

# HACID - Deliverable

# Final Exploitation Plan

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101070588. UK Research and Innovation (UKRI) funds the Nesta and Met Office contributions to the HACID project.

<b>Deliverable number:</b>	D8.8
<b>Due date:</b>	28/02/2026
<b>Nature<sup>1</sup>:</b>	R
<b>Dissemination Level<sup>2</sup>:</b>	PU
<b>Work Package:</b>	WP8
<b>Lead Beneficiary:</b>	CNR
<b>Contributing Beneficiaries:</b>	all

---

<sup>1</sup> The following codes are admitted:

- R: Document, report (excluding the periodic and final reports)
- DEM: Demonstrator, pilot, prototype, plan designs
- DEC: Websites, patents filing, press & media actions, videos, etc.
- DATA: Data sets, microdata, etc.
- DMP: Data management plan
- ETHICS: Deliverables related to ethics issues.
- SECURITY: Deliverables related to security issues
- OTHER: Software, technical diagram, algorithms, models, etc.

<sup>2</sup> The following codes are admitted:

- PU – Public, fully open, e.g. web (Deliverables flagged as public will be automatically published in CORDIS project's page)
- SEN – Sensitive, limited under the conditions of the Grant Agreement
- Classified R-UE/EU-R – EU RESTRICTED under the Commission Decision No2015/444
- Classified C-UE/EU-C – EU CONFIDENTIAL under the Commission Decision No2015/444
- Classified S-UE/EU-S – EU SECRET under the Commission Decision No2015/444

# Document History

<b>Version</b>	<b>Date</b>	<b>Description</b>	<b>Author</b>	<b>Partner</b>
V1	30/03/2025	Creation of draft document	Vito Trianni	CNR
V2	28/10/2025	First draft of plans for MS	Alessandro Russo	CNR
V3	25/11/2025	Strategy for MD	Stefan Herzog	MPG
V4	09/01/2026	Definition of plans for foreground knowledge in MD and CS	Vito Trianni	CNR
V5	13/02/2026	Finalisation of plans for MD	Stefan Herzog	MPG
V6	27/01/2026	Definition of plan for other services in CS	Neha Mittal	Met Office
V7	28/01/2026	Finalisation of the document	Vito Trianni	CNR
V8	28/01/2026	Quality check	Ralf Kurvers	MPG
V9	10/02/2026	Quality check	Fiorela Ciroku	CNR
V10	12/02/2026	Final version	Vito Trianni	CNR

# Table of content

<b>Document History</b>	<b>2</b>
<b>Table of content</b>	<b>3</b>
<b>1. Introduction</b>	<b>4</b>
<b>2. Plan for exploitation in the healthcare domain</b>	<b>4</b>
2.1. Backend algorithms for knowledge analysis and aggregation	4
2.1.1. Advanced aggregation algorithms based on the DKG and on LLMs	4
2.1.2. Advanced knowledge-extraction methods for generation of structured knowledge from medical text	5
2.2. Frontend features for elicitation of more accurate diagnoses	6
2.2.1. Exposure to diagnoses generated by other teams	6
2.2.2. Integrated chat with dedicated chatbots	7
2.3. Concluding Remarks	7
<b>3. Plan for exploitation in the climate service domain</b>	<b>7</b>
3.1. Plan for exploitation of foreground knowledge	7
3.1.1. DKG for Climate Services	8
3.1.2. Interactive Visualisation Tool for the DKG	9
3.1.3. HACID Decision Support System (HACID-DSS)	10
3.1.4. Concluding Remarks	10
3.2. Plan for expansion towards new services	11
<b>4. Plan for exploitation in other domains</b>	<b>12</b>
4.1. Community-Led Energy Systems	13
4.2. Horizon Scanning for Water Regulation	14
4.3. Asylum determination management	15
4.4. Concluding remarks	16

# 1. Introduction

Within the HACID project, exploitation will follow two different paths: a) direct exploitation of the project results, and b) expansion of the HACID technology to new services. On the one hand, we plan to exploit foreground knowledge, identifying results of the project that can be directly exploited. On the other hand, the HACID technology can be expanded to offer new services, and plans to explore these opportunities are discussed here.

