Ref. Ares(2025)8171261 - 29/09/2025

COOPERANT

Deliverable 3.6

Digitalisation ontology and tool architecture

WP3 - Smart digitalisation, modelling and tools development

Version 1.0

COOPERANT Key Facts

Acronym	COOPERANT
Project title	Leading-edge cooperative advances towards the next generation of concentrated solar power (CSP) technology
GA nº	101172882
Starting date	01/10/2024
Duration-months	48
Call (part) identifier	HORIZON-CL5-2024-D3-01 (Sustainable, secure and competitive energy supply)
Type of Action	HORIZON-RIA 'Horizon Europe Research and Innovation Action Programme'
Topic identifier	HORIZON-CL5-2024-D3-01-10 (Next generation of renewable energy technologies)
Consortium	7 organisations, 4 EU Member States, 1 UK and 2 Switzerland
Model GA type	HORIZON Action Grant Budget-Based

COOPERANT Consortium Partners

N.	Partner	Acronym	Country
1	IDENER RESEARCH & DEVELOPMENT AGRUPACION DE INTERES ECONOMICO	IDE	ES
2	KRAFTBLOCK GMBH	KB	GE
3	CENTRO DE INVESTIGACIONES ENERGETICAS MEDIOAMBIENTALES Y TECNOLOGICAS	CIEMAT	ES
4	HOLOSS - HOLISTIC AND ONTOLOGICAL SOLUTIONS FOR SUSTAINABILITY, LDA.	HOLOSS	PT
5	PHASE CHANGE MATERIAL PRODUCTS LTD	PCMP	UK
6	FACHHOCHSCHULE ZENTRALSCHWEIZ HOCHSCHULE LUZERN	HSLU	CH
7	ENGICER SA	ENGI	CH

Disclaimer of COOPERANT project: COOPERANT project is funded by the European Union under the Grant Agreement Nº 101172882. This work was supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract Nº 2400402. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or CINEA. Neither the European Union nor the granting authority can be held responsible for them.

Deliverable Details

Deliverable Details	
Deliverable title	Digitalisation ontology and tool architecture
Lead Beneficiary of Dv	IDE
Authors	IDE
Reviewers	All partners
Work package Nª and title	WP3 - Smart digitalisation, modelling and tools development
Work package leader	IDE
Type of document	Document, Report
Dissemination level	PU-Public
Month	M12
Contractual due date	30/09/2025
Delivery date	24/09/2025

Executive summary

This deliverable defines the **digitalisation ontology** and **system architecture** of the **COOPERANT-AI TOOL**, developed under **Subtask 3.5.1.** It provides the semantic and digital foundations necessary to align the physical and digital development of materials and systems. This work fully aligns with the **Data Management Plan (DMP)** requirements set out in **Deliverable D1.2** under **Task 1.2**, adhering to the **FAIR principles** (Findable, Accessible, Interoperable and Reusable) to ensure proper data handling and reproducibility.

At the core of this deliverable lies the definition of a **graph-based semantic data model** tailored to the project's specific needs. The **COOPERANT ontology** serves as the semantic backbone of the tool, which is outlined in Annexe I: "Visualisation of the ontology inputs".

It is based on linked data principles and **standard frameworks** like **OWL** and **RDF**. Ontologies, modelling standards, and semantic web technologies were examined to identify best practices and tools for development. The **ontology** was developed using **Protégé** and enhanced with the **RDFlib** and **PyVis** libraries to incorporate interactive visualisation capabilities. There are three main parts to the semantic framework of COOPERANT: Inputs, Processing Modules, and Outputs. The framework discusses the semantic annotation and integration of various data sources, such as simulations, experiments, and external databases, enabling the system to comprehend the data consistently. This semantic foundation is closely tied to the tool's modular system design and data architecture. It illustrates the significance of ontology in facilitating the integration of various components within a project and work package.

The **ontology** and **system architecture** are concurrently developed, ensuring a close connection between the semantic structure and functional implementation. It establishes a modular, scalable framework that integrates **Artificial Intelligence (AI) algorithms**, control modules, user interfaces, and distributed data repositories. The architecture follows key industrial digitalisation concepts such as **RAMI 4.0** and **IIRA**, and builds upon the architectural principles established in previous IDENER experience, namely the HEU ISLANDER project. In addition, **REST APIs** and shared semantic definitions make interoperability possible by letting all parts of a system interact with each other.

The deliverable also includes information on how to get, store, and access data, all in line with the project's Data Management Plan (DMP) - Deliverable D1.2. These steps consider the need for long-term data transparency, reproducibility, and public access.

In conclusion, this deliverable sets a **comprehensive and interoperable digital framework** that supports the intelligent, **modular development of the COOPERANT-AI TOOL**. It fulfils the objectives of Subtask 3.5.1 by providing a robust semantic and architectural foundation for future activities, particularly **Subtask 3.5.2: Databases and Modules Coupling**. Defining a shared digital language and a scalable infrastructure enables consistent classification, integration, and semantic interaction with diverse data sources. These outcomes create the conditions for advanced data-driven functionalities, seamless collaboration across domains, and strong support for upcoming integration efforts, while also contributing to broader digitalisation initiatives in materials and process engineering.

Table of Contents

List of Figures

List of Tables

List of Acronyms

AAS	Asset Administration Shell
API	Application Programming Interface
CA	Consortium Agreement
CSP	Concentrated Solar Power
DI	Data Integration
DMP	Data Management Plan
DOIs	Digital Object Identifiers
DS	Digital Skills
DSS	Decision Support Systems
DT	Digital Twin
DX.x	Deliverable X.x
ETSI	European Telecommunications Standards Institute
FAIR	Findable, Accessible, Interoperable and Reusable
GDPR	General Data Protection Regulation
HTF	Heat Transfer Fluid
IIRA	Industrial Internet Reference Architecture
KPIs	Key Performance Indicators (KPIs)
MDO	Multidisciplinary Design Optimisation
ML	Machine Learning
OADR	Open Access to Data and Research
OPC UA	Open Platform Communications Unified Architecture
ORM	Object-Relational Mapping
OVR	Open Volumetric Receiver
OWA	Open-World Assumption
OWL	Web Ontology Language
RAMI 4.0	Reference Architecture Model for Industry 4.0
RBAC	Role-Based Access Control
RDF	Resource Description Framework
RDF4J	Resource Description Framework for Java
REST	Representational State Transfer
RIF	Rule Interchange Format
SAREF	Smart Applications REFerence
SH	Sensible Heat Storage
SHACL	Shapes Constraint Language
SPARQL	SPARQL Protocol and RDF Query Language

SRB	Stakeholder Replicability Board
SWRL	Semantic Web Rule Language
TES	Thermal Energy Storage
TX.x	Task X.x
UNA	Unique Name Assumption
W3C	World Wide Web Consortium
WP	Work Package

1 Introduction

1.1 Purpose of the document

This deliverable outlines the initial results of Task 3.5, which focuses on the development and deployment of the COOPERANT-AI TOOL. Specifically, it addresses **Subtask 3.5.1**, which is dedicated to defining the digitalisation ontology that forms the foundation for the **COOPERANT-AI TOOL's overall ontology and system architecture**. This ontology is essential for aligning the physical and digital development of materials and systems, as well as for defining the procedures for data acquisition, storage and management, based on the requirements from the **Data Management Plan (DMP)-Deliverable D1.2** under **Task 1.2**.

1.2 Objectives

This task directly contributes to achieving COOPERANT's specific objectives, particularly SO3.1, which focuses on defining the ontology and the internal architecture of the decision support tool and the energy management. The **scope and target of SO3.1** are defined as:

- COOPERANT **ontology and tool core architecture** under a modular and flexible REST API.

In addition, the digitalisation ontology serves as an essential enabler for the broader objectives of Work Package 3 (WP3), by providing a shared semantic framework that facilitates the following:

- Modelling, designing and optimising the thermal performance of the new CSP-TES prototype (SO3.2).
- Modelling alternative TES operation modes to predict full system performance for long-term analyses to increase dispatchability (SO3.3).
- Increasing the dispatchability of the CSP plant performance with an easy upscaling approach, including potential PV hybridisation (SO3.4).

This deliverable demonstrates progress by ensuring that the COOPERANT ontology and architecture enable seamless **interoperability** between models and data sources, support **modular integration** of AI algorithms, simulation models, and digital twin, and allow for **scalable**, **reusable** components that can be replicated across different CSP-TES configurations. It is asserted that this foundation ensures that the COOPERANT-AI TOOL supports **hybrid renewable integration** and **advanced energy management**, thereby enabling intelligent CSP-TES orchestration across both operational and scaling scenarios.

2 Digitalisation Ontology for COOPERANT-AI TOOL

2.1 Introduction to Digitalisation Ontology

In the context of COOPERANT, a **digitalisation ontology** is a formal, machine-readable approach to conveying domain knowledge that systematically defines and connects concepts, items, and processes that are important for designing, controlling, and optimising CSP-TES (Concentrating Solar Power with Thermal Energy Storage) systems. The ontology's clear representation of the connections and limitations between concepts, enabling semantic interoperability. The **primary goal** of creating a digital ontology in COOPERANT is to establish a **common semantic framework that integrates the project’s physical and digital parts**. This framework guarantees the utilisation of a unified vocabulary for discussing material creation, component behaviour, system control, and optimisation strategies. Thus, this facilitates the transparent dissemination of information, the utilisation of sophisticated analytics, and the amalgamation of data across various domains. This agreement is essential for achieving the project's goals of modularity, interoperability, and replicability.

COOPERANT utilises the **Web Ontology Language (OWL)** for the **formal representation of the ontology**. OWL functions as a Semantic Web standard that enables the definition of complex class hierarchies, properties, and logical axioms. OWL is well-suited for dynamic environments such as COOPERANT due to its foundation in automated reasoning, consistency verification, and knowledge inference. These environments must facilitate real-time decision-making processes.

Furthermore, the digitalisation ontology is designed to **interface seamlessly with other COOPERANT modules**, including the distributed control system, the hybrid digital twin, the replicability module, and the energy management system. By serving as the semantic backbone of the entire COOPERANT-AI TOOL architecture, the ontology ensures that all modules operate based on a consistent and shared understanding of the system’s structure and behaviour, enabling coordinated optimisation and control strategies across scales and domains.

2.1.1 Graph-based semantic data model

A popular framework for structuring and connecting information is the graph-based semantic data model. It represents knowledge using data triples, which establish directed relationships between entities and serve as the foundation for semantic graphs.

Edges or links connect nodes, also referred to as vertices, to form a graph. Nodes represent entities or concepts, while edges illustrate the connections between them. Both nodes and edges can possess descriptive metadata. Edges typically indicate the nature of the relationship and its directionality. This representation is particularly advantageous for modelling complex, interrelated domains such as communication networks, systemic architectures, and knowledge systems, where relationships hold equal significance to the data itself.

At the heart of the model are **data triples,** each consisting of three elements:

- **Subject:** the entity or resource being described.
- **Predicate:** the relationship or property connecting the subject and object.
- **Object:** the entity or value related to the subject.

These triples illustrate **Subject** → **Predicate** → **Object**, where the subject and object function as nodes, and the predicate serves as the directed edge linking them. Figure 1 illustrates a triple consisting of Apple as the subject, hasColour as the predicate, and Red as the object. These triples constitute the fundamental components of a directed semantic graph. The significance of each triple is contingent upon the types, attributes, and relationships of the elements within the graph.

Figure 1. Visual representation of a semantic triple as a directed graph. The subject Apple is connected to the object Red via the predicate hasColour.

Semantic graphs are typically governed by **ontologies** to provide structure and semantic clarity. These are formal frameworks that define the types of entities and relationships represented in the graph. Ontologies specify:

- **Classes:** high-level categories or types of entities (e.g., *Fruit, Colour, Taste*), which group individual instances under a shared concept. For example, *Apple* and *Banana* are instances of the class *Fruit*.

- **Object properties**: relationships linking entities. For example, *Apple → hasColour → Red* and *Apple → hasTaste → Sweet* use object properties to connect one entity (a Fruit) to others (a Colour and a Taste).
- **Data properties:** attributes or characteristics of classes using literal values. For example, *Apple* may have the data properties *hasWeight = 150g* and *hasSweetnessLevel = High*.

Figure 2 illustrates the ontology schema used in the earlier semantic graph examples. It shows how classes, object properties, and data properties provide structure and consistency to the graph model.

In this schema:

- The **nodes** represent individual entities, such as *Apple, Banana, Red* and *Sweet*.
- The **edges**, which are labelled with predicates such as *hasColour*, show the relationships between these entities.
- Each triple can be considered a directional link in the graph, forming a coherent knowledge structure.

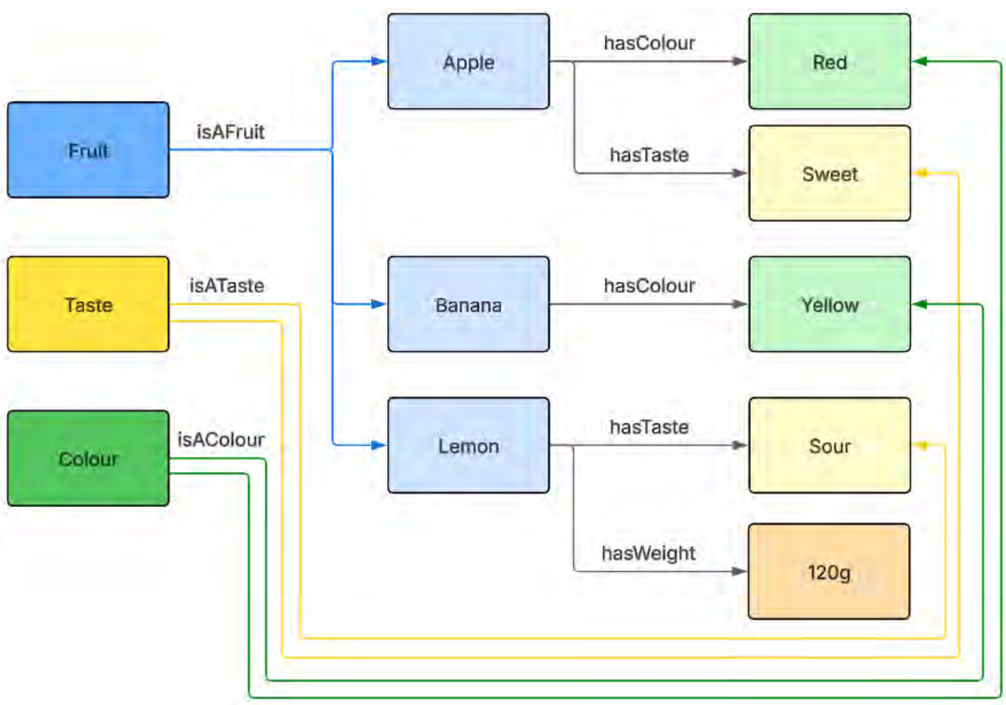

Figure 2. Ontology schema illustrating classes (Fruit, Colour, Taste), object properties (hasColour, hasTaste), and data properties (hasWeight).

