*. IEEE, 2022, pp. 101–113. DOI: 10.1109/RE54965.2022.00016.

[75] Claudia Negri-Ribalta, Marius Lombard-Platet, and Camille Salinesi. "Understanding the GDPR from a Requirements Engineering Perspective—A Systematic Mapping Study on Regulatory Data Protection Requirements". In: *Requirements Engineering* 29.4 (2024), pp. 523–549. DOI: 10.1007/s00766-024-00423-4.

[76] Oleksandr Kosenkov et al. "Systematic Mapping Study on Requirements Engineering for Regulatory Compliance of Software Systems". In: *Information and Software Technology* 178 (2025), p. 107622. DOI: 10.1016/j.infsof.2024.107622.

[77] Petros Terzis and O. E. Enrique Santamaria Echeverria. "Interoperability and Governance in the European Health Data Space Regulation". In: *Medical Law International* 23.4 (2023), pp. 368–376. DOI: 10.1177/09685332231165692.

[78] Emily M. Bender and Batya Friedman. "Data Statements for Natural Language Processing: Toward Mitigating System Bias and Enabling Better Science". In: *Transactions of the Association for Computational Linguistics* 6 (2018), pp. 587–604. DOI: 10.1162/tacl_a_00041.

[79] Timnit Gebru et al. "Datasheets for Datasets". In: *Communications of the ACM* 64.12 (2021), pp. 86–92. DOI: 10.1145/3458723.

[80] Margaret Mitchell et al. "Model Cards for Model Reporting". In: *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, 2019, pp. 220–229. DOI: 10.1145/3287560.3287596.

[81] Mahima Pushkarna, Andrew Zaldivar, and Oddur Kjartansson. "Data Cards: Purposeful and Transparent Dataset Documentation for Responsible AI". In: *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, 2022, pp. 1776–1826. DOI: 10.1145/3531146.3533231.

[82] Albana Celepija et al. "AI Product Cards: A Framework for Code-Bound Formal Documentation Cards in the Public Administration". In: *Data & Policy* 7 (2025), e1. DOI: 10.1017/dap.2024.55.

[83] Ben Hutchinson et al. "Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure". In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, 2021, pp. 560–575. DOI: 10.1145/3442188.3445918.

[84] Marina Micheli et al. "The Landscape of Data and AI Documentation Approaches in the European Policy Context". In: *Ethics and Information Technology* 25.4 (2023). DOI: 10.1007/s10676-023-09725-7.

[85] Xinyu Yang, Weixin Liang, and James Zou. *Navigating Dataset Documentations in AI: A Large-Scale Analysis of Dataset Cards on Hugging Face*. Accepted at ICLR 2024; accessed 17 July 2026. 2024. eprint: 2401.13822. URL: https://arxiv.org/abs/2401.13822.

[86] Weixin Liang et al. *What's Documented in AI? Systematic Analysis of 32K AI Model Cards*. Accessed 2026-07-17. 2024. eprint: 2402.05160. URL: https://arxiv.org/abs/2402.05160.

[87] Shayne Longpre, Robert Mahari, Anthony Chen, et al. "A Large-Scale Audit of Dataset Licensing and Attribution in AI". In: *Nature Machine Intelligence* 6.8 (2024), pp. 975–987. DOI: 10.1038/s42256-024-00878-8.

[88] Jack Bandy and Nicholas Vincent. “Addressing “Documentation Debt” in Machine Learning: A Retrospective Datasheet for BookCorpus”. In: *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*. Accessed 2026-07-17. 2021. URL: https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/9a1158154dfa42caddbd0694a4e9bdc8-Abstract-round1.html.

[89] World Wide Web Consortium. *Data Catalog Vocabulary (DCAT)—Version 3*. W3C Recommendation. Accessed 2026-07-17. 2024. URL: https://www.w3.org/TR/vocab-dcat-3/.

[90] Nishad Thalhath, Mitsuharu Nagamori, and Tetsuo Sakaguchi. “Metadata Application Profile as a Mechanism for Semantic Interoperability in FAIR and Open Data Publishing”. In: *Data and Information Management* 9.1 (2025), p. 100068. DOI: 10.1016/j.dim.2024.100068.

[91] Mingfang Wu et al. “An Analysis of Crosswalks from Research Data Schemas to Schema.org”. In: *Data Intelligence* 5.1 (2023), pp. 100–121. DOI: 10.1162/dint_a_00186.

[92] Erhard Rahm and Philip A. Bernstein. “A Survey of Approaches to Automatic Schema Matching”. In: *The VLDB Journal* 10.4 (2001), pp. 334–350. DOI: 10.1007/s007780100057.

[93] Inès Osman, Salvatore Flavio Pileggi, and Sadok Ben Yahia. “Uncertainty in Automated Ontology Matching: Lessons from an Empirical Evaluation”. In: *Applied Sciences* 14.11 (2024), p. 4679. DOI: 10.3390/app14114679.

[94] European Commission, Interoperable Europe, SEMIC. *MLDCAT-AP 3.1.0*. SEMIC Candidate Recommendation. Version 3.1.0. Accessed 2026-07-17. 2026. URL: https://semiceu.github.io/MLDCAT-AP/releases/3.1.0/.

[95] European Commission, Directorate-General for Health and Food Safety. *HealthDCAT-AP Release 7*. European Commission Draft Specification. Version Release 7. Accessed 2026-07-17. 2026. URL: https://healthdataeu.pages.code.europa.eu/healthdcat-ap/releases/release-7/.