Exploitation will take place not only within the application domains currently targeted, namely services for healthcare and for climate change adaptation management, but also towards new application domains that have been identified through activities with external stakeholders. In the following, we provide the current exploitation plan grouping the activities according to the target application domains. This plan will be executed, in full or in part, subject to the availability of resources to bring forward the mentioned activities.

## 2. Plan for exploitation in the healthcare domain

We have identified foreground knowledge that can find space for exploitation within the medical diagnostics domain, with particular reference to the possible integration within [the Human Dx mobile application](#). In this document, we discuss the backend algorithms produced to analyse data and aggregate solutions into a collective diagnosis, as supported by the developed medical diagnostics domain knowledge graph (DKG), as well as frontend features that have been tested during experimental activities within WP4.

### 2.1. Backend algorithms for knowledge analysis and aggregation

With backend algorithms, we refer to software and methods that run in the backend of the HACID Decision Support System (DSS). We can distinguish between Natural Language Processing (NLP) algorithms for information extraction from text, specifically for concepts related to the medical domain, and algorithms for aggregation of diagnoses produced by a diagnosing agent (human or artificial). All developments resulting from scientific research have been or will be published and described in open access academic papers, as stated in the Data Management Plan.

#### 2.1.1. Advanced aggregation algorithms based on the DKG and on LLMs

**Context of the Exploitation Activity** — Collective medical diagnosis relies on combining multiple diagnostic hypotheses generated by heterogeneous agents, including clinicians with varying expertise and AI systems. Current aggregation approaches typically rely on frequency-based or performance-weighted voting, which fail to fully exploit semantic relationships between diagnoses and contextual medical knowledge. Within HACID, advanced aggregation algorithms have been developed that leverage the medical diagnostics DKG and, optionally, Large Language Models (LLMs). These methods, studied within WP4, enable deeper semantic alignment of diagnostic hypotheses, integration of

metacognitive signals from users, and enrichment of diagnoses with contextual clinical knowledge. This makes them particularly suitable for integration into an existing collective intelligence platform such as the Human Dx mobile application.

**Description of the Exploitation Activity** — The exploitation activity consists of integrating HACID aggregation algorithms into the backend of the medical diagnostic decision support platform. This includes:

- Deployment of DKG-based semantic aggregation methods to identify similarity and support among diagnostic hypotheses.
- Integration of LLM-based enrichment methods to contextualise and expand user-provided diagnoses.
- Configuration of multiple aggregation strategies, including methods based on expert similarity, user performance, and metacognitive elicitation.
- Implementation of user- or platform-level configuration options to enable or disable specific aggregation methods (e.g. LLM-based approaches).
- User acceptance testing and refinement of result restitution to ensure clarity, transparency, and trust.

All algorithms remain openly published, while exploitation focuses on operational integration, configuration, and validation.

#### **Roadmap for Implementation**

- Months 0–3: Integration of aggregation modules into the backend infrastructure; alignment with existing diagnostic workflows.
- Months 4–6: Configuration and parameter tuning of aggregation strategies; implementation of user-configurable options.
- Months 7–9: User acceptance testing with clinicians; refinement of result presentation and explanation.
- Months 10–12: Deployment in a production or pilot environment; monitoring and iterative optimisation.

### 2.1.2. Advanced knowledge-extraction methods for generation of structured knowledge from medical text

**Context of the Exploitation Activity** — Medical knowledge evolves rapidly and is predominantly documented in unstructured textual sources, such as scientific publications and clinical reports. Maintaining an up-to-date and comprehensive medical DKG is therefore a major challenge, yet critical for improving diagnostic aggregation quality. HACID has developed NLP and RAG-based methods for extracting structured knowledge from medical text and generating knowledge graphs that can be linked to the existing DKG. These methods enable continuous enrichment of domain knowledge and provide additional evidence to support more accurate and context-aware diagnostic aggregation.

**Description of the Exploitation Activity** — The exploitation activity focuses on operationalising these knowledge-extraction methods as backend services, including:

- Further testing and validation of NLP pipelines for extracting structured medical knowledge from text.
- Integration of generated knowledge graphs with the existing medical diagnostics DKG.

- Evaluation of the impact of newly extracted knowledge on diagnostic aggregation performance.
- Deployment of the knowledge-extraction modules as backend services supporting continuous DKG enrichment.