Semantic graphs that effectively represent knowledge are predicated on fundamental presumptions that impact the interpretation, extension, and reasoning of information. The

Open-World Assumption (OWA) and the **Unique Name Assumption** (UNA) are two fundamental presumptions in knowledge representation and Semantic Web standards. These presumptions affect how data is modelled, queried, and validated in fields with dynamic or incomplete knowledge.

- **Open-World Assumption (OWA):** Postulates that no observer is fully knowledgeable about a domain. Since the world is regarded as "open", more information than what is currently understood may exist. As a result, listing an entity's known attributes does not mean that it is completely described or that there are no more attributes. This presumption is essential for semantic modelling because it permits incremental validation and flexible data definitions.
- **Unique Name Assumption (UNA):** This assumption states that each entity has a unique identifier. However, most Semantic Web technologies do not adopt UNA, allowing multiple identifiers to refer to the same entity. When such cases occur, equivalences between identifiers can be explicitly declared.

Together, OWA and UNA provide a logical foundation for knowledge representation, enabling reasoning and inference capabilities essential for semantic data models.

2.2 Existing Approaches, Concepts, and Technologies

In the context of digitalisation, data integration, and intelligent system development, a variety of approaches, conceptual models, and technological solutions have been proposed and implemented. This section provides an overview of the current landscape, focusing on ontological frameworks, standardised technologies and languages, widely adopted tools, conceptual methodologies, and the key challenges that persist in the field.

2.2.1 Existing approaches in ontologies

By defining concepts, relationships, and rules, **ontologies** offer a **formalised**, **standardised** way to represent knowledge domains. They make it easier for different systems to **integrate**, **reuse** knowledge, and **communicate** semantically. Building upon existing frameworks and guaranteeing compatibility requires an understanding of current ontology models. Several well-known ontologies are used as reference points or foundational models for manufacturing, industrial automation, and smart energy:

- **SAREF (Smart Applications REFerence ontology):** SAREF is a modular ontology created by the European Telecommunications Standards Institute (ETSI) with a primary focus on smart energy applications. Through specialised extensions like SAREF4ENER (energy) and SAREF4INMA (industry manufacturing automation), it has been expanded to cover other domains, including industry and manufacturing.

- **AAS (Asset Administration Shell):** An emerging standard by Platform Industrie 4.0, the AAS provides a digital twin representation of industrial assets. It standardises communication and interoperability between physical assets and software systems, allowing links to semantic vocabularies for richer knowledge representation and enhanced asset management.
- **OPC UA (Open Platform Communications Unified Architecture) Semantic Extensions:** OPC UA is a widely used protocol for industrial communication. To facilitate semantic reasoning, data integration, and enhanced interoperability in industrial settings, recent studies and initiatives concentrate on mapping OPC UA information models into OWL ontologies.
- **Descriptive Ontology:** A formal representation that aims to organise and capture a system's static elements, such as its constituent parts, their connections, and characteristics. Without attempting to model the behaviour or functionality of the system, it seeks to establish a common, unambiguous understanding of its structure.

Additionally, a wide range of domain-specific ontologies and knowledge graphs have been developed to illustrate manufacturing processes, sensor data, energy flows, and AI-driven diagnostics. These provide industrial and environmental-specific vocabularies that can be added to or modified as needed.

2.2.2 Standard technologies and languages

An examination of contemporary ontologies underscores the significance of understanding semantic web standards and languages that facilitate ontology creation, data modelling, and conceptualisation. The **World Wide Web Consortium (W3C)** has established a framework to facilitate computer inferences and interoperability among various types of semantic data:

- **Resource Description Framework (RDF)**: A flexible, graph-based data model that represents information using subject-predicate-object (or triple) representations. Schema-light data representation and linking between dispersed sources are made possible by RDF.
- **Web Ontology Language (OWL)**: an extension of RDF that incorporates advanced ontology constructs such as logical axioms, property restrictions, and sophisticated

class hierarchies. OWL facilitates the utilisation of automated reasoning and consistency verification in ontologies.

- **SPARQL (SPARQL Protocol and RDF Query Language)**: It is a standardised query language for RDF data retrieval and modification. SPARQL supports advanced pattern matching and federated queries over semantic graphs.
- **SHACL (Shapes Constraint Language)**: Outlines limitations and validation guidelines for RDF data to ensure data quality, compliance, and integrity in integrated semantic systems.
- **Semantic Rule and Reasoning Languages**: RIF (Rule Interchange Format) and SWRL (Semantic Web Rule Language) exemplify languages utilised for establishing intricate logical rules for inference, validation, and rule interchange among semantic frameworks.

These technologies establish a strong basis for modelling, querying, validating, and reasoning about the knowledge contained within ontologies, thus promoting interoperability and intelligent data processing.

2.2.3 Tools and Platforms for Linked Data and Semantic Web

Specialised software tools and platforms are needed to implement **linked data applications** and **semantic ontologies** in real-world systems. These tools facilitate collaborative development across geographically dispersed teams and enable the generation, editing, validation, querying, and reasoning of semantic data.

For instance, SHACL validation has three main components: a *data graph* containing the RDF data that needs to be checked, a *shapes graph* that provides constraints, and a *validation engine* that performs the compliance tests. While reasoners apply logical rules and axioms defined within the ontology to enable semantic inference, ontology editors help users create both shapes and data graphs.

Additionally, tools must support **version control**, **modularity**, **ontology alignment**, and **team collaboration** as semantic systems become more complex and larger, particularly in distributed or collaborative settings. The scope and level of semantic engineering projects vary, and each requires a different set of tools. Long-term project **sustainability**, **maintainability**, and **interoperability** can both be improved by choosing the right combination.

The primary tools used or contemplated in the COOPERANT project are detailed in the ensuing subsections, with an emphasis on their distinct functions and how they can be integrated into a semantic technology stack that is complementary and interoperable.

2.2.3.1 Protégé - Ontology Modelling and Validation

The project employs Protégé, a widely-used open-source ontology editing platform developed and maintained by Stanford University, to organise the development, validation, and management of the COOPERANT digitalisation ontology.

Protégé provides all the necessary tools for the creation and modification of ontologies in OWL. This enables the definition of semantic classes, properties, relationships, and data models in a modular, traceable, and extensible manner. This renders it an excellent option for the challenging semantic modelling tasks required by interdisciplinary systems such as the COOPERANT-AI TOOL.

Key features of Protégé include:

- **Advanced reasoning and querying capabilities**: Integration with reasoners like HermiT, Pellet, and Fact++ supports sophisticated semantic inference and validation.
- **Ontology visualisation**: Tools for exploring and understanding ontology structures improve usability and debugging.
- **Extensible plugin architecture**: Built in Java, Protégé supports custom plugins, enabling integration with additional tools such as alignment modules, specialised reasoners, or tailored user interfaces.
- **Collaborative development support**: **WebProtégé**, a browser-based interface, facilitates distributed collaborative ontology editing. While WebProtégé offers slightly reduced functionality compared to the desktop version, it is ideal for multi-partner projects, such as COOPERANT.

Protégé remains the de facto standard in ontology engineering. It is widely acknowledged and supported by a large global user community and an extensive academic ecosystem.

2.2.3.2 TopBraid Composer – Semantic Integration and Automation

For advanced modelling and the integration of semantic resources, the COOPERANT project also considers **TopBraid Composer**, a professional, commercial development environment that is part of the TopBraid Suite by TopQuadrant. This tool is helpful for projects that require complex integration mechanisms, as it facilitates both individual and group ontology building. Among its primary attributes are:

- Editing RDFS and OWL ontologies, SPARQL queries, and semantic rules in accordance with **standards**.
- **Integration of legacy data**, including strong import capabilities from relational databases, spreadsheets, and XML.

- Support for **reasoning and validation**, including integrated inference engines (like TopSPIN and SHACL) for debugging, advanced semantic reasoning, and consistency checking.
- **Ontology visualisation and diagramming**, providing graphical tools to improve comprehension of ontology structures and support communication among stakeholders.
- **Ontology refactoring**, enabling modular development and effective restructuring of large or interconnected ontologies.
- **SPIN and SPARQLMotion support**, offering declarative data transformation and workflow automation through a visual scripting interface.
- **Plugin extensibility**, because of its Eclipse-based architecture, can be integrated with other modelling tools (such UML and XML) and supports custom extensions like reasoning components and visualisations.

TopBraid Composer is a useful choice for semantic engineering activities that are outside the purview of lightweight or research-focused tools because of its robust feature set, Eclipse integration, and enterprise-oriented extensions.

2.2.3.3 RDF4J - RDF Application Development in Java

RDF4J (Resource Description Framework for Java) is an open-source Java framework for developing **RDF and linked data applications**. It offers a flexible API for working with RDF models and SHACL validation, as well as SPARQL querying and support for multiple backends.

- In-memory stores for testing and lightweight use cases.
- Native disk-based stores for scalable, persistent storage.
- Elasticsearch-based indexing for performance-oriented querying.
- Integration with third-party RDF databases such as GraphDB and Amazon Neptune**.**

RDF4J is best suited to embedding semantic capabilities directly within Java-based applications and microservices and offers both standalone and embedded deployment modes.

2.2.3.4 Apache Jena - Full-Stack Semantic Web Development

Apache Jena is another mature open-source Java framework for Semantic Web and linked data applications, supporting RDF, OWL, SPARQL querying, reasoning, and data storage via its TDB database and SPARQL endpoint [1].

[1] Apache Software Foundation. (s. f.). Apache Jena. https://jena.apache.org/

It is highly mature, with strong community support, and is widely used in both research and industry settings for building scalable semantic services and knowledge graphs. Jena also includes ARQ (SPARQL engine), Fuseki (SPARQL server), and a rule-based inference engine.

2.2.3.5 Eclipse KOMMA - RDF Object Mapping for Java

Eclipse KOMMA builds on RDF4J to provide **an object-relational mapping (ORM)** layer for RDF, OWL, and RDFS. KOMMA is particularly useful in projects where developers need to embed semantic data handling into Java-based applications with tight integration and minimal RDF-specific code [2].

2.3 Requirements for the COOPERANT-AI TOOL Ontology

The ontology must meet several critical requirements to support COOPERANT's goals:

- **Expressivity and Flexibility:** ability to model complex relationships among CSP components, thermal storage, control strategies, and environmental conditions.
- **Interoperability:** seamless integration with external data sources, simulation models, and control systems via standardised interfaces.
- **Scalability:** refers to the ability to extend the ontology with new components or use cases without requiring a major redesign.
- **Modularity:** a structure that allows for the independent development and reuse of ontology modules aligned with system components.
- **Real-Time Compatibility:** efficient representation that enables real-time reasoning and decision-making support within the COOPERANT-AI TOOL.
- **Data Governance:** integration of data acquisition, provenance, and security aspects to ensure reliable data management throughout the system lifecycle.

Protégé was selected as the **primary ontology development tool** for COOPERANT due to its alignment with the project's technical and collaborative requirements. Several justifications for selecting Protégé as the primary ontology tool for COOPERANT are presented below:

1. Full compliance with OWL standards, enabling rich semantic modelling.

[2] KOMMA: RDF Object Mapper. (n.d.). KOMMA documentation. Retrieved August 8, 2025, from https://komma.enilink.net/docs/

2. Integration with advanced reasoning engines for robust semantic validation and querying.
3. A flexible plugin system allowing customisation and extension.
4. Availability of a web-based collaborative environment (WebProtégé), which aligns with the distributed nature of the COOPERANT project consortium.
5. Proven stability, broad adoption, and strong community support.

The COOPERANT project employs Protégé to ensure the development of its digitalisation ontology is **coherent, reusable, and interoperable** with other systems. This ensures the coherence of the project modules' meaning and facilitates integration with external data sources, tools, and standards.

Scalability and adaptability are key considerations in the design of the COOPERANT ontology. Protégé's modular structure makes it possible to incorporate new elements, attributes, and design parameters without compromising existing classes. The ontology can be effectively expanded as new technologies or data sources become available. Long-term maintainability and interoperability with other digital tools and decision-support systems are supported by standardised definitions, transparent class hierarchies, and adaptable property schemas that make updates and modifications easier.

The project employs a specialised Python-based visualisation pipeline utilising the **RDFlib**[3] and **PyVis**[4] libraries to create distinctive, context-aware charts and facilitate visual representation. The RDFlib library enables the modification and integration of semantic data defined in Protégé by providing programmatic access to the ontology's structure and content. The PyVis package can be utilised to dynamically exhibit RDF graphs. It enables modifications to aspects such as colour schemes, node shapes, and layout positioning. This tailored solution surpasses conventional ontology tools by automatically extracting, querying, and formatting ontology elements into adaptable and informative visual representations. Modifying these visualisations dynamically according to user requirements or application contexts can enhance their comprehensibility and usability.

In general, this method significantly enhances the ontology's architectural representation and strengthens the readability, communication, and clarity of its underlying semantics, supporting both internal development and external stakeholder engagement.

[3] *rdflib 7.1.4 — rdflib 7.1.4 documentation*. (s. f.). https://rdflib.readthedocs.io/en/stable/
[4] Pyvis. (n.d.). *Pyvis documentation*. https://pyvis.readthedocs.io/en/latest/

2.3.1 Ontology co-creation with partners for technical validation

Co-creation methodology was used in COOPERANT's ontology development process to ensure technical accuracy and alignment with actual industrial practices. The project proposal and cutting-edge references served as the foundation for the initial design of the ontology, but all project partners' input and approval were crucial to its improvement.