[96] German Federal Institute for Drugs and Medical Devices (BfArM). *Data Access and Coordination Office: Metadata Catalogue*. Official implementation information. Accessed 2026-07-19. 2026. URL: https://www.bfarm.de/EN/BfArM/Tasks/Data-Access-and-Coordination-Office/Metadata-catalog/_node.html.

[97] Sebastian Neumaier, Tobias Dam, and Fabian Kovac. “AIDOC-AP: AI Documentation Application Profile”. Ontology release accompanying an ontology-description manuscript under major revision after open peer review; accessed 2026-07-19. 2026. DOI: 10.5281/zenodo.17787787. URL: https://www.semantic-web-journal.

net/content/aidoc-ap-application-profile-technical-documentation-ai-systems.

[98] ELI Task Force and Publications Office of the European Union. *European Legislation Identifier (ELI) Metadata Ontology, Core Release 1.3*. Official Ontology Release. Version 1.3. Accessed 2026-07-18. 2020. URL: https://interoperable-europe.ec.europa.eu/collection/eli-european-legislation-identifier/solution/eli-ontology/releases.

[99] European Parliament. *ELI-EP Application Profile 1.5*. European Parliament Application Profile. Version 1.5. Profile page records last update 5 December 2023; accessed 17 July 2026. 2022. URL: https://europarl.github.io/eli-ep/1.5/.

[100] Publications Office of the European Union. *European Legislation Identifier (ELI) Metadata Ontology, Core Distribution 1.5*. Official Machine-Readable Ontology Artifact. Version 1.5. Official distribution dated 21 March 2024; distinct from ELI-EP 1.5; accessed 18 July 2026. 2024. URL: https://publications.europa.eu/resource/distribution/eli/20240321-0/owl/owl/eli.owl.

[101] W3C Data Privacy Vocabularies and Controls Community Group. *Data Privacy Vocabulary (DPV) Version 2.3*. W3C Community Group Final Report. Accessed 2026-07-17. World Wide Web Consortium, 2026. URL: https://w3id.org/dpv/2.3.

[102] Jonathan Lawson et al. "The Data Use Ontology to Streamline Responsible Access to Human Biomedical Datasets". In: *Cell Genomics* 1.2 (2021), p. 100028. DOI: 10.1016/j.xgen.2021.100028.

[103] World Wide Web Consortium. *PROV-AQ: Provenance Access and Query*. W3C Working Group Note. Accessed 2026-07-17. 2013. URL: https://www.w3.org/TR/prov-aq/.

[104] OASIS Open. *LegalRuleML Core Specification Version 1.0*. OASIS Standard. Accessed 2026-07-17. 2021. URL: https://docs.oasis-open.org/legalruleml/legalruleml-core-spec/v1.0/os/legalruleml-core-spec-v1.0-os.html.

[105] Piero A. Bonatti et al. "Machine Understandable Policies and GDPR Compliance Checking". In: *KI – Künstliche Intelligenz* 34.3 (2020), pp. 303–315. DOI: 10.1007/s13218-020-00677-4.

[106] Joris Hulstijn. "Evidence-Based Compliance Engineering". In: *Proceedings of the 17th International Workshop on Value Modelling and Business Ontologies*. Vol. 3821. CEUR Workshop Proceedings. Accessed 2026-07-19. 2024, pp. 1–12. URL: https://ceur-ws.org/Vol-3821/paper5.pdf.

[107] Julio Hernandez, Delaram Golpayegani, and Dave Lewis. "An Open Knowledge Graph-Based Approach for Mapping Concepts and Requirements between the EU AI Act and International Standards". In: *AI and Ethics* 5 (2025), pp. 4463–4474. DOI: 10.1007/s43681-025-00708-6.

[108] Livinus Obiora Nweke and Prosper Kandabongee Yeng. "Compliance-as-Code for AI-Driven Identity Systems: Clause-to-Control Traceability and Machine-Readable Evidence". In: *IEEE Access* 14 (2026), pp. 28258–28281. DOI: 10.1109/ACCESS.2026.3665991.

[109] Giedrė Sabaliauskaitė et al. "Automated Regulatory Compliance for AI Systems in the Security Domain: The Case of Dual-Use Deployment". In: *Open Research Europe* 6 (2026). Version 1, p. 83. DOI: 10.12688/openreseurope.23137.1.

[110] Travis D. Breaux and Annie I. Antón. "Analyzing Regulatory Rules for Privacy and Security Requirements". In: *IEEE Transactions on Software Engineering* 34.1 (2008), pp. 5–20. DOI: 10.1109/TSE.2007.70746.

[111] Andrea Zasada et al. "Evaluation of Compliance Rule Languages for Modelling Regulatory Compliance Requirements". In: *Software* 2.1 (2023), pp. 71–120. DOI: 10.3390/software2010004.

[112] Romina Etezadi et al. "Classifier or Prompt: A Case Study on Legal Requirements Traceability". In: *Empirical Software Engineering* (2026). Accessed 2026-07-17. DOI: 10.1007/s10664-026-10827-1. eprint: 2502.04916. URL: https://arxiv.org/abs/2502.04916.

[113] Julia Mucha, Andreas Kaufmann, and Dirk Riehle. "A Systematic Literature Review of Pre-Requirements Specification Traceability". In: *Requirements Engineering* 29.2 (2024), pp. 119–141. DOI: 10.1007/s00766-023-00412-z.

[114] Marcela Ruiz, Jin Yang Hu, and Fabiano Dalpiaz. "Why Don't We Trace? A Study on the Barriers to Software Traceability in Practice". In: *Requirements Engineering* 28.4 (2023), pp. 619–637. DOI: 10.1007/s00766-023-00408-9.