The exploitable value lies in the ability to maintain a living, automatically enriched medical knowledge base that directly improves decision support.

### **Roadmap for Implementation**

- Months 0–6: Further testing and benchmarking of knowledge-extraction methods.
- Months 7–12: Evaluation of the impact of enriched knowledge on aggregation algorithms.
- Months 13–18: Integration of extraction pipelines into the backend infrastructure.
- Months 19–24: Deployment, monitoring, and optimisation of continuous knowledge enrichment workflows.

## **2.2. Frontend features for elicitation of more accurate diagnoses**

Within WP4, several experiments have been conducted to understand how collective intelligence can be improved by means of additional information, which can be delivered by means of frontend features. Besides the elicitation of metacognitive reasoning (e.g., subjective confidence, mention of popular responses), we have focused on the provision of additional information that can lead to improved individual and collective differentials.

### **2.2.1. Exposure to diagnoses generated by other teams**

**Context of the Exploitation Activity** — Experimental results within HACID are exploring the effects on the clinicians' diagnostic performance from exposure to aggregated diagnoses generated by other teams (human-only, AI-only, or hybrid). We believe that, when carefully designed, such exposure can stimulate reflection, reduce anchoring errors, and improve collective outcomes without undermining professional autonomy. This feature is particularly relevant for a mobile application such as Human Dx, where information timing, framing, and transparency are critical.

**Description of the Exploitation Activity** — The exploitation activity consists of refining and deploying frontend features that:

- Present collective differential diagnoses generated by different team compositions.
- Display this information alongside case findings and metadata describing the contributing agents.
- Allow controlled exposure (e.g. before or after case review) to assess cognitive impact.
- Improve the user experience to ensure the feature is perceived as supportive rather than prescriptive.

### **Roadmap for Implementation**

- Months 0–3: Deeper evaluation of performance and behavioural effects of the feature.
- Months 4–6: UX redesign and refinement of information presentation.

- Months 7–9: Integration and testing within the application.
- Months 10–12: Deployment and monitoring in real-world or pilot usage.

### 2.2.2. Integrated chat with dedicated chatbots

**Context of the Exploitation Activity** — Interactive dialogue can support diagnostic reasoning by encouraging reflection, justification, and learning. HACID has explored the use of integrated chatbots configured as coaches, evaluators, or peers, each providing distinct forms of support during the diagnostic process. Such conversational interfaces are well suited to mobile diagnostic platforms and can enhance user engagement and diagnostic quality when carefully designed.

**Description of the Exploitation Activity** — The exploitation activity focuses on deploying and refining chatbot-based frontend features that:

- Enable interaction with role-specific chatbots (coach, evaluator, peer).
- Support different diagnostic needs, such as guidance, feedback, or peer-like discussion.
- Are configurable to match user preferences and expertise levels.
- Are designed to complement, not replace, clinician reasoning.

#### Roadmap for Implementation

- Months 0–3: Deeper evaluation of the impact of different chatbot roles.
- Months 4–6: UX improvements and refinement of conversational flows.
- Months 7–9: Integration and testing within the mobile application.
- Months 10–12: Deployment and monitoring of usage, satisfaction, and performance effects.

## 2.3. Concluding Remarks

By clearly separating backend and frontend exploitation activities and aligning them with realistic timelines, this plan provides a concrete and implementable pathway for transferring HACID foreground knowledge into operational medical diagnostic platforms. The combined exploitation of semantic aggregation, automated knowledge enrichment, and carefully designed interaction features strengthens hybrid collective intelligence while preserving clinician’s agency, trust, and accountability.

# 3. Plan for exploitation in the climate service domain

## 3.1. Plan for exploitation of foreground knowledge

Similarly to the medical diagnostics domain, the outputs of the project mainly consist of a DKG for climate services, an interactive dashboard for exploration of structured knowledge, and a web application that integrates these components implementing the full HACID-DSS for climate services. The climate service domain is still young and rather unstructured. The bulk component is provided by climate scientists focusing on climate modelling. In this respect, different organisations provide a diverse set of methodologies owing to individual

expertise gained over the years, often constrained in the usage of resources and models that are internally developed. By widening the range of tools available for climate services, there is an opportunity to foster the generation of information sources and methodologies that, once combined, can offer better solutions to specific cases. Additionally, linking the climate modelling aspects to the user needs does not always follow a straight and safe path. Understanding risk factors and determining fitting future climate scenarios require a deep understanding of the case specificities, which can be improved through integration of diverse perspectives.