Each partner provided data on the main components and properties relevant to their domain. These data inputs were consolidated and mapped into the ontology classes, relations, and attributes. Values for key properties were defined combining quantitative data from the proposal with expert feedback from the consortium.

The technical validation was carried out iteratively through several co-creation workshops and bilateral exchanges. Partners contributed as follows:

Materials Properties

- **CIEMAT** – Provided data on material properties for OVR components and guidance on the definition of relevant attributes.
- **ENGICER** – Provided porosity data for OVR materials, contributing to accurate definition of their properties.
- **PCM** – Provided data and guidance on defining properties of Storage Materials, including material costs and cyclability.
- **KRAFTBLOCK** – Provided data on material porosity for Storage Materials.

Design Parameters

- **HSLU** – Provided recommendations for the structure of the TES schema, helping to define a clear and coherent representation of the system.
- **CIEMAT** – Offered recommendations and guidance on OVR design parameters, helping to define operational constraints and component relationships.
- **ENGICER** – Advised on design parameters for OVR components, supporting practical configuration and operational constraints.
- **KRAFTBLOCK** – Delivered technical parameters and specifications for TES-related design.

Policy Recommendations and Sustainability KPIs

- **HOLOSS** – Provided advice on the structuring and distribution of outputs related to policy recommendations, ensuring clarity and alignment with project objectives.

Final Review and Feedback

- All partners participated in the final ontology review, providing feedback on class structure, property naming, and consistency with project terminology.

Through cooperation, the ontology was made technically sound, interoperable, and representative of the real data and constraints pertinent to COOPERANT's use cases. The outcome is a strong framework for knowledge representation that can serve as a guide for future developments in digital tools and decision-support systems.

2.4 COOPERANT's semantic approach

The **COOPERANT-AI Tool** facilitates the integration of data and models across diverse domains. This facilitates optimal control, oversight, and replication of **Concentrated Solar Power (CSP)** systems integrated with **Thermal Energy Storage (TES)**. This integration is essential to address the inherent complexity and diversity of data sources, models, and operational requirements in CSP-TES systems. It enables the formulation of robust decisions and the adaptation to real-life changes.

- The ontology comprises three primary components: inputs, processing modules, and outputs. This ensures that the digital ecosystem is modular, traceable, and interoperable. This framework provides a coherent and effective method for systems to communicate and exchange data. Inputs are categorised into groups, including Design Parameters, Operational Parameters, Environmental Conditions, and Material Properties.
- Processing Modules encompass a range of functionalities, including the following: Data Preprocessing, AI Algorithms, Multiphysics and Mathematical Modelling, and the hybrid reconfigurable Digital Twin (DT).
- Outputs are structured into elements such as Policy Recommendations, Key Performance Indicators (KPIs), Replicability Strategies, and Real-time System Commands.

The ontology establishes a comprehensive network of designated semantic relationships that connect various types of information, including materials, operational strategies, system components, and modelling techniques. Organising these semantic links into functional groups elucidates their roles and interdependencies. This framework facilitates the collaboration of external models and data sources within the organisation. It ensures that all components are distinct and delineated. Figure 3 illustrates the primary components of the ontology and their interrelations concerning meaning. This provides an overview of the system's configuration and the interoperation of its components.

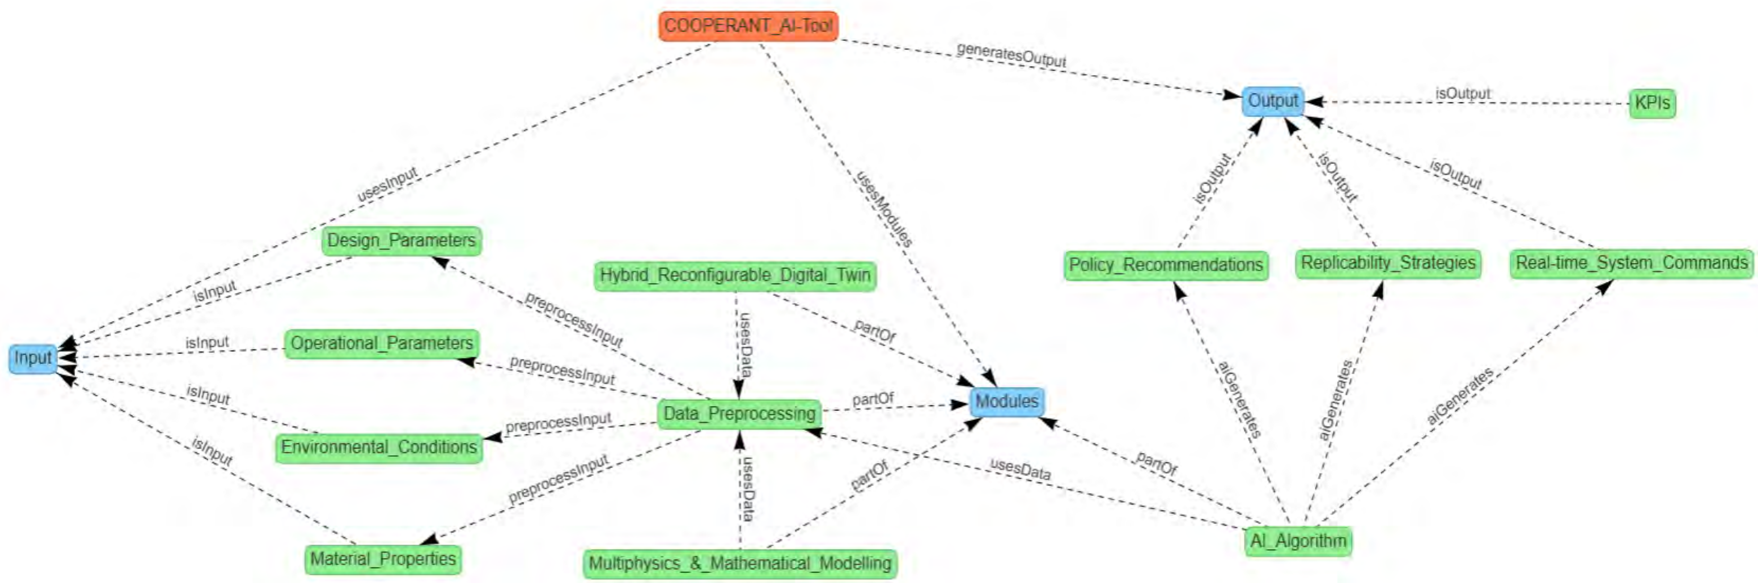

Figure 3. Visualisation of the main ontology modules and their relationships.

As seen in Figure 3, elements such as **Design Parameters, Operational Parameters, Environmental Conditions**, and **Material Properties** are connected to the general **Input** category through the *isInput* relationship, signifying their role as recognised data inputs. The COOPERANT-AI TOOL references these inputs via the *usesInput* relationship, establishing that it actively consumes and integrates these data categories into its processing workflow.

Additionally, the COOPERANT-AI TOOL references Processing Modules via the *usesModules* relationship. These modules, which represent different processing types such as **Data Preprocessing, AI Algorithms, Digital Twin (DT), and Multiphysics and Mathematical Modelling**, are all linked by a *partOf* relationship to the overarching **Module** category.

Focusing on the **Data Preprocessing** module, it is linked through the *preprocessInput* relationship to modules that supply raw input data, whose role is to prepare and cleanse this data for further analysis. This preprocessing module, in turn, connects via the *usesData* relationship to modules that utilise the cleaned and structured data, such as the AI algorithms, digital twin simulations, and mathematical modelling modules.

The **outputs** generated by the COOPERANT-AI TOOL (Policy Recommendations, Replicability Strategies, and Real-time System Commands) are linked via the *generatesOutput* relationship and are themselves connected by *isOutput* to the general Output category. Moreover, the AI-Tool module ismodule is specifically linked to these outputs through the *aiGenerates* relationship, reflecting its role in producing actionable insights and control commands. To fully understand the scope of the ontology, it is important to examine the other semantic relations that support the framework. The following subsections will examine each module in detail, along with its subcomponents and associated properties. To facilitate a clear understanding, the specific types of semantic relationships present within each module are described at the outset.

Material classification relationships (e.g., *isAStorageMaterial*, *isAReceiverMaterial*) specify the functional role of each material within the system, while **property-based relationships** (e.g., *isASensibleHeatProperty, isAPCMProperty*) associate materials to their thermal, physical, and operational performance attributes, ensuring accurate and comprehensive system characterisation (see Figure 4).

Design parameter relationships (e.g., *isAHTFParameter*, *isAOVRParameter*, *isATESParameter*, *isEquipment*, *isGeometryData*) link system components to their corresponding design parameters. These include operational conditions, geometric configurations, and specific equipment specifications, facilitating the precise modelling of system architecture and functionality (see Figure 5).

Environmental and contextual relationships (e.g., *isAEnvCondition*) ensure that all system elements are assessed within realistic operating conditions and scenarios, including factors such as time of day, solar radiation and geographic location (see Figure 6).

Moving towards dynamic aspects, **operational relationships** (e.g., *isAOperationalParameter, isASetPoint*) describe dynamic behaviours under various control scenarios, providing essential information to the distributed control system (see Figure 6).

Processing modules relationships (e.g., *isAPreprocess*, *isAMLModel*, *isAAIAlgorithm*) define the roles of each module, while interaction relationships (e.g., *worksWith*, *trainsWith*) specify the collaboration and data exchange patterns between modules, supporting modular and scalable system design (see Figure 7).

Lastly, **implementation**, **design**, and **optimisation relationships** (such as isAKPI, generatedBY, and providesOptimisation) link performance indicators, data provenance, and optimisation capabilities to support adaptive control, decision-making, and continuous system improvement (see Figure 8).

Altogether, these semantic connections guarantee that the ontology functions as a formally defined, logically organised, and interoperable framework. It offers a common language and reasoning framework that unifies models, data, and decision-making procedures throughout the COOPERANT-AI TOOL ecosystem.

2.4.1 Input

The **Input layer** encompasses a wide array of heterogeneous data and parameters essential for accurate modelling, intelligent control, and data-driven decision-making. These inputs are categorised in the following key domains, which will be analysed in detail in this section. A comprehensive overview of the overall Input layer structure is available in Annexe I (see Figure 11).

2.4.1.1 Material properties

As can be seen in Figure 4 the tool evaluates the **material properties** of all key system components. This includes the **Open Volumetric Receiver (OVR)** materials, including the absorber and cup, composed of silicon-infiltrated silicon carbide (Si-SiC).

The analysis also looks at **storage materials**, assessing **Sensible Heat Storage (SH)** and **Phase-Change Materials (PCM)**. While density, thermal conductivity, operating temperature, and cyclability are among the common characteristics of both classes, their fundamental thermal storage mechanisms are where they diverge most.

PCMs allow energy storage during transitions between physical states (such as solid to liquid) by storing and releasing thermal energy through **latent heat** linked to a specific **phase-change temperature**. Several important material parameters affect this thermal buffering capacity:

- The quantity of thermal energy stored or released during the phase transition is known as **latent heat**. A higher latent heat value enables the storage of more energy per unit mass or volume, which determines efficiency.
- PCM's ability to absorb or release heat is influenced by **thermal conductivity**. Materials with higher thermal conductivity, like metallic PCMs, enable faster charging and discharging rates.
- The PCM's effective thermal window is defined by the **operating temperature** and **phase-change temperature**. These must match the thermal requirements of the application (e.g., industrial heat processes).
- The term "**cyclability**" describes a material's capacity to repeatedly melt and solidify without degrading. Long-term stability is crucial because PCMs are designed to function over a wide range of heat cycles.
- The volumetric energy storage capacity is influenced by **density**. Compact system designs benefit from the ability of denser materials to store more energy in a smaller volume.

In contrast, **SH materials** store energy through **sensible heat**, where temperature changes occur without a phase transition. Such energy is captured through a set of properties or parameters:

- The **sensible heat capacity** of a material is directly related to how much energy it can store per unit of mass and temperature increase.
- Heat absorption and release rates are influenced by **thermal conductivity**; higher values lead to more efficient storage, which is essential for effective energy buffering. Materials with higher conductivity ensure faster heat exchange, which is particularly beneficial in dynamic thermal systems.
- The thermal window that the SH material can be used in without degrading is defined by the **operating temperature range**. Stability and performance depend on selecting a material that can withstand the 700–1300 °C target process temperature.
- The energy stored per unit volume is affected by **density**. Systems with higher densities are more compact, which is useful in situations where space is limited.
- The durability of a material over numerous heating/cooling cycles is measured by its **cyclability**. Long-term dependability with little performance loss is guaranteed by materials with high cyclability (~10,000 cycles).

- The function of **porosity** is complicated. High porosity can enhance weight characteristics or component integration, but it may decrease mechanical strength and thermal conductivity. Porosity must therefore be optimised for each application.