[115] Patrick Mäder and Alexander Egyed. "Do Developers Benefit from Requirements Traceability When Evolving and Maintaining a Software System?" In: *Empirical Software Engineering* 20.2 (2015), pp. 413–441. DOI: 10.1007/s10664-014-9314-z.

[116] Data Spaces Support Centre. *Data Spaces Blueprint*. Reference Blueprint. Version 3.0. Accessed 2026-07-17. 2026. URL: https://blueprint.dssc.eu/.

[117] International Data Spaces Association. *IDS Reference Architecture Model*. Reference Architecture. Version 4. Accessed 2026-07-17. 2022. URL: https://docs.internationaldataspaces.org/ids-knowledgebase/ids-ram-4/.

[118] Jovan Powar et al. “From Policy to Practice in Data Governance and Responsible Data Stewardship: System Design for Data Intermediaries”. In: *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. New York, NY, USA: Association for Computing Machinery, 2025, pp. 2450–2463. DOI: 10.1145/3715275.3732164.

[119] Gaia-X European Association for Data and Cloud AISBL. *Gaia-X Compliance Document*. Tech. rep. Version 2.5.0. Accessed 2026-07-17. Gaia-X European Association for Data and Cloud AISBL, 2026. URL: https://docs.gaia-x.eu/policy-rules-committee/compliance-document/latest/Introduction_and_scope/.

[120] European Telecommunications Standards Institute. *Data Solutions (DATA); Data Act (art. 33) Requirement and References Analysis*. Technical Report ETSI TR 104 409 V1.1.1. Accessed 2026-07-17. ETSI, 2025. URL: https://www.etsi.org/deliver/etsi_TR/104400_104499/104409/01.01.01_60/tr_104409v010101p.pdf.

[121] European Telecommunications Standards Institute. *Data Solutions (DATA); Data Act (art. 33) Standardization Suggestions*. Technical Report ETSI TR 104 410 V1.1.1. Accessed 2026-07-17. ETSI, 2025. URL: https://www.etsi.org/deliver/etsi_TR/104400_104499/104410/01.01.01_60/tr_104410v010101p.pdf.

[122] European Parliament and Council of the European Union. *Regulation (EU) 2024/903 Laying Down Measures for a High Level of Public Sector Interoperability Across the Union*. Official Journal of the European Union. Accessed 2026-07-18. 2024. URL: https://eur-lex.europa.eu/eli/reg/2024/903/oj.

[123] Florian Kelbert and Alexander Pretschner. “Data Usage Control for Distributed Systems”. In: *ACM Transactions on Privacy and Security* 21.3 (2018), pp. 1–32. DOI: 10.1145/3183342.

[124] Malte Hellmeier et al. “Implementing Data Sovereignty: Requirements and Challenges from Practice”. In: *Proceedings of the 18th International Conference on Availability, Reliability and Security*. ACM, 2023, pp. 1–9. DOI: 10.1145/3600160.3604995.

[125] Sebastian Copei and Linda Rülicke. “Best Practices to Overcome Challenges and Barriers during the Implementation of a Data Space for the Energy Domain: An Experience Report”. In: *Data in Brief* 61 (2025), p. 111838. DOI: 10.1016/j.dib.2025.111838.

[126] Tobias Dam et al. “A Survey of Dataspace Connector Implementations”. In: *Proceedings of ITADATA 2023*. Accessed 2026-07-17. 2023. eprint: 2309.11282. URL: https://ceur-ws.org/Vol-3606/paper41.pdf.

[127] OneTrust. *Data Use Governance*. Product documentation. Vendor-authored description; accessed 2026-07-19. 2026. URL: https://www.onetrust.com/solutions/data-use-governance/.

[128] Jessica Morley et al. "From What to How: An Initial Review of Publicly Available AI Ethics Tools, Methods and Research to Translate Principles into Practices". In: *Science and Engineering Ethics* 26.4 (2020), pp. 2141–2168. DOI: 10.1007/s11948-019-00165-5.

[129] Brent Mittelstadt. "Principles Alone Cannot Guarantee Ethical AI". In: *Nature Machine Intelligence* 1.11 (2019), pp. 501–507. DOI: 10.1038/s42256-019-0114-4.

[130] Andrew D. Selbst et al. "Fairness and Abstraction in Sociotechnical Systems". In: *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, 2019, pp. 59–68. DOI: 10.1145/3287560.3287598.

[131] Kenneth Holstein et al. "Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need?" In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 2019, pp. 1–16. DOI: 10.1145/3290605.3300830.

[132] Sara Kijewski, Elettra Ronchi, and Effy Vayena. "The Rise of Checkbox AI Ethics: A Review". In: *AI and Ethics* 5.3 (2025), pp. 1931–1940. DOI: 10.1007/s43681-024-00563-x.

[133] Riccardo Nanni et al. "AI Ethics Canvas: A Graphical Tool to Design and Deploy Ethical Artificial Intelligence for Public Administrations". In: *Data & Policy* (2026). Forthcoming; accepted and peer reviewed.

[134] Saleema Amershi et al. "Guidelines for Human–AI Interaction". In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 2019, pp. 1–13. DOI: 10.1145/3290605.3300233.

[135] Upol Ehsan et al. "Expanding Explainability: Towards Social Transparency in AI Systems". In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, 2021, pp. 1–19. DOI: 10.1145/3411764.3445188.

[136] Gagan Bansal et al. "Does the Whole Exceed Its Parts? The Effect of AI Explanations on Complementary Team Performance". In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, 2021, pp. 1–16. DOI: 10.1145/3411764.3445717.