### 3.1.1. DKG for Climate Services

**Context of the Exploitation Activity** — The HACID DKG addresses a critical gap in the climate services domain: the lack of formalised, machine-readable, and interoperable representations of knowledge related to climate change adaptation management. Published as an open-access resource, the DKG has the potential to become a reference infrastructure for climate services, supporting evidence-based decision making, interoperability among tools, and reproducibility of climate information workflows. Given its nature as a living resource, the DKG requires continuous curation, updates, and governance to remain relevant and trustworthy. This creates opportunities for service-oriented exploitation rather than exclusive ownership-based commercialisation.

**Description of the Exploitation Activity** — The exploitation strategy for the DKG focuses on knowledge infrastructure services, including:

- Provision and long-term maintenance of reliable SPARQL endpoints and APIs for programmatic access.
- Curation, extension, and quality control of the DKG to integrate new scientific evidence, methods, and climate service use cases.
- Customisation and domain adaptation of the DKG for specific users (e.g. regional climate services, sectoral adaptation planning).
- Training, consulting, and support services for organisations wishing to adopt or integrate the DKG into their own systems.

Target users include climate service providers, research institutions, public authorities, consultancies, and EU or international projects requiring structured climate adaptation knowledge.

#### **Roadmap for Implementation**

Short term (0–12 months):

- Stabilise and document the DKG schema, vocabularies, and governance model.
- Deploy and monitor robust public SPARQL endpoints and APIs.
- Produce technical documentation, tutorials, and example queries.
- Identify early adopters and pilot users within the climate services community.

Medium term (1–3 years):

- Establish a formal curation and update workflow, involving domain experts.
- Introduce tiered access or service-level agreements (SLAs) for advanced usage (e.g. high availability, custom endpoints).
- Integrate the DKG with external knowledge bases and standards (e.g. WMO, INSPIRE, GEOSS).

- Secure funding through service contracts, institutional support, or project-based contributions.

Long term (3+ years):

- Position the DKG as a reference knowledge infrastructure for climate adaptation services.
- Build a community-driven governance model with contributors and validators.
- Expand the DKG to adjacent domains (e.g. mitigation, disaster risk reduction).
- Ensure long-term sustainability through institutional anchoring or a dedicated legal entity.

### 3.1.2. Interactive Visualisation Tool for the DKG

**Context of the Exploitation Activity** — The interactive dashboard developed within HACID enables intuitive exploration, visual querying, and discovery of complex semantic knowledge embedded in the DKG. This capability is particularly valuable for climate scientists and practitioners who may not be experts in semantic technologies but need to navigate structured knowledge efficiently. Beyond climate services, the tool represents a generic, domain-agnostic solution for visual interaction with knowledge graphs, addressing a broader need across data-intensive scientific and professional domains.

**Description of the Exploitation Activity** — The exploitation of the visualisation tool focuses on software reuse and adaptation, including:

- Offering the dashboard as a hosted service connected to the HACID DKG.
- Providing a configurable and deployable version of the tool for other domains and knowledge graphs.
- Custom UI/UX adaptations for specific user groups or decision contexts.
- Training and support services for organisations adopting the tool internally.
- Integration of the tool into broader decision-support or data platforms.

Potential users include research organisations, public agencies, companies managing complex knowledge assets, and digital infrastructure providers.

#### Roadmap for Implementation

Short term (0–12 months):

- Refine usability, performance, and accessibility based on pilot feedback.
- Modularise the software to separate domain-specific and domain-independent components.
- Produce user manuals, tutorials, and demo instances.
- Deploy a stable hosted version connected to the HACID DKG.

Medium term (1–3 years):

- Enable configuration for arbitrary knowledge graphs via standard interfaces.
- Introduce advanced visual analytics and filtering capabilities.
- Offer customisation services (branding, workflows, domain views).
- Establish partnerships with projects or organisations managing large knowledge graphs.

Long term (3+ years):

- Position the tool as a reference solution for visual knowledge graph exploration.
- Develop a plugin ecosystem or extension framework.
- Integrate AI-assisted exploration and explanation features.