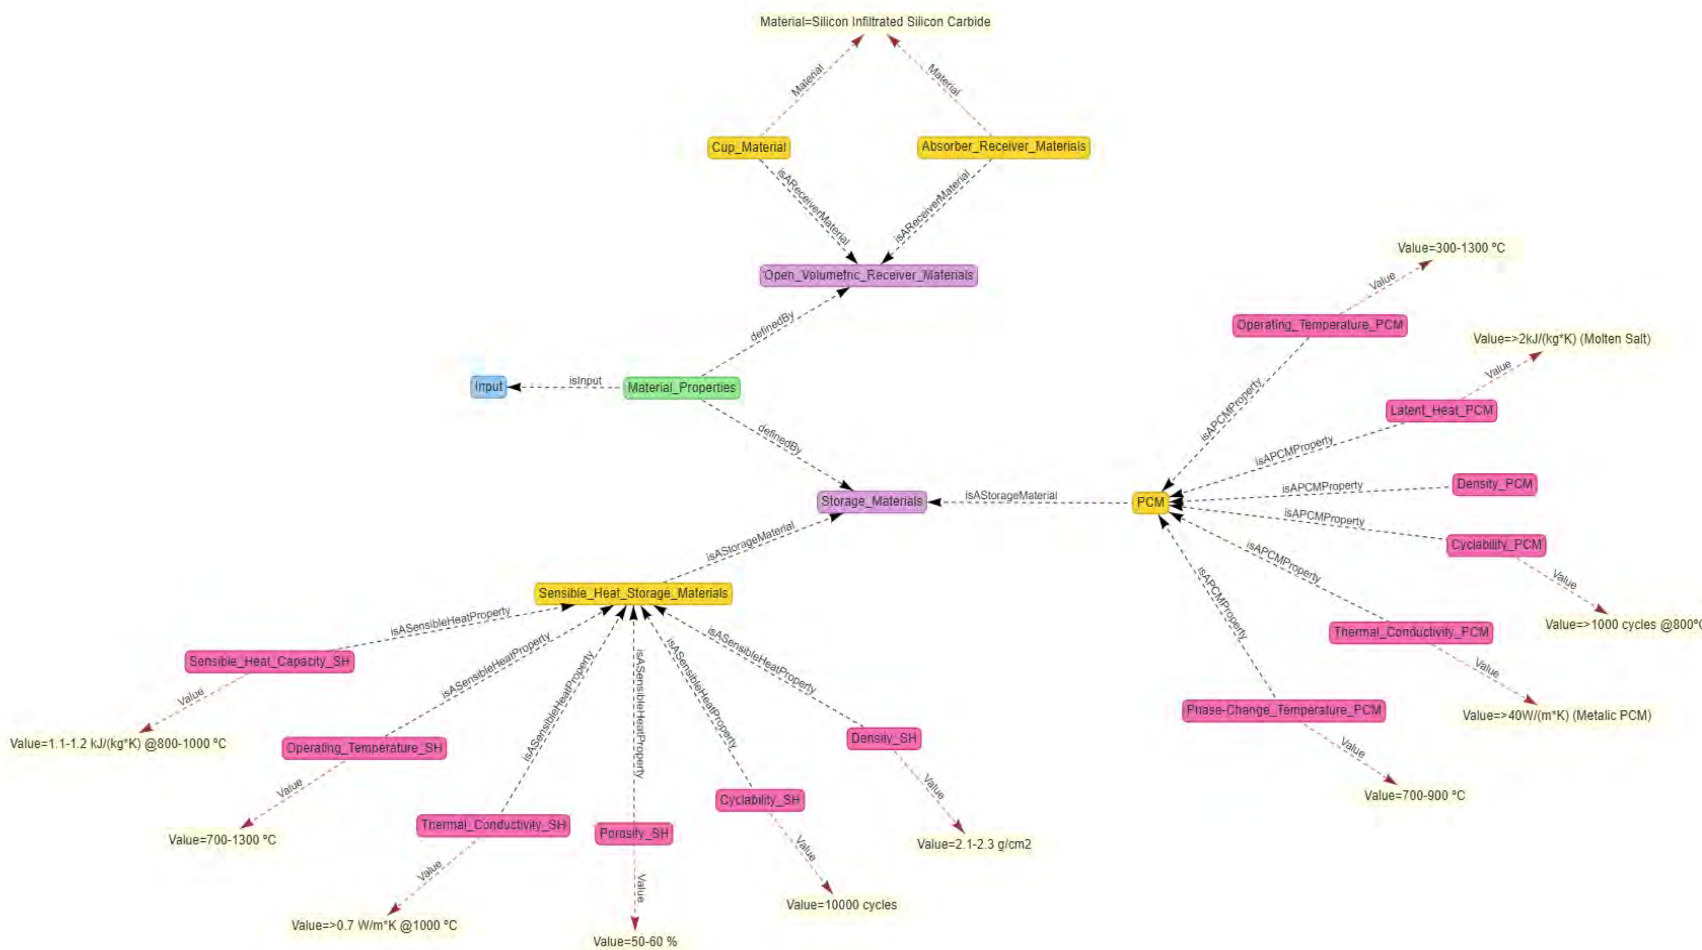

Figure 4. Visualisation of the material properties module.

2.4.1.2 Design parameters

As shown in Figure 5, the COOPERANT-AI Tool requires a thorough set of **design parameters** that specify the system's operational logic and physical structure. These parameters must be ontologically arranged to represent the interdependencies and relationships between the system's core subsystems.

The **Open Volumetric Receiver (OVR)**, a crucial part that absorbs and transforms concentrated solar radiation into thermal energy, is the brain of the **Concentrated Solar Power (CSP)** system. This system absorbs and changes concentrated solar radiation into thermal energy. The **OVR's power capacity**, **operating temperature range**, **geometric design**, and **thermal efficiency** are just a few of the important performance factors that compose it. The ability of COOPERANT design to function at temperatures as high as 1000ºC is a crucial differentiator. Interestingly, because each cup adds a specific percentage to the system's total thermal capacity, its power capacity is directly correlated with the number of cups or receivers.

Another essential aspect is the **pressure drop across the receiver**, which influences fluid dynamics and system pumping requirements. This pressure drop is influenced by several geometric parameters, including cell size range, cell distribution, and **total porosity** of the receiver structure. Narrower or longer flow paths, more complex internal geometries or tighter spacing between elements increase flow resistance, resulting in higher pressure losses. These increased losses have direct implications for the design and power demand of the pumping system, affecting the overall energy efficiency of the CSP plant.

The heliostat field is an important part of collecting solar energy, in addition to the **OVR**. It has a set of mirrors that can move around and follow the trajectory of the sun, sending sunlight to the receiver. The **number of heliostats**, the **area of their reflective surfaces**, the **optical reflectance of the mirrors**, and the **tracking accuracy** (usually shown as slope error) all determine how effectively this subsystem works. These parameters work together to determine the optical efficiency and effective solar concentration ratio. Both are important for getting the most energy out of the CSP plant.

To capture its dynamic behaviour and interaction with the CSP, the **Hybrid Thermal Energy Storage (TES)** subsystem is also heavily parameterised. Its geometrical configuration, which includes **mass**, **volume**, **shape**, and **internal layout**, including the number of flow channels, as well as its energy storage density, which shows the ability to store thermal energy per unit volume, defines it. These physical and geometrical characteristics are closely related to the heat exchange surface area and thermal inertia of the system.

The **air flow dynamics** within the TES system are also modelled, especially when operating with high-temperature air as a heat transfer medium. This makes it possible to accurately simulate the distribution of heat within the storage volume as well as cycles of charge and discharge. Parameters such as storage duration, thermal response time, and inlet and outlet temperatures provide a more thorough depiction of its operational behaviour. behaviour is given by parameters like storage duration, thermal response time, and inlet and outlet temperatures.

Specifically, the air flow rate and the **inlet and outlet temperatures** work together to shape the packed bed's thermal behaviour. The outlet temperature shows how much heat has been transferred, while the inlet temperature creates the initial thermal gradient. The air flow rate determines the residence time of the fluid and directly influences how deeply and uniformly heat penetrates the storage medium.

The **Heat Transfer Fluid (HTF)** is another core component whose properties must be precisely defined. The model considers the **fluid type** as well as its **thermal conductivity** and **heat capacity**, which control the plant's heat transfer efficiency and rate. When simulating transient states and energy flows between the OVR and TES subsystems, these characteristics are particularly important. The model also considers the **pressure drop** across the HTF loop, which has an impact on pump sizing and system losses.

The ontology also incorporates all pertinent **Ancillary Equipment** to guarantee a realistic depiction of the plant. This includes the construction and operating specifications of auxiliary parts that support fluid circulation, thermal control, and energy recovery, such as pumps, cylinders, and heat exchangers.

Figure 5 provides a visual overview of how these design parameters are interconnected within the COOPERANT ontology.

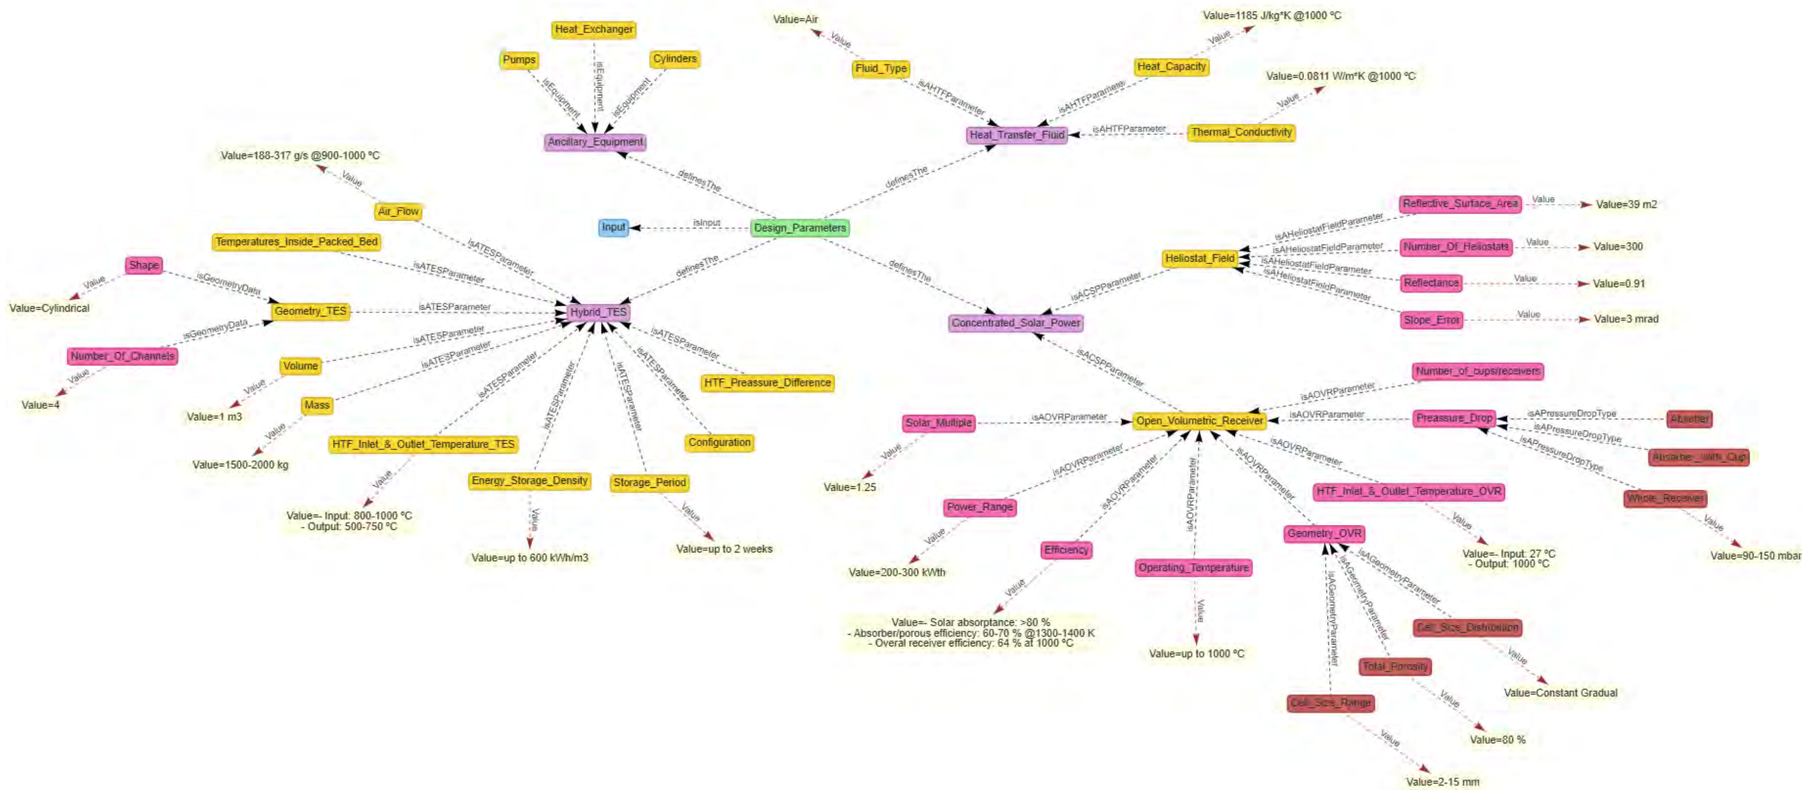

Figure 5. Visualisation of the design parameters modules.

Funded by
the European Union

2.4.1.3 Environmental conditions

Figure 6 illustrates the **environmental conditions** module. This module must process a variety of environmental data to model solar energy accurately. It documents significant elements that directly affect the operation and efficiency of the Concentrated Solar Power system. These include **temporal information**, the time of day, **solar radiation**, and **geographic coordinates** (latitude and longitude). To enable accurate estimation of the available solar resources, the module integrates these inputs to compute potential energy inputs and supports dynamic simulation, control, and optimisation tasks.

2.4.1.4 Operational parameters

The **operational parameters** can also be observed in Figure 6. These define the dynamic, real-time, or scheduled conditions under which the system operates. They act as essential control inputs and boundary conditions within the simulation and optimisation framework of the COOPERANT-AI Tool. These parameters include specific control setpoints, such as **target temperatures**, **charging and discharging rates**, and other values that directly influence system behaviour. In addition to these, the module incorporates broader operational configurations, including **storage cycling strategies**, **demand profiles, and hybridisation settings** with auxiliary systems, such as photovoltaics. The structured representation of these parameters is fundamental for enabling realistic system modelling, supporting predictive and adaptive control strategies, and ensuring effective energy management.