[137] Harmanpreet Kaur et al. "Interpreting Interpretability: Understanding Data Scientists' Use of Interpretability Tools for Machine Learning". In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, 2020, pp. 1–14. DOI: 10.1145/3313831.3376219.

[138] Zana Buçinca, Maja Barbara Malaya, and Krzysztof Z. Gajos. "To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-Assisted Decision-Making". In: *Proceedings of the ACM on Human-Computer Interaction* 5.CSCW1 (2021), pp. 1–21. DOI: 10.1145/3449287.

[139] Michelle Vaccaro, Abdullah Almaatouq, and Thomas Malone. "When Combinations of Humans and AI Are Useful: A Systematic Review and Meta-Analysis". In: *Nature Human Behaviour* 8.12 (2024), pp. 2293–2303. DOI: 10.1038/s41562-024-02024-1.

[140] Kate Goddard, Abdul Roudsari, and Jeremy C. Wyatt. "Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators". In: *Journal of the American Medical Informatics Association* 19.1 (2012). Accessed 2026-07-17, pp. 121–127. DOI: 10.1136/amiajnl-2011-000089. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC3240751/.

[141] Daniela Sele and Marina Chugunova. "Putting a Human in the Loop: Increasing Uptake, but Decreasing Accuracy of Automated Decision-Making". In: *PLOS ONE* 19.2 (2024), e0298037. DOI: 10.1371/journal.pone.0298037.

[142] Alexia Gaudeul et al. "Understanding the Impact of Human Oversight on Discriminatory Outcomes in AI-Supported Decision-Making". In: *ECAI 2024*. Ed. by Ulle Endriss et al. Vol. 392. Frontiers in Artificial Intelligence and Applications. IOS Press, 2024, pp. 1067–1074. DOI: 10.3233/FAIA240598.

[143] Susanne Gaube et al. "Do as AI Say: Susceptibility in Deployment of Clinical Decision-Aids". In: *npj Digital Medicine* 4.1 (2021). DOI: 10.1038/s41746-021-00385-9.

[144] Saar Alon-Barkat and Madalina Busuioc. "Human–AI Interactions in Public Sector Decision Making: "Automation Bias" and "Selective Adherence" to Algorithmic Advice". In: *Journal of Public Administration Research and Theory* 33.1 (2023), pp. 153–169. DOI: 10.1093/jopart/muac007.

[145] Maria De-Arteaga, Riccardo Fogliato, and Alexandra Chouldechova. "A Case for Humans-in-the-Loop: Decisions in the Presence of Erroneous Algorithmic Scores". In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, 2020, pp. 1–12. DOI: 10.1145/3313831.3376638.

[146] Dan Hendrycks et al. "CUAD: An Expert-Annotated NLP Dataset for Legal Contract Review". In: *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*. Vol. 1. NeurIPS, 2021. URL: https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/hash/6ea9ab1baa0efb9e19094440c317e21b-Abstract-round1.html.

[147] Pietro Giovanni Bizzaro et al. "Annotation and Classification of Relevant Clauses in Terms-and-Conditions Contracts". In: *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. Ed. by Nicoletta Calzolari et al. Accessed 2026-07-17. Torino, Italy: ELRA and ICCL, May 2024, pp. 1209–1214. URL: https://aclanthology.org/2024.lrec-main.108/.

[148] Marco Lippi et al. "CLAUDETTE: An Automated Detector of Potentially Unfair Clauses in Online Terms of Service". In: *Artificial Intelligence and Law* 27.2 (2019), pp. 117–139. DOI: 10.1007/s10506-019-09243-2.

[149] Maria Assunta Cappelli et al. "Evaluating GenAI for Automated EU AI Act Compliance against Human Experts". In: *Discover Artificial Intelligence* 6.1 (2026). DOI: 10.1007/s44163-026-01196-1.

[150] Gijs van Dijck, Carlos Aguilera, and Shashank M. Chakravarthy. "Deciphering Disagreement in the Annotation of EU Legislation". In: *Artificial Intelligence and Law* 34.1 (2024), pp. 191–226. DOI: 10.1007/s10506-024-09423-9.

[151] Joe Collenette, Katie Atkinson, and Trevor Bench-Capon. "Explainable AI Tools for Legal Reasoning about Cases: A Study on the European Court of Human Rights". In: *Artificial Intelligence* 317 (2023), p. 103861. DOI: 10.1016/j.artint.2023.103861.

[152] Eclipse Dataspace Components. *Developer Handbook*. Accessed 2026-07-17. 2026. URL: https://github.com/eclipse-edc/docs/blob/main/developer/handbook.md.

[153] Julia Pohle and Thorsten Thiel. "Digital Sovereignty". In: *Internet Policy Review* 9.4 (2020), pp. 1–19. DOI: 10.14763/2020.4.1532.

[154] Andrea Calderaro and Stella Blumfelde. "Artificial Intelligence and EU Security: The False Promise of Digital Sovereignty". In: *European Security* 31.3 (2022), pp. 415–434. DOI: 10.1080/09662839.2022.2101885.

[155] Standing Committee of the National People's Congress of the People's Republic of China. *Personal Information Protection Law of the People's Republic of China*. English translation by DigiChina, Stanford University. Passed 20 August 2021; effective 1 November 2021; accessed 2026-07-20. 2021. URL: https://digichina.stanford.edu/work/translation-personal-information-protection-law-of-the-peoples-republic-of-china-effective-nov-1-2021/.

[156] Standing Committee of the National People's Congress of the People's Republic of China. *Data Security Law of the People's Republic of China*. English translation by DigiChina, Stanford University. Passed 10 June 2021; effective 1 September 2021; accessed 2026-07-20. 2021. URL: https://digichina.stanford.edu/work/translation-data-security-law-of-the-peoples-republic-of-china/.