- Ensure sustainability through licensing, service contracts, or integration into larger platforms.

### 3.1.3. HACID Decision Support System (HACID-DSS)

**Context of the Exploitation Activity** — The HACID-DSS introduces a novel, structured methodology for designing workflows that support climate adaptation management, combining domain knowledge, human expertise, and AI components. Its emphasis on hybrid collective intelligence differentiates it from traditional decision-support tools and aligns with emerging needs for transparent, collaborative, and explainable decision processes. While demonstrated in relevant environments, the HACID-DSS requires further development and validation to reach operational maturity, particularly regarding AI integration and scalability.

**Description of the Exploitation Activity** — The exploitation strategy for the HACID-DSS is centred on advanced decision-support services, including:

- Deployment of the HACID-DSS as a web-based platform for climate service provision.
- Customisation of workflows for specific adaptation contexts (e.g. regional planning, sectoral risk assessment).
- Integration of AI components to assist workflow generation, evidence retrieval, and recommendation.
- Facilitation of expert networks and training communities around the platform.
- Use of the HACID-DSS as a backbone for participatory and collaborative decision-making processes.

Primary users include climate service providers, public authorities, consultants, and international organisations involved in adaptation planning.

#### Roadmap for Implementation

Short term (0–12 months):

Consolidate the existing prototype into a stable, documented platform.

- Conduct additional pilots with selected stakeholders and real-world use cases.
- Identify priority AI enhancements aligned with user needs.
- Develop training materials and onboarding workflows.

Medium term (1–3 years):

- Integrate AI-assisted workflow design and recommendation features.
- Expand the system to support collaborative, multi-user scenarios.
- Formalise service offerings (e.g. hosted DSS, tailored deployments).
- Build a network of experts, trainers, and early adopters around the platform.

Long term (3+ years):

- Establish the HACID-DSS as a reference framework for hybrid collective intelligence in climate services.
- Strengthen institutional uptake and long-term funding mechanisms.
- Evolve the DSS into a scalable ecosystem integrating tools, knowledge, and communities.

### 3.1.4. Concluding Remarks

Together, the DKG, the visualisation tool, and the HACID-DSS form a coherent exploitation portfolio that combines **open knowledge infrastructure**, **innovative interaction**

**paradigms**, and **advanced decision-support methodologies**. Their exploitation relies less on traditional product commercialisation and more on service provision, community building, and long-term sustainability—aligned with the strategic needs of the climate service community and beyond.

## 3.2. Plan for expansion towards new services

The user research performed within the HACID project has highlighted numerous opportunities to support climate services by means of advanced technologies based on AI and/or on hybrid collective intelligence. Expansion of the functionality offered by the HACID-DSS towards specific sectors can contribute to surfacing user needs and risk appetite, integrating additional knowledge and providing explanations about uncertainty ranges.

**Context of the Exploitation Activity** - The HACID-DSS builds on extensive user consultation supporting national climate-service development, which highlighted strong demand for decision-support tools that better link climate information to sector-specific impact pathways. Hydrology impact modellers, energy-system planners, infrastructure risk managers, and other climate-sensitive sectors all expressed the need for transparent delivery of climate data and information that can be easily interpreted. The HACID-DSS provides a structured methodology that combines domain knowledge, human expertise, and selected AI components to design workflows that guide users through climate-risk and adaptation analysis. Its hybrid-intelligence approach, integrating expert reasoning with curated evidence, positions HACID as an innovative alternative to traditional climate-service tools. While demonstrated effectively in relevant environments, the system must now undergo targeted expansion and validation to ensure its workflows are fully adapted to the realities of sectoral impact modelling and risk assessment.

**Description of the Exploitation Activity** - Deploy HACID-DSS as a web-based platform supporting structured climate-risk and adaptation workflows.

- Develop customised sector-specific workflows, starting with:
  - Hydrology and flood-risk modelling (e.g., extreme thresholds, model inputs).
  - Energy-sector planning (e.g., temperature-dependent demand, renewable variability, extreme-event stressors).
  - Infrastructure and built-environment risk assessment (e.g., heat, storm impacts, interdependencies).
- Integrate sector-relevant knowledge modules, such as operational decision thresholds, modelling assumptions, regulatory requirements, and key performance indicators
- Embed lightweight AI components to:
  - Suggest workflow steps tailored to each sector's analytical process.
  - Retrieve relevant evidence sources and recommended datasets.
  - Support structured reasoning and explainability.
- Enable knowledge-sharing and learning, including curated case studies, practical guidance, and community-driven best-practice examples.
- Support cross-sectoral collaboration (e.g., hydrology + energy + local authorities) through shared workflows and transparent decision records.