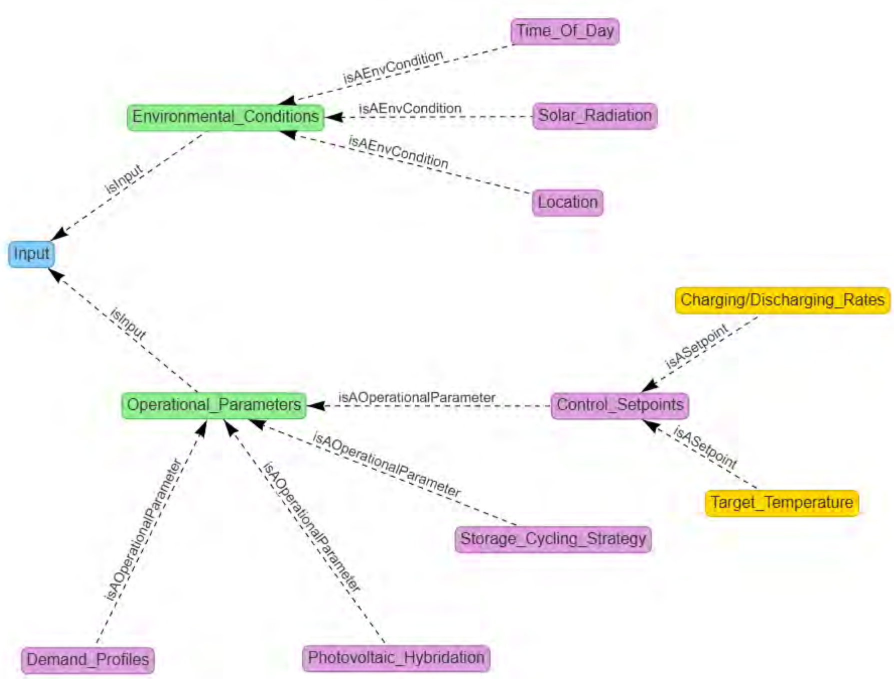

Figure 6. Visualisation of the modules of both operating parameters and environmental conditions.

2.4.2 Processing modules

The COOPERANT-AI TOOL integrates a set of advanced **processing modules** designed to enable seamless data acquisition, modelling and decision support across multiple system scales (see Figure 7). The modules are described as follows:

2.4.2.1 Data Preprocessing

This module is responsible for **cleaning, normalising, and generating new features** from data. It ensures that data obtained from experiments, sensors, and simulations is coherent and applicable for subsequent utilisation in AI and modelling processes.

2.4.2.2 AI Algorithms

A suite of **Machine Learning (ML)** techniques is applied to support various functionalities:

- ML for **material modelling**: Simulates and proposes alternative Phase Change Materials (PCMs) with improved thermal performance.
- Neural networks are employed to support the **dynamic behaviour modelling** of the digital twin, facilitating accurate, real-time, and predictive simulations.
- Reinforced Deep Learning: Reinforcement learning approaches are integrated with deep learning architectures to enhance the **optimisation algorithms** used in Multidisciplinary Design Optimisation (MDO), enabling intelligent design-space exploration and decision-making.

2.4.2.3 Multiphysics and Mathematical Modelling

The system modelling follows a hierarchical structure with **three levels** of abstraction:

- **Level 1 - Material Capsule or Subcomponent Level**: Focuses on the thermophysical and structural behaviour of materials and individual components.
- **Level 2 – TES System Level**: Models the complete Thermal Energy Storage system, including energy flows, control strategies and operational dynamics.
- **Level 3 – Integrated Plant Level**: Integrates the Thermal Energy Storage (TES) with the complete Concentrated Solar Power (CSP) plant configuration, examining its interaction with other components of the plant, such as the solar field and power block, and assessing the overall system efficiency.

2.4.2.4 Hybrid Digital Twin

The **COOPERANT CSP-TES** prototype interfaces with the **hybrid Digital Twin (DT)**, which acquires data from control systems such as valves, actuators, and electrical components, in addition to monitoring instruments like thermocouples and flow meters. It integrates high-fidelity physical models with AI-driven predictive algorithms to enable real-time monitoring, prediction, and scenario testing. This ensures that the system operates more efficiently, is more manageable, and can be readily replicated and modified for various configurations and applications.

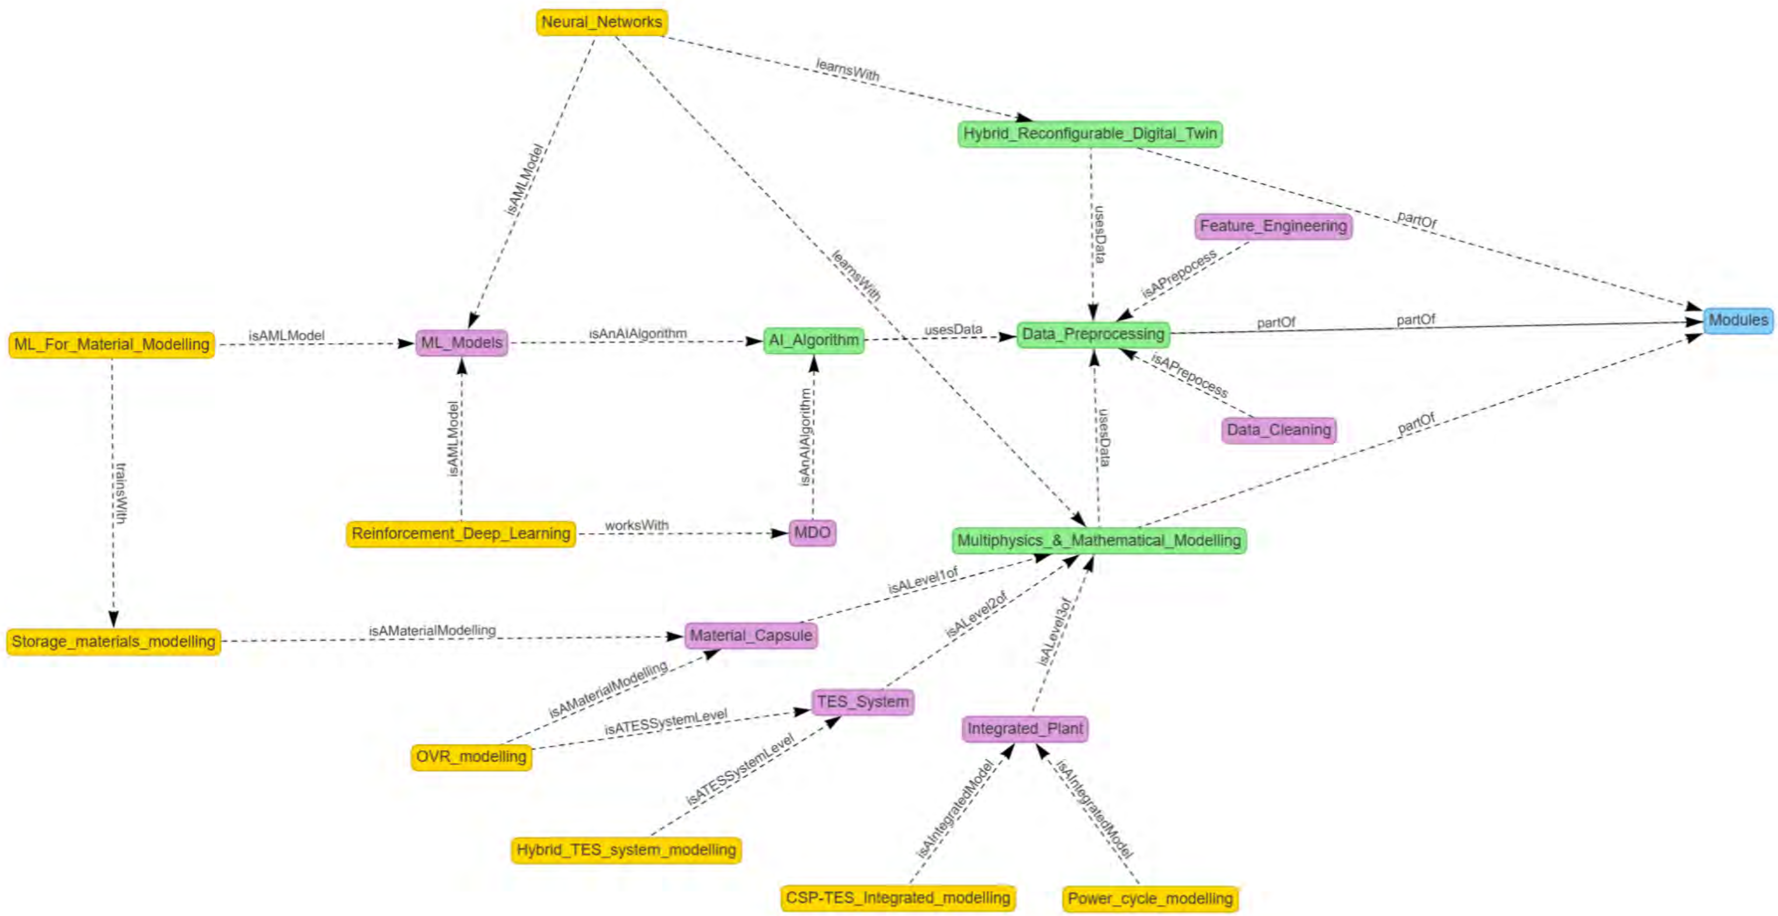

Figure 7. Visualisation of the processing modules within the ontology, including data preprocessing, multiphysics and mathematical modelling, AI algorithms (ML and MDO), and the hybrid reconfigurable digital twin, along with their interconnections.

2.4.3 Output

The COOPERANT-AI TOOL generates a range of actionable **outputs** designed to support decision-making, system optimisation, and knowledge transfer across different application contexts (see Figure 8). These **outputs** are integrated into the **COOPERANT-Transfer** module, which supports **reporting** and **visualisation** functionalities. This integration promotes clear stakeholder communication, fosters transparency throughout all processes, and provides a solid basis for technical improvement and policy alignment. The outputs include:

2.4.3.1 Policy Recommendations

Based on comprehensive simulations and AI-driven analyses, the system provides what-if scenario evaluations that inform decision-makers on the implications of different design, operational, and policy choices. These recommendations support evidence-based planning that is aligned with sustainability, efficiency, and circular economy objectives.

2.4.3.2 Key Performance Indicators (KPIs)

A quantitative set is established to monitor, evaluate, and compare the system's performance across various configurations and conditions. Certain KPIs serve as indicators of circularity, energy efficiency, temperature mapping, cost and profit, environmental impacts, and socioeconomic effects.

2.4.3.3 Replicability Strategies

This output includes design and operational guidelines that enable the replication of CSP-TES systems across different geographical, climatic, and industrial contexts. It covers:

- **Feasibility assessments** and design recommendations for diverse use cases.
- **Replication pathways** for similar systems in varying environments and application domains.
- **Operational strategies** tailored for flexibility and performance under different boundary conditions.
- **Optimised design parameters** generated through AI-based design space exploration.

2.4.3.4 Real-Time System Commands

The **hybrid digital twin** generates **real-time control recommendations**, including:

- **Optimised control strategies** to enhance system responsiveness and efficiency.
- **Time-evolving operational parameters** based on predictive modelling and continuous feedback from sensors and simulations.