[157] Cyberspace Administration of China, Ministry of Industry and Information Technology, Ministry of Public Security, and State Administration for Market Regulation. *Internet Information Service Algorithmic Recommendation Management Provisions*. English translation by DigiChina, Stanford University. Issued 31 December 2021; effective 1 March 2022; accessed 2026-07-20. 2021. URL: https://digichina.stanford.edu/work/translation-internet-information-service-algorithmic-recommendation-management-provisions-effective-march-1-2022/.

[158] Standing Committee of the National People's Congress of the People's Republic of China. *Cybersecurity Law of the People's Republic of China*. English translation by DigiChina, Stanford University. Passed 7 November 2016; effective 1 June 2017; accessed 2026-07-20. 2016. URL: https://digichina.stanford.edu/work/translation-cybersecurity-law-of-the-peoples-republic-of-china-effective-june-1-2017/.

[159] State Council of the People's Republic of China. *Critical Information Infrastructure Security Protection Regulations*. English translation by DigiChina, Stanford University. State Council Order No. 745; effective 1 September 2021; accessed 2026-07-20. 2021. URL: https://digichina.stanford.edu/work/translation-critical-information-infrastructure-security-protection-regulations-effective-sept-1-2021/.

[160] Cyberspace Administration of China. *Outbound Data Transfer Security Assessment Measures*. English translation by DigiChina, Stanford University. Issued 7 July 2022; effective 1 September 2022; accessed 2026-07-20. 2022. URL: https://digichina.stanford.edu/work/translation-outbound-data-transfer-security-assessment-measures-effective-sept-1-2022/.

[161] Cyberspace Administration of China, Ministry of Industry and Information Technology, and Ministry of Public Security. *Provisions on the Administration of Deep Synthesis Internet Information Services*. English translation by China Law Translate. Issued 25 November 2022; effective 10 January 2023; accessed 2026-07-20. 2022. URL: https://www.chinalawtranslate.com/en/deep-synthesis/.

[162] Cyberspace Administration of China and six other authorities. *Interim Measures for the Management of Generative Artificial Intelligence Services*. English translation by China Law Translate. Issued 10 July 2023; effective 15 August 2023; accessed 2026-07-20. 2023. URL: https://www.chinalawtranslate.com/en/generative-ai-interim/.

[163] Rogier Creemers. "China's Emerging Data Protection Framework". In: *Journal of Cybersecurity* 8.1 (2022), tyac011. DOI: 10.1093/cybsec/tyac011.

[164] European Parliament and Council of the European Union. *Directive (EU) 2022/2555 on Measures for a High Common Level of Cybersecurity Across the Union*. Official Journal of the European Union. Accessed 2026-07-18. 2022. URL: https://eur-lex.europa.eu/eli/dir/2022/2555/oj.

[165] Anu Bradford. *The Brussels Effect: How the European Union Rules the World*. New York: Oxford University Press, 2020. ISBN: 9780190088583. DOI: 10.1093/oso/9780190088583.001.0001.

[166] Oskar Josef Gstrein and Andrej J. Zwitter. "Extraterritorial Application of the GDPR: Promoting European Values or Power?" In: *Internet Policy Review* 10.3 (2021), pp. 1–30. DOI: 10.14763/2021.3.1576.

[167] European Commission. *AI Continent Action Plan Delivers Major Milestones*. European Commission news item, 9 April 2026; accessed 20 July 2026. 2026. URL: https://digital-strategy.ec.europa.eu/en/news/ai-continent-action-plan-delivers-major-milestones.

[168] European Parliament and Council of the European Union. *Directive (EU) 2019/1024 on Open Data and the Re-Use of Public Sector Information*. Official Journal of the European Union. Accessed 2026-07-18. 2019. URL: https://eur-lex.europa.eu/eli/dir/2019/1024/oj.

[169] Alexander Osterwalder and Yves Pigneur. *Business Model Generation: A Handbook for Visionaries, Game Changers, and Challengers*. Hoboken, NJ: John Wiley & Sons, 2010. ISBN: 9780470876411.

[170] Open Data Institute. *The Data Ethics Canvas*. Online tool and user guide. Published 28 June 2021; accessed 2026-07-20. June 2021. URL: https://theodi.org/article/the-data-ethics-canvas-2021/.

[171] Samar Fatima et al. "Public AI Canvas for AI-Enabled Public Value: A Design Science Approach". In: *Government Information Quarterly* 39.4 (2022), p. 101722. DOI: 10.1016/j.giq.2022.101722.

[172] Jeroen Van Der Hoven and Noemi Manders-Huits. "Value-Sensitive Design". In: *The Ethics of Information Technologies*. Ed. by Keith W. Miller and Mariarosaria Taddeo. 2nd ed. Routledge, 2020, pp. 329–332. DOI: 10.4324/9781003075011-23.

[173] Flora Cornish et al. "Participatory Action Research". In: *Nature Reviews Methods Primers* 3.1 (2023), p. 34. DOI: 10.1038/s43586-023-00214-1.

[174] High-Level Expert Group on Artificial Intelligence. *Ethics Guidelines for Trustworthy AI*. Tech. rep. Published 8 April 2019; accessed 2026-07-20. Brussels: European Commission, 2019. URL: https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai.

[175] Garante per la protezione dei dati personali. *Regole deontologiche per trattamenti a fini statistici o di ricerca scientifica*. Provvedimento n. 515 del 19 dicembre 2018. Published in Gazzetta Ufficiale No. 11 of 14 January 2019; accessed 2026-07-20. 2018. URL: https://www.garanteprivacy.it/home/docweb/-/docweb-display/docweb/9069637.