- Target primary users include climate-service providers, environmental regulators, water agencies, energy-system planners, consultants, and local authorities.

## **Roadmap for Implementation**

### **Short-Term (0–12 months): Foundations for Sector Integration**

- Consolidate the HACID-DSS prototype into a stable, documented platform.
- Engage hydrology, energy, infrastructure, and environmental-risk stakeholders to identify priority workflows and required sectoral knowledge.
- Run small pilot demonstrations with selected organisations (e.g., water utilities, energy analysts, local authorities).
- Develop introductory training materials and onboarding tailored to sector-specific use cases.

### **Medium-Term (1–3 years): Embedding Sectoral Decision Pathways**

- Integrate lightweight AI-assisted workflow design and sector-specific evidence retrieval.
- Add basic multi-user features to support small teams working across sectors or organisations.
- Develop hosted and configurable versions of HACID-DSS for particular sectors (e.g., hydrology-focused instance).
- Build a network of expert users and trainers in hydrology, energy, and infrastructure domains.

### **Long-Term (3+ years): HACID-DSS as a Cross-Sector Decision-Support Framework**

- Establish HACID-DSS as a recognised reference approach for structured, traceable, and cross-sector climate-impact assessment.
- Support institutional adoption of HACID workflows within national climate-service processes.
- Extend the DSS into a sustainable platform with modular sectoral knowledge components and a community of practitioners.

## **4. Plan for exploitation in other domains**

HACID defines a methodology for decision support that can be generalised to different application domains beyond health diagnostics and climate services. The exploitation activities carried out during the project have focused on the identification of potential domains of application, as described in the deliverable D8.6 Project Catalogue. Specifically, we highlighted two application domains of particular interest: community-led energy systems and water regulation management. For each domain, a participatory design activity has led to blueprints that highlight potential usage of the HACID technology in the specific context. Here, we propose an exploitation plan towards these two application domains.

Additionally, we have addressed another application—although without going as deep as for those presented in D8.6, that is, asylum determination management. This additional domain has been investigated through interviews with experts from the UNHCR. The interview resulted in a high potential for HACID technologies for asylum determination, because also in this case the human expertise and the AI contribution would be strongly complementary. Hence, here we provide an exploitation plan also for this additional domain.

## 4.1. Community-Led Energy Systems

**Context of the Exploitation Activity** — Energy systems are undergoing a structural transition towards decentralisation, with growing emphasis on community-led initiatives, cooperative ownership models, and local participation in decision-making. These settings are characterised by fragmented data, diverse stakeholder perspectives, and the need to balance technical optimisation with social legitimacy and trust. The HACID approach is particularly well aligned with this context because it enables combining AI-driven analysis of technical, environmental, and economic data with situated human knowledge from communities, practitioners, and local experts. The participatory design process carried out within HACID has shown that such systems are not only technically desirable, but also socially meaningful when designed around literacy, inclusion, and facilitation rather than automation alone.

**Description of the Exploitation Activity** — The exploitation of HACID technologies in the energy domain would focus on decision support services for community energy planning and governance, rather than a standalone software product. Key exploitation activities include:

- Deploying HACID-inspired decision support platforms to support community energy cooperatives, municipalities, and regional energy agencies in decisions such as infrastructure siting, investment prioritisation, and energy mix planning.
- Using HACID workflows and knowledge graph structures to integrate heterogeneous data sources (technical data, environmental constraints, regulatory requirements, and community inputs).
- Providing facilitation and capacity-building services that combine the decision support platform with human moderation, enabling meaningful participation and trust.
- Adapting HACID components to support energy literacy and transparency, making complex trade-offs understandable to non-experts.

The primary beneficiaries would be community energy organisations, local authorities, NGOs, and public bodies supporting just and inclusive energy transitions.