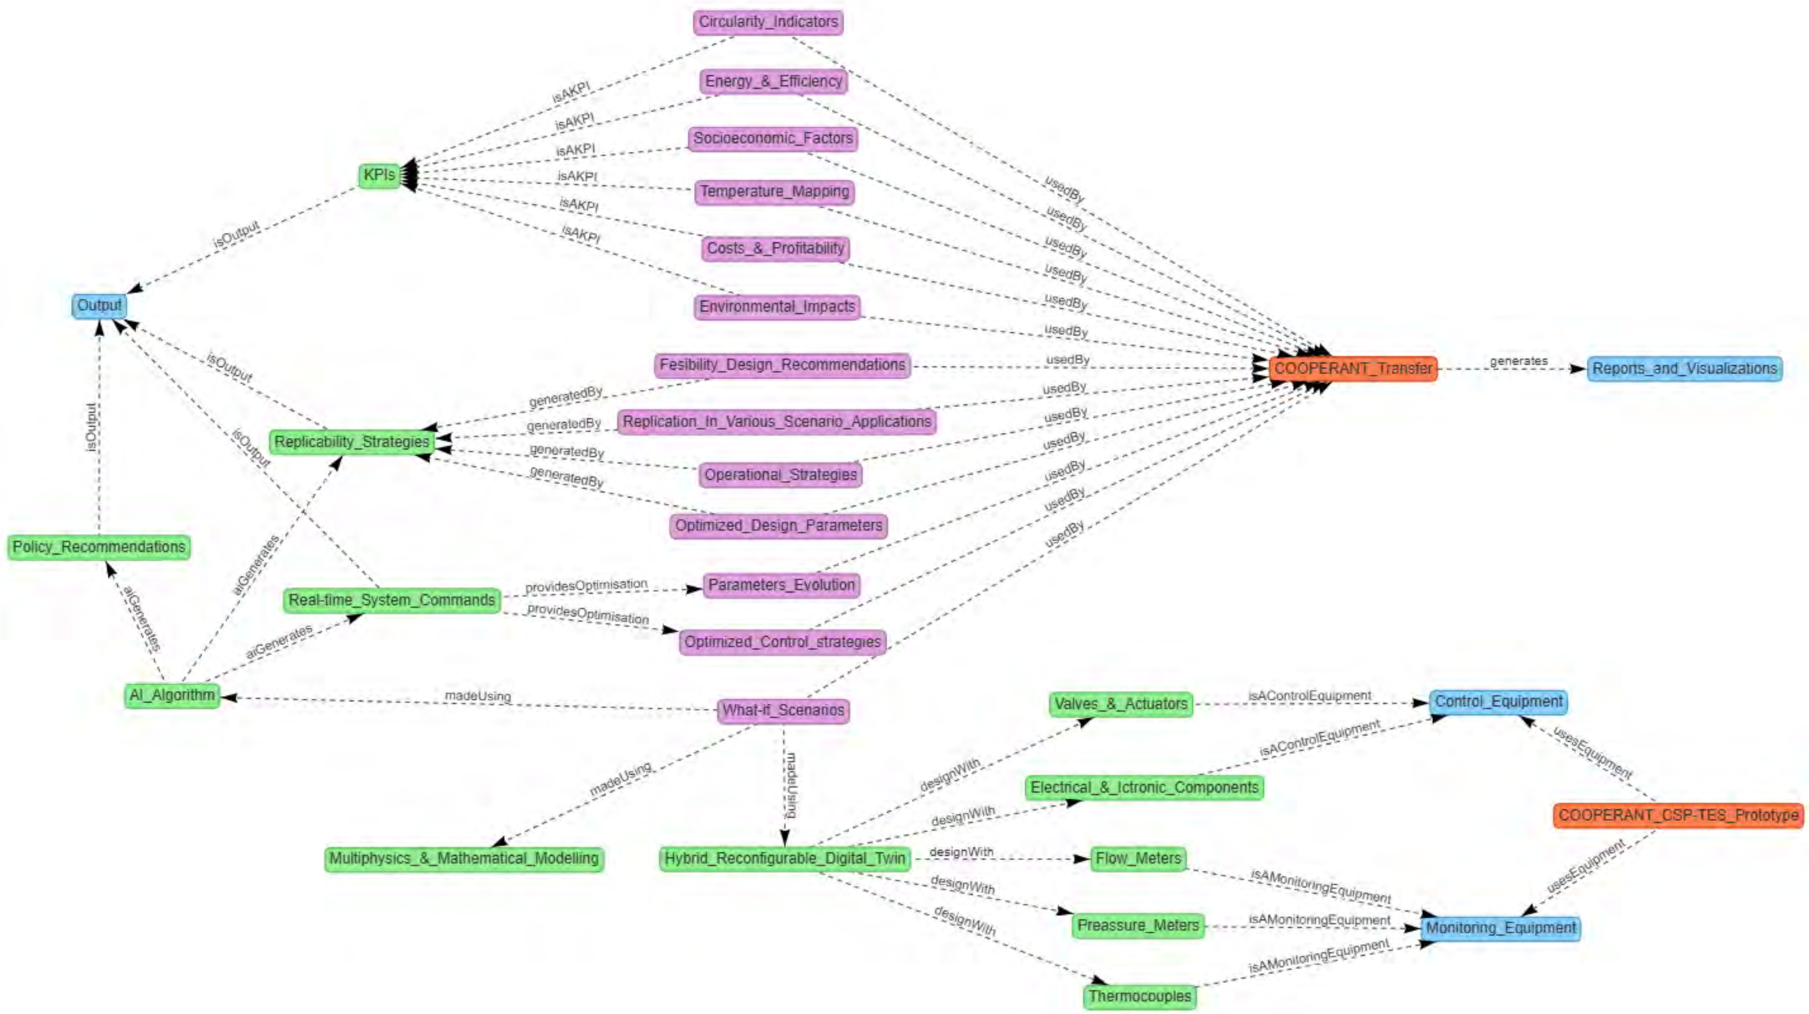

Figure 8. Visualisation of the ontology outputs.

3 Architecture – COOPERANT-AI TOOL

This section provides a detailed breakdown of the responsibilities and interactions of the main software and hardware modules that constitute the **COOPERANT architecture**, structured across three integrated tiers: **Operational Environment & Data Acquisition, Interoperability & Control Coordination**, and the **COOPERANT-AI TOOL**. Each of these tiers provides a group of modules designed to ensure secure, scalable, and standardised data flows.

All these data flows are fully aligned with the semantic principles of the ontology, as defined in Section 2, meaning every exchange follows a shared vocabulary and structure. This ensures that the architecture enables **semantic interoperability**, allowing for the consistent interpretation of data across all tiers.

The role of each software module is described within its tier, outlining how it contributes to the platform's objectives of interoperability, replicability, and informed decision-making. A **REST-based, secure, and modular architecture** is assumed throughout, ensuring reliable communication and extensibility across the platform lifecycle while preserving the semantic coherence established by the ontology.

3.1 Requirements

Based on the technical objectives analysis, the following **functional and non-functional architecture requirements** have been defined for the COOPERANT-AI TOOL. These requirements describe both the core **capabilities** the platform must deliver (what it should do) and the **quality attributes** or constraints under which it must operate (how it should perform).

Table 1 presents a summary of these requirements, grouped under **five functional groups** representing core capabilities of the platform. The first group focuses on the key architecture components, which form the backbone for the other services. The remaining groups concern multiple services and applications that will be deployed and reliant on a common data spine, as they will either consume or produce data via the COOPERANT-AI TOOL. Requirements are labelled with **module-specific codes** for traceability: **DI** for Data Infrastructure, **DT** for Digital Twin, **RP** for Replication, **DSS** for Decision Support System, and **PM** for Platform Management.

Table 1. COOPERANT-AI TOOL architecture requirements summary, organised into five functional groups.

Software Architecture: Specification and Requirements	
DI-RQ-01	Enable real-time telemetry data collection from the CSP-TES Prototype via the On-site Communication Module.
DI-RQ-02	Ensure secure, bi-directional communication between on-site hardware and the central telemetry collector.
DI-RQ-03	Supporting binding between devices distributed across different physical networks.
DI-RQ-04	Common Data Model to promote data sharing and interoperability.
DI-RQ-05	Provide adaptive stream processing.
DI-RQ-06	Ensure data integrity.
DI-RQ-07	Provide scalable distribution and orchestration of data processing (Hybrid Cloud).
DI-RQ-08	Alignment of the solution with existing standards.
DI-RQ-09	Ensure robust, flexible and secure connection in compliance with the REST API approach.
Distributed Control and Digital Twin	
DT-RQ-01	Integrate multiscale and multiphysics simulation models (e.g., COMSOL, Open Modelica) to inform decision-making.
DT-RQ-02	Maintain seamless integration of the Hybrid Digital Twin with real-time data sources for predictive diagnostics and scenario simulations.
DT-RQ-03	Provide bi-directional data flow between the Distributed Control Module, Replication Module, Decision Support System, and Distributed Control Service.
DT-RQ-04	Ensure secure control loop between the Distributed Control Service and On-site Communication Module for feedback-based actuation.
Performance Monitoring & Replication	
RP-RQ-01	Include a Replication Module capable of adapting to different sectors with varying energy requirements.

RP-RQ-02	Implement a flexible MDO algorithm integrating physical equations and sector-specific KPIs.
RP-RQ-03	Incorporate sustainability and economic assessments into the MDO process.
RP-RQ-04	Support AI-driven predictive analytics and optimisation through integration with AI techniques.
Decision Support System	
DSS-RQ-01	Enable KPI calculation and reporting.
DSS-RQ-02	Support real-time and predictive decision-making using AI algorithms.
DSS-RQ-03	Deliver intuitive and user-friendly interfaces to end-users for system monitoring and control.
DSS-RQ-04	Facilitate bi-directional communication with the Replication and Distributed Control modules for feedback adaptation.
Platform Management & Orchestration	
PM-RQ-01	Include a Platform Management Service to orchestrate application deployment, data routing, and module communication.
PM-RQ-02	Guarantee scalable deployment across pilot sites and support TRL4 validation.
PM-RQ-03	Deliver intuitive and user-friendly interfaces to end-users for system monitoring and control.
PM-RQ-04	Facilitate bi-directional communication with the Replication and Distributed Control modules for feedback adaptation.

3.2 COOPERANT-AI TOOL Architecture Design

Based on the requirements listed in the previous section, a **high-level reference architecture** design for the **COOPERANT-AI TOOL** is presented in Figure 9. This architecture intends to outline the main building blocks and functional components of the COOPERANT-AI TOOL platform in the current design phase while abstracting the implementation details. This architecture is aligned with the reference architectures detailed in Subsection 3.4 and extended to support the specific needs of COOPERANT.

The architecture is composed of three vertical tiers, from left to right: 1) the **operational environment and data acquisition tier**, which contains the physical and virtual components where the COOPERANT-AI TOOL platform operates; 2) the **interoperability and control coordination tier**, which contains the link between the CSP-TES Prototype and the platform itself, governed by an edge security layer and the semantic model for data interoperability; and 3) the **COOPERANT-AI TOOL**, which implements the data management platform, together with the internal services, the AI CORE and the Decision Support System for the users and applications to interact with it.

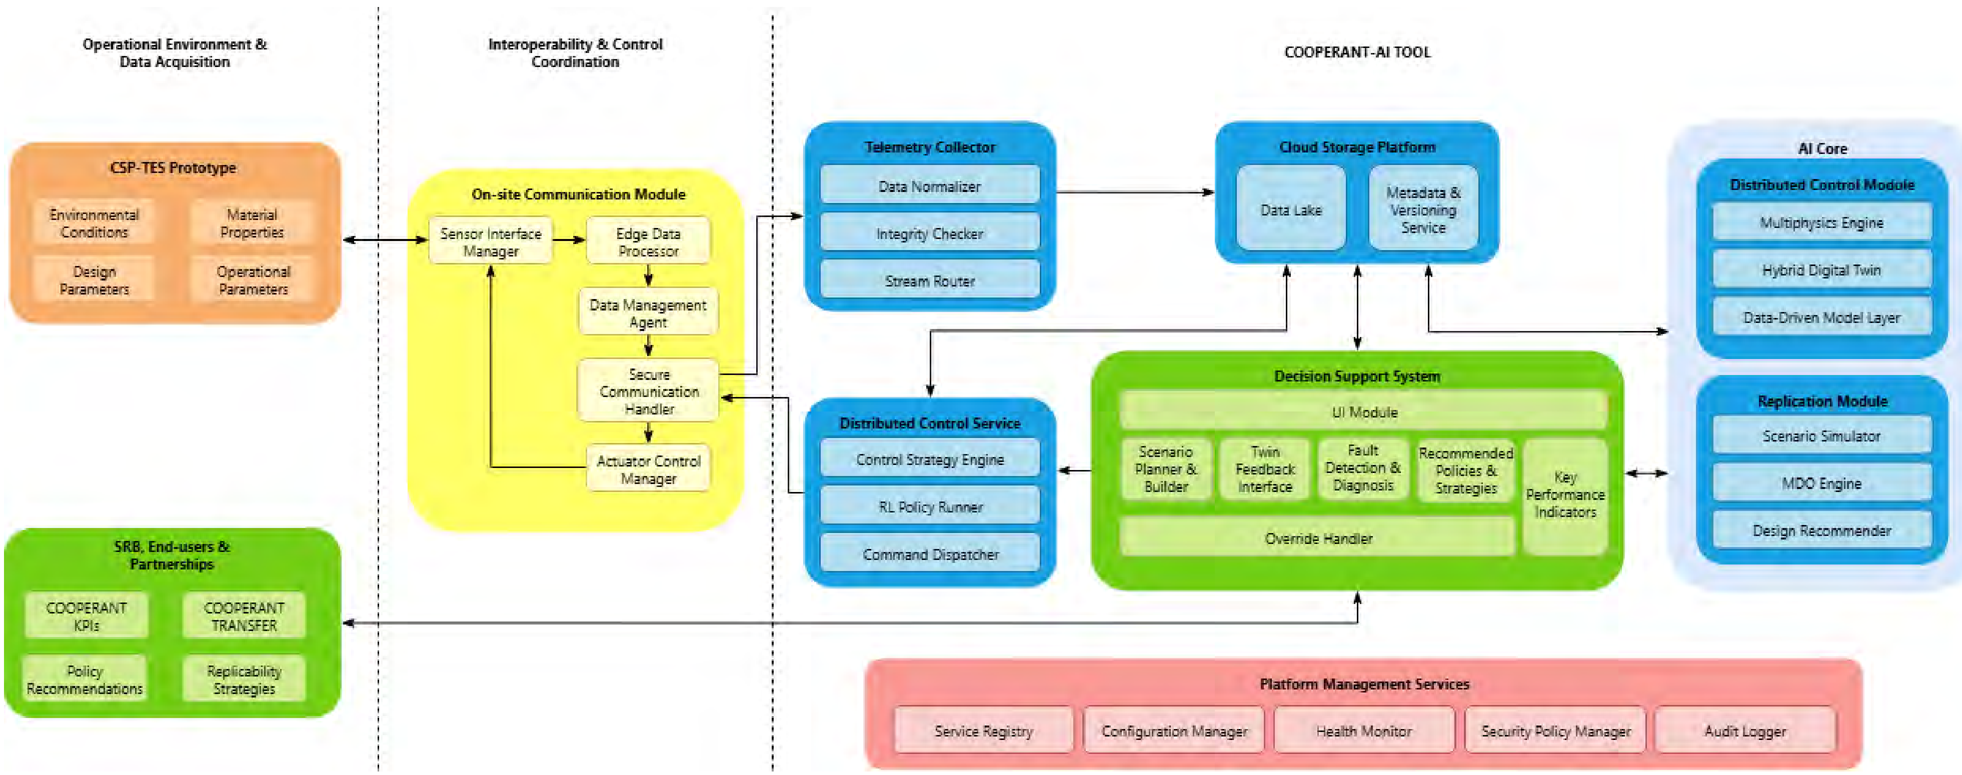

Figure 9. COOPERANT-AI TOOL software architecture design.

3.2.1 Operational Environment & Data Acquisition

This foundational tier serves as the first critical touchpoint in the COOPERANT platform, ensuring the rigorous and trustworthy generation, capture, preprocessing, and transmission of experimental and operational data from physical testbeds.

3.2.1.1 COOPERANT CSP-TES Prototype

The **CSP-TES prototype** is an integrated system including a solar thermal energy collector, hybrid packed-bed thermal storage, thermal fluid circuits, and the associated mechanical components. Several high-fidelity sensors embedded within these components continuously monitor a diverse array of performance indicators, including temperature gradients, pressure profiles, flow regimes, phase-change behaviours, and material degradation under various ambient, loading, and control circumstances.

Localised **edge computing modules ingest field data** directly and act as **preprocessing nodes for standardising, formatting, and finding anomalies in the data**. These edge modules provide translation services that convert device-specific protocols and payloads into a unified data format that follows COOPERANT's canonical ontology schema. Each communication is signed with a cryptographic signature and a timestamp, then sent through the **RESTful API service** that uses multi-factor authentication, including token renewal mechanisms, and transport-layer encryption to keep the message private and prove its source.

To ensure high resilience in off-grid or intermittently connected environments, each edge device includes a ring-buffer storage system capable of caching significant volumes of unsent data until acknowledgement is received from central repositories. Concurrent background services, such as drift monitors, communication watchdogs, and data quality agents, operate on-device to maintain real-time evaluations of sensor fidelity, uptime, and calibration drift. Any deviations beyond predefined tolerances automatically trigger upstream alerts, flagging the decision support system and initiating validation workflow.