[176] Experts Working Group on Data Protection and Privacy. *Data Protection and Privacy Ethical Guidelines*. Tech. rep. Version 5. Document dated 18 September 2009; archived mirror accessed 2026-07-20. European Commission, 2009. URL: https://www.kowi.de/Portaldata/2/Resources/fp7/ethik/ethics-privacy.pdf.

[177] European Commission, Directorate-General for Research. *Ethics for Researchers: Facilitating Research Excellence in FP7*. Tech. rep. ISBN 978-92-79-05474-7; accessed 2026-07-20. Luxembourg: Publications Office of the European Union, 2007. URL: https://op.europa.eu/en/publication-detail/-/publication/413ea00b-a3b0-4dfa-9a76-2b695e2fb2cd/language-en.

[178] Amnesty International and Access Now. *The Toronto Declaration: Protecting the Rights to Equality and Non-Discrimination in Machine Learning Systems*. Human-rights declaration. Published 16 May 2018; accessed 2026-07-20. 2018. URL: https://www.amnesty.org/en/documents/pol30/8447/2018/en/.

[179] Ada Lovelace Institute. *Algorithmic Impact Assessment: A Case Study in Healthcare*. Tech. rep. ISBN 978-1-8382567-9-1; published 8 February 2022; accessed 2026-07-20. Ada Lovelace Institute, 2022. URL: https://www.adalovelaceinstitute.org/report/algorithmic-impact-assessment-case-study-healthcare/.

[180] Digital Catapult. *Lessons in Practical AI Ethics: Taking the UK's AI Ecosystem from 'What' to 'How'*. Tech. rep. Accessed 2026-07-20. Digital Catapult, Apr. 2020. URL: https://www.digicatapult.org.uk/wp-content/uploads/2020/05/20200430_DC_143_EthicsPaper__1_.pdf.

[181] Elham Tabassi. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Tech. rep. NIST AI 100-1. Gaithersburg, MD: National Institute of Standards and Technology, 2023. DOI: 10.6028/NIST.AI.100-1.

[182] Riccardo Nanni et al. *AI Ethics Canvas, Version 1.0.0*. Zenodo. Supplementary material; published 9 July 2026. 2026. DOI: 10.5281/zenodo.21282490.

[183] Marijn Janssen et al. "Data Governance: Organizing Data for Trustworthy Artificial Intelligence". In: *Government Information Quarterly* 37.3 (2020), p. 101493. DOI: 10.1016/j.giq.2020.101493.

[184] Nithya Sambasivan et al. ""Everyone Wants to Do the Model Work, Not the Data Work": Data Cascades in High-Stakes AI". In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2021, pp. 1–15. DOI: 10.1145/3411764.3445518.

[185] United Nations Economic Commission for Europe. *Data Stewardship and the Role of National Statistical Offices in the New Data Ecosystem*. Tech. rep. ECE/CES/STAT/2023/4. Accessed 2026-07-20. United Nations, 2024. URL: https://unece.org/info/publications/pub/387294.

[186] Silke Cuno et al. "Data Governance and Sovereignty in Urban Data Spaces Based on Standardized ICT Reference Architectures". In: *Data* 4.1 (2019), p. 16. DOI: 10.3390/data4010016.

[187] Aidan Hogan et al. "Knowledge Graphs". In: *ACM Computing Surveys* 54.4 (2021), 71:1–71:37. DOI: 10.1145/3447772.

[188] Harshvardhan J. Pandit, Declan O'Sullivan, and Dave Lewis. "Queryable Provenance Metadata for GDPR Compliance". In: *Procedia Computer Science* 137 (2018), pp. 262–268. DOI: 10.1016/j.procs.2018.09.026.

[189] World Wide Web Consortium. *Shapes Constraint Language (SHACL)*. W3C Recommendation. Accessed 2026-07-17. 2017. URL: https://www.w3.org/TR/shacl/.

[190] Anouk Oudshoorn, Magdalena Ortiz, and Mantas Šimkus. "SHACL Validation in the Presence of Ontologies: Semantics and Rewriting Techniques". In: *Artificial Intelligence* 352 (2026), p. 104483. DOI: 10.1016/j.artint.2026.104483.

[191] MLCommons. *Croissant Responsible AI Specification 1.0*. Community Specification. Accessed 2026-07-17. 2024. URL: https://docs.mlcommons.org/croissant/docs/croissant-rai-spec.html.

[192] Nicolas Prat, Isabelle Comyn-Wattiau, and Jacky Akoka. "A Taxonomy of Evaluation Methods for Information Systems Artifacts". In: *Journal of Management Information Systems* 32.3 (2015), pp. 229–267. DOI: 10.1080/07421222.2015.1099390.

[193] John Venable, Jan Pries-Heje, and Richard Baskerville. "FEDS: A Framework for Evaluation in Design Science Research". In: *European Journal of Information Systems* 25.1 (2016), pp. 77–89. DOI: 10.1057/ejis.2014.36.

[194] UK Government Digital Service and Department for Science, Innovation and Technology. *Guidelines and Best Practices for Making Government Datasets Ready for AI*. Guidance. Published 19 January 2026; accessed 2026-07-20. Jan. 2026. URL: https://www.gov.uk/government/publications/making-government-datasets-ready-for-ai/guidelines-and-best-practices-for-making-government-datasets-ready-for-ai.