### Roadmap to Implement the Exploitation

Short term (0–12 months):

- Translate the existing design blueprints into a minimal, domain-specific HACID prototype focused on a small number of recurring community energy decisions.
- Engage early adopters (e.g. community energy cooperatives, city councils) to validate feasibility and governance models.
- Define roles for human facilitators and clarify how AI outputs are presented, discussed, and contested.

Medium term (1–3 years):

- Expand the prototype into a configurable service that can be adapted to different local and regulatory contexts.
- Integrate additional AI components for scenario comparison, uncertainty exploration, and pattern detection across cases.
- Establish partnerships with public programmes and intermediary organisations that support community energy initiatives.
- Develop sustainable funding models (e.g. service contracts, public funding, cooperative membership fees).

Long term (3+ years):

- Position the HACID approach as a valid framework for participatory, data-informed energy decision-making.
- Build a shared repository of cases, workflows, and lessons learned to support cross-community learning.
- Extend the approach to regional and national planning contexts, while preserving bottom-up participation principles.

## 4.2. Horizon Scanning for Water Regulation

**Context of the Exploitation Activity** — Water regulation operates in a highly complex and uncertain environment, shaped by climate change, ageing infrastructure, evolving public expectations, and fragmented institutional responsibilities. Regulators increasingly need to move from reactive oversight towards proactive horizon scanning, identifying emerging risks, trends, and interdependencies before they materialise as crises. The HACID design exploration has highlighted that this challenge is not primarily a lack of data or models, but rather the difficulty of integrating diverse forms of expertise—scientific, regulatory, operational, and societal—into coherent decision processes. Hybrid human-AI systems are well suited to this context, provided they are designed to support transparency, accountability, and professional judgement.

**Description of the Exploitation Activity** — The exploitation of HACID technologies in the water regulation domain would focus on strategic decision support for regulators, particularly in horizon scanning and long-term planning. Exploitation activities include:

- Developing HACID-based horizon scanning services that allow regulators to pose open-ended questions and explore plausible future scenarios.
- Using knowledge graphs to structure regulatory knowledge, policy objectives, environmental indicators, and expert inputs.
- Integrating AI methods to detect weak signals, emerging patterns, and cross-sectoral dependencies, while keeping humans in control of interpretation and judgement.
- Supporting cross-organisational collaboration by making assumptions, evidence, and uncertainties explicit and shareable.

The main users would be national and regional water regulators, with potential spillover benefits for environmental agencies, policy units, and oversight bodies.

### Roadmap to Implement the Exploitation

Short term (0–12 months):

- Refine the speculative blueprints into a focused pilot addressing a clearly defined horizon scanning process.
- Identify a limited set of priority data sources and expert contributors to ensure feasibility and trust.
- Test the approach internally within a regulatory organisation (e.g., Ofwat) as a strategic foresight and learning tool.

Medium term (1–3 years):

- Expand the pilot into an operational decision support service embedded in regulatory workflows.
- Introduce AI-assisted scenario generation and comparison, with explicit mechanisms for expert validation and challenge.

- Develop governance frameworks addressing accountability, explainability, and data stewardship.
- Facilitate cross-regulator and cross-agency collaboration using shared HACID cases and outputs.

Long term (3+ years):

- Establish HACID-inspired horizon scanning as a core capability within water regulation.
- Extend the approach to other regulatory functions (e.g. performance assessment, investment planning).
- Promote international exchange of methods and insights, positioning HACID as a transferable model for regulation under uncertainty.

### 4.3. Asylum determination management

Given the sensitivity, legal constraints, and ethical implications of the domain, the exploitation plan is deliberately framed as a progressive, safeguarded exploration, grounded in the expert interviews with UNHCR and consistent with HACID's core principle of \*decision support without loss of human agency\*.

**Context of the Exploitation Activity** — Asylum determination is a high-stakes, knowledge-intensive decision process characterised by uncertainty, time pressure, and significant human consequences. Decision makers must integrate heterogeneous sources of information—legal frameworks, country-of-origin information, geopolitical developments, and individual narratives—while exercising professional judgement and empathy during interactions with asylum seekers. Current asylum systems face structural challenges: increasing caseloads, rapidly evolving geopolitical contexts, fragmented information sources, and strong requirements for procedural fairness, transparency, and consistency. At the same time, full automation of asylum decisions is widely regarded as unacceptable, both ethically and legally. This makes asylum determination a natural fit for a hybrid collective intelligence approach. The interviews conducted with UNHCR experts indicate high potential for HACID technologies precisely because human and AI capabilities are complementary: humans excel at contextual interpretation, credibility assessment, and ethical judgement, while AI can support large-scale information integration, pattern recognition across cases, and systematic application of best practices.

