HACID - Deliverable

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.

Deliverable number:	D8.7
Due date:	29.02.2024
Nature[1]:	R
Dissemination Level[2]:	PU
Work Package:	WP8
Lead Beneficiary:	CNR
Contributing Beneficiaries:	all

---

[1] 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.

[2] 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

Version	Date	Description	Author	Partner
0.1	19/01/2024	Creation of draft document	Vito Trianni	CNR
0.2	29/2/2024	Description of exploitation for Medical Diagnostics	Gioele Barabucci	HDX
0.3	13/3/2024	Exploitation for Medical Diagnostics	Stefan Herzog	MPG
0.4	14/3/2024	Stakeholder inclusion activities	Aleks Berditchevskaia	NESTA
0.5	11/4/2024	Advanced draft	Vito Trianni	CNR
1.0	17/04/2024	Final Version	Vito Trianni	CNR

Table of content

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 to be identified through activities with external stakeholders. In the following, we provide the current exploitation plan grouping the activities according to the target application domains.

2. Plan for exploitation in the healthcare domain

2.1. Plan for exploitation of foreground knowledge

According to the current and prospective results, we have identified foreground knowledge—possibly still under development—that can find space for exploitation within the medical diagnostics domain, with particular reference to the possible integration within the Human Dx online platform. 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).

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 and methods 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 will be published and described in open access academic papers, as stated in the Data Management Plan. In addition, the plan for exploitation foresees that Human Dx EU identifies the most promising results to be further developed for integration in their online platform.
Problem: There is currently widespread interest in automated systems for decision support in healthcare settings. Within the context of medical diagnostics, there is a need to reduce diagnostics errors in a cost-efficient manner. AI-based solutions, most recently based on Large Language Models, are emerging as valuable tools to support the diagnostic process. However, their exploitation into medical practice still needs to be thoroughly evaluated.
Description: HACID proposes a hybrid collective intelligence approach to surface differential diagnoses from medical experts and to aggregate them into a single collective differential. Differently from fully automated systems, HACID places the medical expert at the centre, therefore fostering effectiveness and trust. The backend algorithms are instrumental to this objective, as they tackle two aspects of the envisaged process.

1. The automatic identification of relevant medical concepts from case descriptions and scientific/grey literature. This task produces relevant information to determine which subset of the domain knowledge graph is relevant for a specific case. To this end, we

exploit NLP methods for named entity recognition, linking and disambiguation. Additionally, approaches for information extraction are also studied and adapted to the project needs.

2. The automatic enrichment and aggregation of medical concepts provided by medical experts (either human or artificial). To this end, we exploit methods from knowledge engineering, cognitive network science, research on judgement aggregation and collective intelligence.

Both of these methodologies exploit the DKG for the medical diagnostics domain, which integrates different sources of information, mainly mediated from SNOMED CT and from WikiData. This DKG can be expanded by linking sources of evidence from scientific and grey literature to concepts and relations within the DKG.

Roadmap: After successful experimentation within the project, all or a part of the backend functionalities can get integrated into the Human Dx platform and released in production. Human Dx EU will perform an evaluation of the costs and benefits expected from such an integration, also in the light of the constraints of the platform with respect to offering training and support services. In a period of 6 months to 1 year after the project ends, the selected functionalities will be tested and possibly improved for integration and deployment. The expected impact of these additional functionalities is improved user experience due to enhanced and richer feedback about complex diagnostics cases.

2.2. Plan for expansion towards new services

HACID aligns with a number of recent global developments and initiatives emerging in response to experiences and lessons learned from the COVID-19 pandemic. These include the WHO Hub for Pandemic and Epidemic Intelligence with its EPI-BRAIN, aimed at (i) the development, evolution, maintenance and expansion of a semantic network of distributed data through the development of fit-for-purpose taxonomies, ontologies and related standards; (ii) the establishment of a sustainable pandemic and epidemic intelligence global data architecture; and (iii) the creation of semantically linked data for analytics and insights.

Additionally, other organisations in healthcare demonstrated interest in the project following initial dissemination activities. Specifically, Nesta held several meetings with the AI team within the UK National Health Service (NHS) to introduce the technology and explore collaboration opportunities. The NHS team provided informal feedback on the HACID concept. In the following, we provide a plan for action with respect to the above initiatives.

WHO Hub for Pandemic and Epidemic Intelligence

In the following, we detail the different phases that can lead to expansion towards new services in collaboration with the WHO Hub for Pandemic and Epidemic Intelligence.

Phase 1: Initial Engagement and Scoping (M19-M24)

- Objectives
  1. Understand the WHO Hub's strategic objectives and operational challenges.
  2. Identify how the HACID framework can support the WHO Hub's mission.
- Actions:
  1. *Stakeholder Meeting*: Organise discussion with WHO Hub representatives to gain insights into their strategic goals and the role of the EPI-BRAIN initiative.
  2. *Literature Review*: Review WHO Hub's strategy papers and related documents to align HACID's capabilities with the Hub's needs.

iii. *Scoping Workshop*: Conduct workshop with stakeholders from MPG, CNR, and Nesta to define the scope of HACID's potential application within the WHO Hub.
iv. *Feasibility Analysis*: Assess the technical and operational feasibility of integrating HACID's decision support system with the WHO Hub's data architecture and processes.