Appendix A

Literature review audit protocol

A.1 Purpose and scope

The bounded search-and-screening audit documented here supports Chapter 2. It records the database queries, eligibility criteria, screening procedure, and deduplication method used to examine the principal literature strands and seek relevant critical or negative findings. The audit supports an interdisciplinary literature review; it is not presented as systematic or exhaustive and does not verify the legal sources or technical specifications considered by the conceptual framework.

The searches covered Common European Data Spaces and comparable data-space ecosystems; AI-training-data governance; legal requirements engineering and compliance by design; documentation, metadata, provenance and traceability; machine-readable policy and usage control; human oversight, overreliance and automation bias; and criticism relevant to the conceptual framework. Record-level screening logs are not included in this appendix.

A.2 Databases and access

Scopus was used as the primary database and OpenAlex as a supplementary source.

A.3 Complete search strings

A.3.1 Scopus

The following strings were executed through Scopus document search using the title, abstract, and keyword fields. No date, discipline, document-type, or language filter was applied.

1. `TITLE-ABS-KEY("Common European Data Space*" AND ("AI" OR "machine learning"))`

2. `TITLE-ABS-KEY("data space*" AND ("data governance" OR "compliance" OR "traceability"))`
3. `TITLE-ABS-KEY(("AI training data" OR "machine learning data") AND ("lawful access" OR "compliance" OR "governance"))`
4. `TITLE-ABS-KEY(("legal requirements engineering" OR "compliance by design") AND ("traceability" OR "automation"))`
5. `TITLE-ABS-KEY(("human oversight" OR "human in the loop") AND ("automation bias" OR ("compliance" AND ("governance" OR "decision making" OR "decision support"))))`

A.3.2 Supplementary OpenAlex search

The supplementary OpenAlex search used the Works API, with results ordered by relevance. Its broader full-text search and interpretation of Boolean-like expressions differ from Scopus `TITLE-ABS-KEY`, so the two sets of retrieval yields are not directly comparable. The submitted strings were:

1. `"Common European Data Space" AND (AI OR "machine learning")`
2. `"data space" AND ("data governance" OR compliance OR traceability)`
3. `("AI training data" OR "machine learning data") AND ("lawful access" OR compliance OR governance)`
4. `("legal requirements engineering" OR "compliance by design") AND (traceability OR automation)`
5. `("human oversight" OR "human in the loop") AND ("automation bias" OR (compliance AND (governance OR "decision making" OR "decision support")))`

A.4 Contrary-evidence searches

Three additional Scopus searches targeted criticism, failures, barriers, and negative findings:

1. `TITLE-ABS-KEY("data space*" AND (failure OR limitation* OR barrier* OR "governance failure" OR "compliance burden" OR "interoperability barrier*" OR "accountability gap*"))`
2. `TITLE-ABS-KEY(("automated compliance" OR "machine-readable polic*" OR ("policy enforcement" AND compliance) OR ("requirements traceability" AND compliance) OR "documentation debt") AND (failure OR limitation* OR barrier* OR burden))`
3. `TITLE-ABS-KEY(("human oversight" OR "human in the loop") AND (overreliance OR "automation bias" OR failure OR limitation* OR "decreasing accuracy"))`

The corresponding OpenAlex contrary-evidence strings were:

1. `"data space" AND (failure OR limitations OR barriers OR "governance failure" OR "compliance burden" OR "interoperability barriers" OR "accountability gaps")`

2. `("automated compliance" OR "machine-readable policies" OR ("policy enforcement" AND compliance) OR ("requirements traceability" AND compliance) OR "documentation debt") AND (failure OR limitations OR barriers OR burden)`

3. `("human oversight" OR "human in the loop") AND (overreliance OR "automation bias" OR failure OR limitations OR "decreasing accuracy")`

Table A.1 reports the pre-screening retrieval yields for the final queries. Included-study counts are outside the table's scope. The breadth of the OpenAlex results was retained as a feature of the audit rather than repeatedly narrowed to obtain more favourable yields.

Table A.1: Retrieval yields for the final literature-review queries

Search group	**Query**	**Scopus**	**OpenAlex**
Core	1	11	548
Core	2	354	13,724
Core	3	45	4,882
Core	4	20	2,206
Core	5	347	25,767
Contrary evidence	1	672	26,843
Contrary evidence	2	137	18,840
Contrary evidence	3	1,801	66,794

A.5 Eligibility criteria

A publication was included at title-and-abstract screening when it directly contributed to at least one of the following: data-space governance or architecture; governance, access, provenance or compliance controls for AI or machine-learning data; legal requirements engineering or compliance by design; machine-readable policy enforcement; documentation, provenance, lineage, auditability or traceability; human oversight, overreliance or automation bias; or relevant empirical, conceptual, technical or legal criticism.

A record was excluded when it was irrelevant, insufficiently connected to the chapter's research problem, a duplicate, or non-scholarly without a materially relevant official or technical function.

No language filter was applied at retrieval.

A.6 Screening procedure

Scopus records were ordered by relevance, and up to the first 20 displayed records from each core query were screened; all records were considered when fewer than 20 were returned. Screening used the titles and abstracts displayed by Scopus. OpenAlex records were likewise ordered by relevance. After within-database duplicates were removed, 97 unique Scopus records and 95 unique OpenAlex records were screened.

A.7 Deduplication procedure

Within each database, records were deduplicated primarily by database identifier and DOI. Normalised title served as the secondary key, with author, year, venue, and abstract used to distinguish materially different versions where necessary. The recorded cross-database comparison identified at least seven overlaps and yielded 185 unique records across the two bounded samples.