3.2.1.2 Stakeholders, End-users & Partnerships

This operational tier also engages directly with **Stakeholder Replicability Board (SRB)** end users, industrial partners, and regional stakeholders, providing them with **high-level summaries** such as operational KPIs, energy yield trends, system health overviews, outputs from policy assessment modules, and COOPERANT TRANSFER reports. These data flows initiate the COOPERANT intelligence lifecycle and establish the factual backbone for downstream inference, simulation, and control decisions.

3.2.2 Interoperability & Control Coordination

3.2.2.1 On-Site Communication Module

This tier serves as the bridge between **field data acquisition and advanced analytical services**. The On-site communication module ensures seamless protocol interoperability, semantic alignment, and secure bidirectional communication between the physical and digital subsystems. It incorporates modular communication stacks organised into composable services hosted on-site within hardened edge enclosures. These include the **Sensor Interface Manager** for field-bus parsing, the **Edge Data Processor** for real-time normalisation, the **Secure Communication Handler** for authenticated transport, and the **Actuator Control Manager** responsible for localised execution and command reconciliation.

Incoming telemetry, regardless of source or format, is passed through a semantic translation engine, aligning each data stream with the COOPERANT Common Data Model using a combination of W3C Thing Descriptions [5] and linked-data ontologies. This **semantic layer performs mapping and validation** against expected schema definitions, enabling uniform processing of heterogeneous signals originating from multi-vendor equipment, legacy infrastructure, or novel research devices. Once semantically aligned, telemetry is classified and routed based on policy-defined logic to appropriate microservice endpoints or event-processing pipelines.

Control operations are treated with equal precision. **Commands originating from AI-driven modules or human-in-the-loop interactions** are packaged as REST-compliant, idempotent operations, each carrying atomic identifiers and digital execution tokens. This guarantees single execution semantics and allows downstream systems to reject duplicates, stale orders, or conflicting directives. Control signals are delivered across authenticated transport channels, with integrated confirmation loops guaranteeing the receipt of responses and the execution within operational boundaries. Actions are recorded and coupled to the originating event, thus enabling end-to-end traceability for audit.

This tier is crucial to accomplish the **objective of safe and transparent autonomy** within the COOPERANT-AI TOOL. It achieves this by ensuring that data and control signals are semantically consistent, secure, and operationally validated, thereby guaranteeing the integrity and quality of the information lifecycle, from sensor to decision to action execution.

3.2.3 COOPERANT-AI TOOL Services, Platforms and Modules

[5] *World Wide Web Consortium (W3C) (2023); Web of Things (WoT) Thing Description*. https://www.w3.org/TR/wot-thing-description11/

The **COOPERANT-AI TOOL** tier architecture embodies the **integrated intelligence engine** of the system, combining advanced data processing, cloud-native analytics, reinforcement learning, simulation, and decision support into a cohesive and highly modular framework. It groups a collection of interconnected **microservices** deployed across **hybrid edge-cloud environments**, allowing scalable, distributed execution of analytics and control routines.

In addition, the COOPERANT-AI TOOL is explicitly designed to integrate with existing systems and workflows by adhering to the semantic principles and interoperability standards established in the ontology (see Section 2). All services expose REST-based interfaces aligned with the ontology's canonical schemas, allowing data and insights generated within COOPERANT to be seamlessly exchanged with legacy monitoring tools, enterprise platforms, and partner infrastructures without requiring disruptive transformation.

3.2.3.1 Telemetry Collector

At the entry point, the **Telemetry Collector** serves as the **ingestion nexus** for both streaming and batch telemetry from field assets. Incoming data is normalised into platform-standard formats, verified for syntactic and semantic integrity, and then routed to the **Cloud Storage Platform**. This module interfaces tightly with the **Secure Communication Handler** at the edge layer, enforcing message schemas and source validation.

3.2.3.2 Cloud Storage Platform

The **Cloud Storage Platform** houses **structured data lakes** and rich **metadata registries** for modelling. Versioning services track lineage across raw, processed, and model-consumed datasets, supporting reproducibility and scenario traceability. This data repository acts as the source of data for the entire AI workflow, supporting both exploratory and operational analytics.

3.2.3.3 Distributed Control Module

The **AI-Core** is the analytical powerhouse of the COOPERANT-AI TOOL, grouping the **Distributed Control Module and Replication Module** under a unified orchestration layer. The Distributed Control Module integrates a robust **multiphysics simulation engine**, a **hybrid digital twin** environment, and **machine learning pipelines** for both supervised and reinforcement learning agents. Together, these components enable real-time predictions, fault forecasting, and optimised control signal synthesis under dynamic conditions.

3.2.3.4 Replication Module

In parallel, the **Replication Module** models the broader ecosystem. This module hosts scenario simulators, **multi-objective design optimisers (MDO engine)**, and a recommendation engine that explores replicability options, regional adaptation strategies, and design alternatives under evolving regulatory or environmental constraints.

3.2.3.5 Decision Support System

Insights produced by the AI Core are presented via the **Decision Support System** (**DSS**), a user-oriented platform designed to empower operators, analysts, and policymakers. The DSS integrates visual dashboards, KPI analysis, scenario planners, override interfaces, policy evaluators, and COOPERANT TRANSFER reports into a unified interface. Users can explore what-if scenarios, approve or reject AI-driven suggestions, and inject domain expertise into strategic decision-making processes. The fault detection, automated diagnosis, and policy recommendation subsystems enable contextualised and timely responses to a wide range of emerging conditions.

3.2.3.6 Distributed Control Service

Control signals validated by the **DSS** are sent to the **Distributed Control Service**, which enforces operational envelopes, safety thresholds, and AI policy compliance. This service transmits commands to the Interoperability layer via robust message brokers, guaranteeing reliable and secure field execution.

3.2.3.7 Platform Management Services

The **Platform Management Services module** allows **meta-governance** for the entire **digital ecosystem,** spanning responsibilities such as service discovery, configuration consistency, security policy enforcement, audit trails, node health monitoring, and lifecycle automation. This module supports horizontal scaling, effective orchestration, and continuous integration practices, ensuring the COOPERANT-AI TOOL remains adaptable and production-ready over its lifespan. Tools dedicated to the coordination and orchestration of all the platform elements, as well as the security components, are also integrated within this block.

3.3 COOPERANT-AI TOOL Data Architecture

COOPERANT-AI TOOL's data architecture fundamentally balances rapid operational choices with comprehensive retrospective analyses. All raw device output and unstructured log entries are first sent to a write-once **data lake**, ensuring an immutable record for further scientific validation. Subsequently, overnight transformation processes integrate **ERP and MES records** with refined sensor data, providing a high-capacity storage optimised for complex queries. A centralised registry monitors schemas and information, with straightforward HTTP endpoints for publishing new formats, reverting to prior versions, and tracking data flows. **Real-time ETL workers** consistently augment event-bus traffic, while planned batch operations address any deficiencies from offline sources. Ultimately, an audit-grade catalogue contains the full graph of transformations, enabling compliance teams to scrutinise the exact evolution of each dataset.

Current practices relating to data platforms can be classified into two main groups: those providing the data architecture with all the elements installed but without any configuration, and those providing services using the data architecture. Within the first group, it is possible to highlight companies that offer this service, such as Cloudera, AWS, IBM, Microsoft Azure, Oracle, etc. These companies provide the necessary components to create a data architecture, but it is up to the company's data engineers to configure it to provide the required service. Most of the available solutions are cloud-based, and their services are not free, except for a limited number of them (e.g. Cloudera Hortonworks Data Platform (HDP)).

Concerning the second group, fewer companies provide these services, as it depends on the desired use case. One can highlight the data architecture provided by OSIsoft, which allows the collection, enhancement, and delivery of time-series data in real-time from different devices (OSIsoft | Operational Intelligence | PI System). The use of this data platform is not free, and the price is calculated based on the number of devices or services required. It also allows integration with SAP HANA; however, this integration is an additional expense.

3.3.1 Data Management Architecture Requirements

A data architecture is a key element in developing a digital twin-oriented platform. Data platforms consist of models, policies, rules, or standards that govern what data is collected and how it is stored, processed, integrated, and used in data systems and organisations. Additionally, the data management architecture must support the ingestion, processing, and storage of large amounts of data, known as Big Data. The main characteristics of Big Data are commonly referred to as the four Vs:

- **Variety** refers to the diversity of types of ingested data, which can be structured, unstructured, or a combination of both.
- **Velocity** is the speed at which data is generated, ingested, processed, and used.
- **Volume** of data corresponds to the size of data that will be analysed, stored, and processed.
- **Veracity** of data reflects data's trustworthiness and reliability. Data comes from a wide range of different sources, and it is essential to understand the chain of custody, the metadata, and the context in which the data was collected.

The **requirements for the COOPERANT-AI TOOL data architecture are based on the Lambda-Architecture** concept [6]. The architecture must be able to ingest and process any type of data, both structured or unstructured (or some combination of both). The data is generated from

[6] Databricks *(2025) Lambda Architecture Basics*
https://www.databricks.com/glossary/lambda-architecture

multiple sources, so the architecture must be able to support the connectivity of on-site devices by using messaging protocols (e.g., MQTT, Modbus), ingestion from existing databases or external databases, and historical data uploaded from offline files. In addition, the architecture must be able to scale without code changes based on the amount of data ingested without the system degrading.

3.4 Reference Software Architectures and Mapping to the COOPERANT-AI TOOL

Two main reference architectures act as standardisation approaches for Industry 4.0 platforms: the **Reference Architecture Model for Industry 4.0 (RAMI 4.0)** and the **Industrial Internet Reference Architecture (IIRA)**.

The RAMI 4.0 architecture can be seen as a three-dimensional space whose axes represent (i) the value chain and lifecycle, (ii) the hierarchies of a production system, and (iii) the layers that describe the physical asset, the integration of software and hardware components, communication capabilities, information creation through data, functional properties, and business processes.

The IIRA architecture model presents five functional domains mapped against system characteristics and cross-cutting functions. These two reference architectures are highly correlated with each other, being complementary more than conflicting. The IIRA focuses on applicability and interoperability, while RAMI 4.0 is more focused on the digitalisation of Smart Manufacturing.

Both reference architectures were designed to support standardised methodologies and help developers in operational domains understand the most important concerns related to architectural requirements and create interoperable products across the industry.

The COOPERANT-AI TOOL is underpinned by the core data architecture, on which the COOPERANT tools, services, and applications are built. The result is a flexible, functional architecture whose main purpose is to provide a digital thread to support the development of digital twin-driven models and tools focused on energy and environmental impact assessment. The mapping of the COOPERANT-AI TOOL architecture with the RAMI 4.0 and IIRA reference models is based on the following considerations:

- Real-time monitoring data from sensors and actuators that the platform interacts with correspond to the **asset layer** in the RAMI 4.0.

- The **communication and integration layers** of RAMI 4.0 provide communication standards for data, services, and control commands that link the physical domain with the digital capability. In the IIRA, the connectivity function refers to standards such as OneM2M [7]. In COOPERANT, this maps to the **secure edge integration** and **interoperability manager** and the **semantic model manager**.
- The **information layer** in RAMI 4.0 refers to the services and data used, generated, or modified through the asset. The IIRA has a dedicated data management function. This layer contains the **data consultation and subscription manager**.
- The **functional layer** in RAMI 4.0 is dedicated to logical functions and corresponds to the analytics function in IIRA. In the COOPERANT-AI TOOL, these functional and analytical capabilities are mainly addressed by the **AI CORE**, together with the internal data platform services dedicated to data analytics and processing.
- The **business layer** present in RAMI 4.0 enables the creation of business-level functionalities and business models under specific system constraints. In IIRA, the intelligent and resilient control function responds to a similar objective. The COOPERANT-AI Tool will address this need through the **Platform Management Services** (which allow for the control and orchestration of the whole tool) and the **Decision Support System.** With a view to extensibility, an **application deployment service** is used to enable COOPERANT developers to deploy their own services on top of the core data architecture and perform business-level analysis and operations, while enhancing the core data services with both general and use case-specific applications. Furthermore, this business layer includes additional functionalities that are external to the data platform.

As demonstrated above, the COOPERANT platform architecture has a close relationship with both the IIRA and RAMI 4.0 reference models.

The COOPERANT-AI TOOL builds upon the developments from previous experience on international projects, namely the **ISLANDER project** [8], in addition to these industrial reference models. This project proposes smart grid solutions that combine renewable energy production with electricity storage and district heating in real-life demonstrations. The knowledge directly

[7] *ETSI* (2020) *SmartM2M; Smart Applications; Reference Ontology and oneM2M Mapping*. https://www.etsi.org/deliver/etsi_ts/103200_103299/103264/03.01.01_60/ts_103264v030101p.pdf

[8] ISLANDER (2025)*; High level view of the overall system and components*. https://islander-project.eu/publications/

supports the development and validation of the COOPERANT-AI TOOL. For this reason, the COOPERANT-AI TOOL will take these architectures as a reference for the initial design phase of the COOPERANT-AI TOOL architecture to be subsequently committed.

4 Data Management Plan Compliance

The digitalization ontology serves as a semantic integrator between components (e.g., digital twins, simulation models, and AI algorithms) and provides a framework for classifying, labelling, and linking all project data. It ensures seamless **interoperability, transparency,** and **long-term usability** of data assets and enables **intelligent automation** across the COOPERANT AI-TOOL architecture. These capabilities highlight the need for robust and coherent data management practices throughout the project.