**Description of the Exploitation Activity** — The exploitation of HACID technologies in asylum determination would focus on decision process support, not decision automation. The core activity would be the development of a HACID-based decision support framework that assists case officers throughout the asylum determination workflow while preserving human responsibility and accountability. Key elements of the exploitation activity include:

- Structuring asylum determination as a transparent, auditable workflow supported by HACID, where each step is informed by relevant knowledge and evidence.
- Developing a domain-specific knowledge graph that integrates legal guidelines, country-of-origin information, jurisprudence, geopolitical developments, and procedural best practices.
- Using AI components to surface relevant precedents, identify gaps or inconsistencies in evidence, and suggest process paths based on comparable cases worldwide.

- Enabling expert contributions (e.g. legal advisors, country experts) to enrich and validate the knowledge base over time.
- Supporting explainability and traceability, allowing decision makers to justify decisions clearly to applicants, courts, and oversight bodies.

Potential exploiters and beneficiaries include international organisations (e.g., UNHCR), national asylum authorities, and NGOs involved in refugee protection and legal assistance.

### **Roadmap to Implement the Exploration and Exploitation**

Short term (0–12 months): Feasibility and Safeguarded Exploration

- Conduct additional expert-led co-design activities with UNHCR and national asylum authorities to map decision workflows, pain points, and ethical constraints.
- Identify a non-operational pilot scope, such as training, retrospective case analysis, or decision rehearsal, to avoid direct impact on live cases.
- Define strict ethical, legal, and governance requirements, including human-in-the-loop guarantees and non-automation principles.
- Develop an initial conceptual blueprint of a HACID-supported asylum determination process.

Medium term (1–3 years): Prototyping and Controlled Pilots

- Build a prototype HACID decision support tool focused on specific subtasks (e.g. country-of-origin information integration, precedent retrieval).
- Populate and validate a preliminary asylum domain knowledge graph with expert oversight.
- Test the prototype in controlled environments (e.g. simulated cases, internal reviews, training programmes).
- Evaluate impacts on decision quality, consistency, transparency, and user satisfaction.

Long term (3+ years): Operationalisation and Responsible Scaling

- Deploy HACID-based decision support tools in limited, well-governed operational contexts, subject to regulatory approval.
- Expand the knowledge graph and AI components through continuous expert contribution and validation.
- Establish governance structures for accountability, auditing, and redress.
- Explore transferability across jurisdictions while respecting legal and cultural differences.
- Position HACID as a reference framework for responsible AI-supported asylum decision making, rather than a proprietary system.

## **4.4. Concluding remarks**

Across energy systems, water management, and asylum determination, the HACID project demonstrates the broad relevance of hybrid human–AI decision support in domains characterised by complexity, uncertainty, and high societal impact. Despite their differences, these three domains share common challenges: fragmented and evolving knowledge, distributed expertise, strong ethical and accountability requirements, and the need for decisions that are not only technically sound but also socially legitimate and transparent.

The HACID approach offers a unifying response to these challenges by embedding artificial intelligence within structured, participatory decision processes that preserve human agency. In community-led energy systems, HACID enables inclusive and informed collective choices

that integrate technical optimisation with local knowledge and values. In water management and regulation, it supports proactive horizon scanning and long-term planning by combining diverse data sources and expert perspectives into coherent, explainable foresight processes. In asylum determination, it provides a safeguarded framework for supporting case officers with integrated legal, geopolitical, and contextual knowledge, while ensuring that responsibility and judgement remain firmly with human decision makers.

Taken together, these application domains illustrate that the value of HACID lies not in automating decisions, but in augmenting professional and collective judgement through transparent workflows, shared knowledge structures, and accountable human–AI collaboration. Exploitation of HACID technologies therefore depends as much on governance, facilitation, and trust-building as on technical performance. By positioning hybrid collective intelligence as a service, a methodology, and a shared infrastructure, HACID has the potential to establish a transferable and responsible model for decision support in a wide range of complex professional contexts.