Appendix B

Legal and Technical Source Inventory for the Conceptual Framework

B.1 Purpose and scope

The inventory records the legal instruments and technical specifications that informed the conceptual framework developed in Chapter 5, making their selection and analytical function visible without presenting a complete account of European data regulation, a comprehensive survey of metadata standards, or a completed mapping between them.

The legal instruments were selected to illustrate different aspects of data access, sharing, reuse, and downstream AI-related activity in Common European Data Spaces; the technical specifications were chosen because their published semantics address recurring informational functions identified by the framework, including legal-source identification, catalogue description, provenance, quality, policy expression, machine-learning documentation, and structural validation.

B.2 Legal sources

The GDPR has a distinct role in this selection. It informs the reading of activities involving personal data but is not treated as an autonomous application. The Data Act example and EHDS illustration begin again from their own provisions and retain the definitions, cross-references, actors, objects, purposes, and temporal conditions needed for the particular analysis.

Table B.1: Selected legal instruments and their function in the conceptual framework

Instrument	**Function in the analysis**	**Selected focus**
General Data Protection Regulation [3]	Contextual legal dependency where the relevant activity involves personal data.	Scope, legal basis, purpose limitation, and the relationship with requirements arising under other instruments.
DSM Copyright Directive [57]	Illustrates how rights and exceptions may qualify downstream data use.	Text and data mining, lawful access, rights reservations, reproductions, and extractions.
Data Governance Act [56]	Supplies horizontal governance arrangements for protected public-sector data and data-sharing services.	Reuse conditions, competent public bodies, secure processing, data intermediation, and institutional roles.
Data Act [32]	Provides the bounded application used to illustrate the five-stage framework.	Article 33(1)(a), with micro-applications concerning collection-methodology provenance and descriptive sufficiency.
Artificial Intelligence Act [2]	Shows how downstream AI requirements depend on system classification, intended purpose, and dataset function.	Article 10 data-governance and quality requirements for relevant high-risk AI systems.
European Health Data Space Regulation [1]	Provides the sectoral illustration of how the framework must be reapplied when actors, objects, and procedures change.	Dataset catalogues, health data access bodies, secondary use, and Article 77(3).

B.3 Technical specifications

Table B.2 sets out the specifications discussed as modular building blocks. Version and formal status remain visible because the meaning and conformance implications of an element depend on the specification actually examined.

Table B.2: Technical specifications considered by the conceptual framework

Specification and version	**Issuer and status**	**Indicative informational function**
Core ELI 1.3 [98]	ELI Task Force and Publications Office; official ontology release	Identification of legal resources, expressions, formats, and subdivisions.
DCAT 3 [89]	W3C Recommendation	Description and distinction of catalogues, datasets, distributions, catalogue records, and data services.
DCAT-AP 3.0.1 [26]	SEMIC Recommendation	European application profile for interoperable catalogue and dataset descriptions.
HealthDCAT-AP Release 7 [95]	European Commission draft specification	Health-sector catalogue descriptions and selected EHDS-oriented relations.
PROV-O [28]	W3C Recommendation	Provenance relations among entities, activities, agents, derivations, uses, and roles.
Data Quality Vocabulary [29]	W3C Working Group Note	Quality dimensions, metrics, measurements, annotations, and related policies.
ODRL Information Model 2.2 [27]	W3C Recommendation	Expression of asserted permissions, prohibitions, duties, parties, actions, assets, and constraints.
SHACL [189]	W3C Recommendation	Validation of declared graph structures, datatypes, controlled values, and cardinalities.
MLDCAT-AP 3.1.0 [94]	SEMIC Candidate Recommendation	Description of machine-learning datasets, their characteristics, provenance, and related resources.
Croissant 1.1 [31]	MLCommons community specification	Description of files, record sets, fields, sources, transformations, and dataset splits.

Table B.2: Technical specifications considered by the conceptual framework (continued)

Specification and version	**Issuer and status**	**Indicative informational function**
Croissant Responsible AI 1.0 [191]	MLCommons community specification	Documentation of collection, preprocessing, limitations, bias, and sensitive information.

B.4 How the inventory should be read

Taken together, the entries document the materials used to formulate and illustrate the conceptual framework. The legal sources retain their own scope, actors, conditions, and effects, while the technical specifications retain the meanings and formal status assigned by their issuing bodies. Their inclusion identifies a relevant function for examination; it does not imply that the listed specifications collectively represent the selected legal instruments or that the selection measures legal or technical coverage.

Publications

Chapters 3 and 4 incorporate work produced with co-authors. Each chapter begins with a disclosure that distinguishes the published or accepted work from the revisions, updates, and synthesis developed specifically for this thesis.

Chapter 3 is based on Riccardo Nanni, Pietro G. Bizzaro and Maurizio Napolitano, "The False Promise of Individual Digital Sovereignty in Europe: Comparing Artificial Intelligence and Data Regulations in China and the European Union," *Policy & Internet* 16, no. 4 (2024): 711–726, `https://doi.org/10.1002/poi3.424`.

Chapter 4 is based on Riccardo Nanni, Pietro Giovanni Bizzaro, Albana Celepija, Munazza Usmani and Maurizio Napolitano, "AI Ethics Canvas: A Graphical Tool to Design and Deploy Ethical Artificial Intelligence for Public Administrations," *Data & Policy* (forthcoming, 2026).

Data Availability

The canonical release of the AI Ethics Canvas is version 1.0.0, available at `https://github.com/dclfbk/canvasaiethics` under the CC BY-SA 4.0 licence and archived at `https://doi.org/10.5281/zenodo.21282490`. Its availability supports inspection, reuse, and adaptation; it does not establish that the artifact is adequate for a particular organisation or use case.