4.1 Data Acquisition, Storage, and Management Procedures

This section presents the integration of data management procedures within the digitalisation ontology and architecture of the COOPERANT-AI TOOL, in alignment with the outcomes of **Task T1.2** as documented in **Deliverable D1.2** – **Preliminary Data Management Plan** and **Open Science** Approach. The foundation of responsible data governance is the Data Management Plan (DMP), which guarantees safe, moral, and efficient data handling throughout the whole data lifecycle. The chosen methodology supports the project's **Open Science and sustainability strategy** and is completely compliant with the Horizon Europe DMP template and the FAIR (Findable, Accessible, Interoperable, and Reusable) principles. The procedures for data management in the COOPERANT-AI TOOL are structured to support secure, efficient, and standardised interactions between the physical and digital components of materials and system development. These procedures build upon the following principles:

- **Data Acquisition**: Data flows originate from heterogeneous sources, including experimental tests, process simulations, monitoring infrastructures (e.g., sensors), and AI-inferred outputs. All data are semantically tagged based on the COOPERANT ontology and enriched with standardised metadata in line with the **COOPERANT Dataset Template**. Dataset types include observational, simulation, experimental, and derived data as described in **D1.2.**

- **Data Storage**: Data is stored within a centrally administered internal repository hosted on **Microsoft SharePoint**, with strict access control, versioning, and update logging. Ontology-based metadata ensures full traceability and enables query-based retrieval and filtering.

Zenodo uses an OpenAIRE-compliant platform to guarantee long-term preservation and, where appropriate, public access. To help the scientific community with discoverability and reuse, it provides archival storage with persistent identifiers (DOIs), adheres to the FAIR data principles, and maintains standardized metadata and version control. The **COOPERANT project community** on Zenodo serves as a central hub for depositing the final versions of scientific publications and is also utilised for **project documentation** and **data storage** (see Figure 10).

https://zenodo.org/communities/cooperant

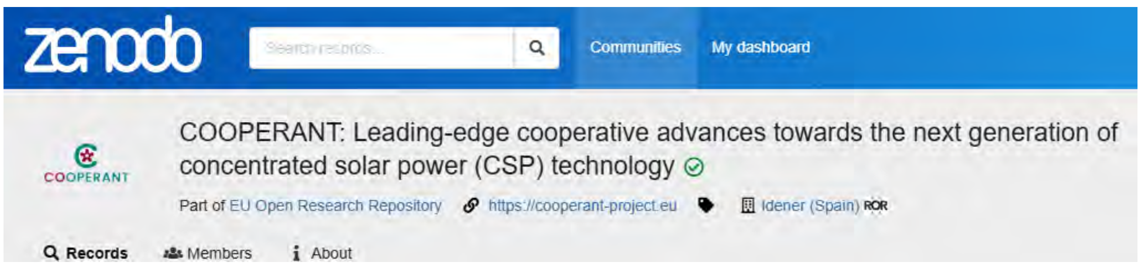

Figure 10. COOPERANT project community on Zenodo.

- **Data Management and Validation**: A collaborative validation process has been established, where partners are responsible for completing the dataset template and ensuring consistency with DMP standards. The digitalisation ontology provides automated support for identifying incomplete records, non-conforming formats, and semantic misalignments. This ensures robustness, quality control, and harmony across WPs.
- **Interoperability and Integration**: Data integration with external tools, platforms, and research infrastructures is made possible by the ontology's alignment with external and domain-relevant standards and vocabularies (such as QUDT, SOSA/SSN, FOAF, and RAMI 4.0). Federated knowledge graphs can be expanded more easily, and semantic query capabilities are further supported by linked data principles.

4.2 Database Alignment with Data Management Plan Principles

The data infrastructure of the COOPERANT-AI TOOL will comply with the guidelines and standards outlined in the COOPERANT DMP, guaranteeing:

- **FAIR Data Compliance**:
    - **Findability**: Findability: Distinct identifiers, comprehensive metadata, and indexing in searchable repositories facilitate data retrieval. This facilitates access, sharing, and reutilization. The COOPERANT project ensures that all datasets possess permanent identifiers (such as DOIs) and are characterised using standardised metadata schemas, facilitating discovery by both individuals and machines.
    - **Accessibility**: All individuals should have maximal access to data, while adhering to privacy, security, and legal regulations such as the General Data Protection Regulation (GDPR). All data is stored on platforms adhering to open standards, providing user manuals, and clearly delineating data access procedures. The definitive versions of scientific papers will be stored in a repository compliant with **OpenAIRE**, such as **Zenodo**, to facilitate **Open Access to Data and Research (OADR)** and enhance long-term accessibility strategies.
    - **Interoperability**: Data is structured utilising standardised protocols, ontologies, and formats that facilitate integration with other datasets and systems. COOPERANT employs domain-specific lexicons, such as those from the European Open Science Cloud, alongside standardised protocols to ensure interoperability among various components and platforms. This enhances semantic clarity and facilitates data readability and exchange for machines.
    - **Reusability**: COOPERANT's data is reusable thanks to thorough documentation, provenance records, and explicit usage guidelines. To maximise legal clarity and promote data reuse in upcoming studies, industrial applications, and cross-sector collaboration, standardised open licences are used. Every dataset is carefully selected to guarantee traceability and long-term preservation.
- **Legal and Ethical Compliance**: GDPR compliance will be maintained using anonymisation techniques, data encryption, and secure transfer protocols. Data ownership and IPR will be managed through the existing Consortium Agreement.

- **Open Science Strategy**: All scientific outputs, including papers, software, and datasets, will be accessible on repositories such as GitHub or Zenodo under open licenses. This approach facilitates public access to the project's outcomes, enables reproduction of these results, and promotes transparency regarding their nature.
- **Sustainability and Updating Mechanism**: To account for any modifications to the dataset landscape, the DMP will be updated midway through the project in **D1.5 (M30)** and again at the end of the project in **D1.6 (M48)**. All pertinent stakeholders will be informed of these updates, which will be coordinated with WP1. This dynamism will be supported by the semantic data structures created in this task, which allow for consistent data mapping even when new data types or formats appear.

5 Conclusions

This deliverable has successfully fulfilled the objectives of **Subtask 3.5.1: Definition of digitalisation ontology**, providing a comprehensive and detailed definition of the digitalisation ontology, as well as establishing the architecture of the COOPERANT-AI TOOL. Fully aligned with the data acquisition, storage, and management requirements specified in Task T1.2 (as detailed in Deliverable D1.2), this work marks a significant milestone in **aligning the physical and digital development of materials and systems.**

To ensure a well-grounded approach, the task began with a comprehensive state-of-the-art review of current digitalisation frameworks and semantic modelling practices. Building on these insights, a **graph-based ontology** was developed and implemented using **Protégé**, a widely used and open-source ontology editing platform. The ontology was enriched with advanced visualisation capabilities using **RDFlib and PyVis Python libraries**. These tools enabled the dynamic generation of intuitive, interactive graphs with colour-coded, context-aware nodes that significantly enhance the description and interpretability of complex relationships within the system.

The **COOPERANT-AI TOOL** ontology is logically **structured into** three main sections**: Inputs, Processing modules, and Outputs**. Each section has comprehensive modules and subcomponents with their respective properties and interrelations. This architecture ensures consistent alignment between physical assets and digital components, facilitating a detailed and comprehensive representation of the system's functionality.

The digital ontology functions not only as a **descriptive model** but also as the **semantic foundation** of the COOPERANT-AI TOOL. It promotes interoperability across various modules, supports automated workflows, and enhances data traceability and reuse. Consequently, it plays a vital role in enabling intelligent automation and semantic integration within the COOPERANT-AI TOOL.

The overall architecture, created concurrently with the ontology, supports this **semantic basis** by outlining the structural and functional relationships among system components. This architecture offers a **unified framework** for the integration of distributed control components, data repositories, artificial intelligence engines, and user interfaces. **Interoperability** and **scalability** are integrated into the tool using the **REST API approach** and the concepts of IIRA, RAMI 4.0, and the experience gained from IDENER in interoperability projects and architectures.

This deliverable also integrates **data management, acquisition,** and **storage procedures** with the digital ontology framework. These protocols, which were developed in compliance with the project's **Data Management Plan (DMP)**, closely follow the **FAIR principles** (Findable, Accessible, Interoperable, and Reusable). This ensures effective, ethical, and safe data handling throughout the project lifecycle and strengthens the project's commitment to **Open Science**.

Importantly, **Subtask 3.5.2: Databases and Modules Coupling** and other upcoming tasks will be directly supported by the ontology and architecture created here. Within the COOPERANT platform, a variety of datasets, including those from simulations, experiments, external sources, and literature, can be consistently classified, integrated, and exchanged thanks to the deliverable's structured semantic layer and precise interface definitions.

By providing a clear semantic structure, the COOPERANT ontology digitalisation has improved information consistency across partners, facilitated collaboration, and accelerated the configuration of the COOPERANT-AI TOOL modules. Its development also adheres to semantic modelling best practices, which obliquely supports European initiatives for interoperability and data standardisation. This strategy fortifies the project's methodological rigour and offers a reusable foundation that can stimulate related research and industry endeavours.

In conclusion, this deliverable effectively satisfies the goals of Subtask 3.5.1 by providing a thorough and cutting-edge digital framework that serves as the foundation for the COOPERANT-AI TOOL's intelligent, modular, and interoperable design. All later technical advancements are made possible by this foundational work, which ensures coherence, scalability, and ongoing innovation throughout the project.

Annexe I

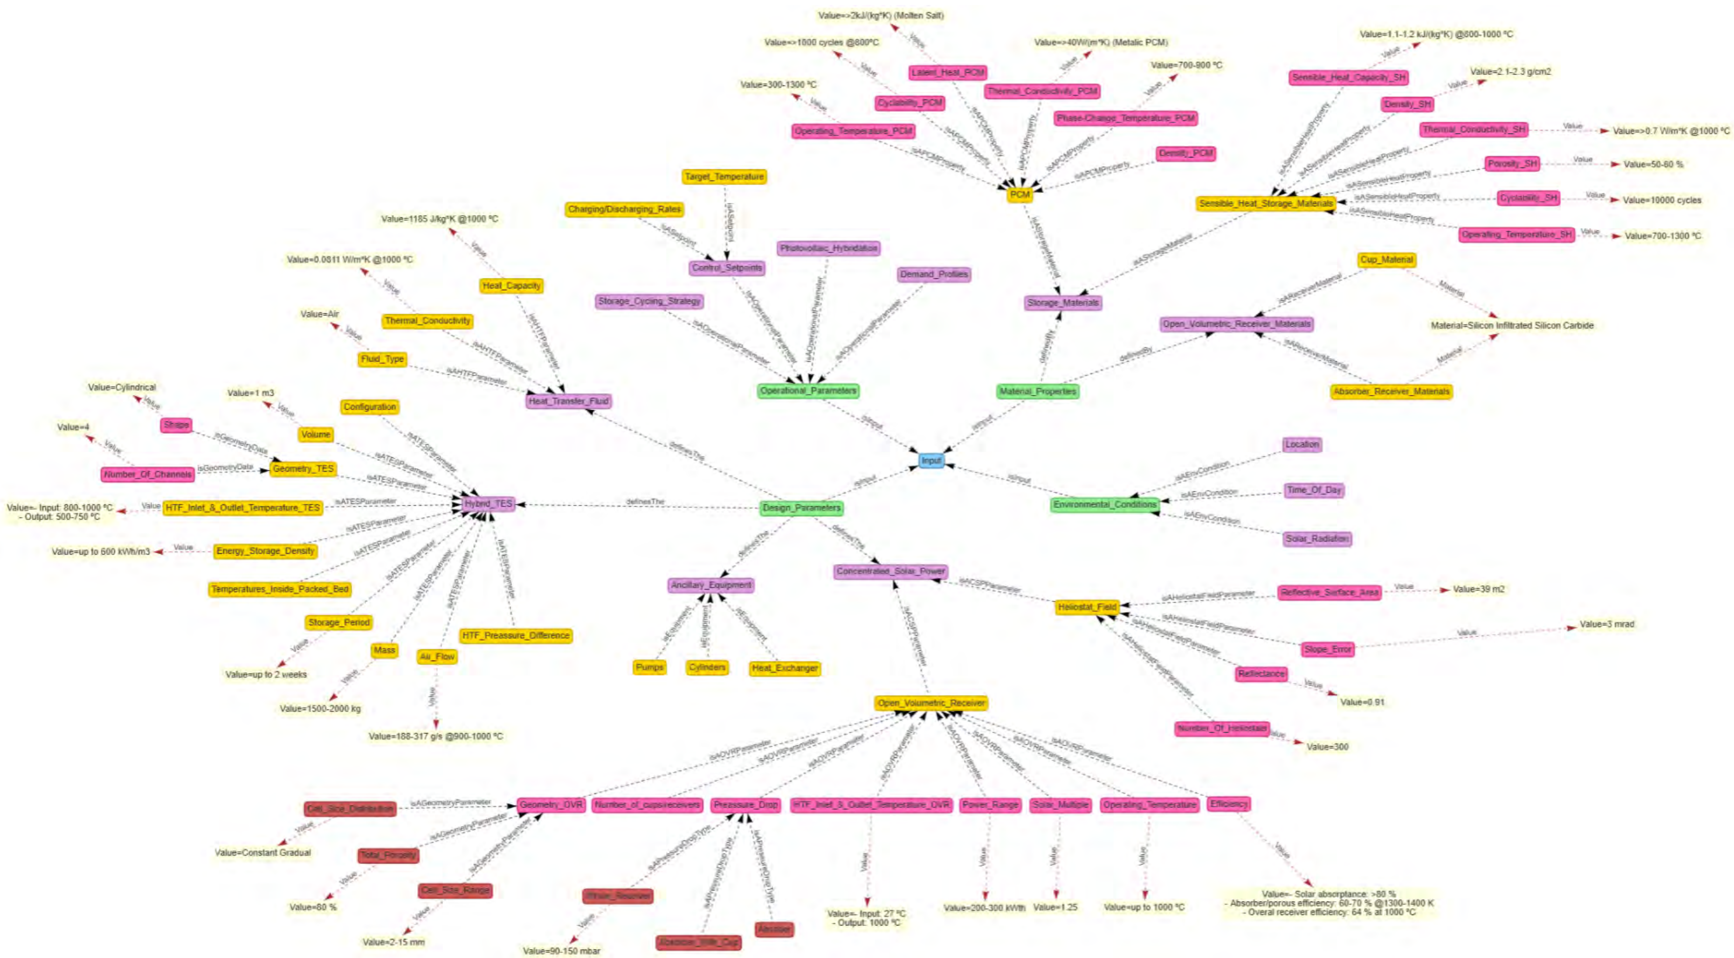

Figure 11. Visualisation of the ontology inputs.