Phase 2: Framework Alignment (if feasibility turns positive, M25-M36)

- Objectives:
  i. Align HACID's framework with the WHO Hub's technological and data requirements.
  ii. Develop a strategic plan for potential integration.
- Actions:
  i. *Alignment Sessions*: Facilitate sessions between HACID developers and WHO Hub's technical team to align the framework's capabilities with the WHO Hub's requirements.
  ii. *Strategic Planning*: Create a strategic plan outlining how HACID can be adapted to support the WHO Hub's objectives, considering the development of taxonomies, ontologies, and data architectures.
  iii. *Documentation*: Prepare detailed documentation of the alignment plan, including roles, responsibilities, and a roadmap for potential future integration.
  iv. *Stakeholder Feedback*: Gather and incorporate feedback from WHO Hub stakeholders to refine the strategic alignment plan.
  v. *Fundraising Strategy Development*: Develop a strategic fundraising plan to secure the necessary resources for adapting the HACID-DSS to the health domain, in collaboration with MPG, CNR, and Nesta, and in alignment with the WHO Hub's objectives

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 such structured knowledge, and a web application that integrates these components implementing the full HACID-DSS for climate services.

The DKG developed within the HACID project will be published open access, following the strategy detailed in the Data Management Plan. Such a public resource can have a strong impact on the climate service community, because there is a lack of formalised knowledge in the domain. Also, the DKG should be considered a living resource that gets updated and refined over time, following new findings and the production of new evidence. Hence, exploitation could be possible by providing services to interested users that need to access and explore the DKG. Such a service can be provided through the HACID-DSS web application, and in particular through the interactive dashboard developed within HACID. Indeed, the way in which users explore and interact with such structured knowledge is crucial for evidence-based decision making. The dashboard is an invaluable software tool to support climate scientists in their daily activities, adding further value to the HACID-DSS solution. Finally, decision support through hybrid collective intelligence can have a large impact on the climate service community, both providing a way to explore solutions to diverse problems, as well as to build and strengthen a network of experts and trainees around a shared product.

Problem: 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 experts that can offer solutions to a specific case, there is an opportunity to find novel 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.

Description: The HACID-DSS for climate services provides decision support thanks to the elicitation of solutions from experts and to the aggregation of feedback from multiple experts. Solution elicitation is supported by the DKG and the linked visualisation tool, which allows tailored exploration routes that can easily lead to discovery of relevant concepts and supporting evidence. Aggregation of expert solutions is performed by means of algorithms exploiting the structured information available in the DKG to find similarities among methodologies and scenarios, ultimately providing a collective solution that exploits all the provided input. Such algorithms are developed integrating the values elicited from users during participatory workshops, and are evaluated in terms of effectiveness and trustworthiness.

Roadmap: To guarantee exploitation of the HACID-DSS beyond the project lifetime, extensive validation of the platform with stakeholder users is mandatory, Met Office provides cases for testing that are exploited both for platform validation and for expert network expansion. Additionally, experts are invited to test the platform submitting new cases to the growing community, inviting colleagues from their entourage to join the platform. For successful validation and exploitation, we plan to include at least 10 test cases and recruit an initial cohort of 30 experts that can contribute to populating the platform and make it a lively place for interaction with fellow climate scientists.

After successful validation, Met Office collaborates with CNR for the development of a business plan that can sustain the platform beyond the project lifetime. Securing funding from private and public sources for continuing research and bringing the platform to the required TRL 8 and 9 will be the major focus of the activities, together with expansions towards new services as discussed below.

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. HACID remains constrained to the core of the activities to provide climate services, especially focusing on the methodologies for providing relevant climate projections. However, new opportunities and challenges arise to target aspects related to the integration of information relevant to the end users requesting climate services. More specifically, expansion of the functionality offered by the HACID-DSS can be evaluated, exploring aspects like surfacing user needs and risk appetite, integrating knowledge from specific application domains (e.g., urban mobility planning), providing explanations about uncertainty ranges, helping users navigate and discover climate information. During the last year of the project, Met Office will explore these and other opportunities in interaction with fellow climate scientists and service providers to understand in what direction there is a larger potential impact for the provision of new services based on the HACID-DSS concept.

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. We plan to identify relevant domains through a structured design process working with decision makers and practitioners across a range of policy domains. We anticipate this process will involve:

Phase 1 - Scoping (M22-M24)

1. Running an internal workshop with HACID partners to specify the design constraints for a “HACID-relevant” domain.
2. Generating a longlist of potential domains (e.g. disaster response and risk reduction, energy systems,) and narrowing this down to a shortlist of the most promising two domains based on criteria (e.g. feasibility, strategic priority, policy impact, existing partnerships etc.)
3. Carrying out additional rapid domain research where needed to improve the problem brief. This may include small ideation workshops within the consortium and critical friends.
4. Identifying relevant stakeholders and 1-2 key “partner” organisations with legitimacy in the domain. Scope out opportunities with partners.

Phase 2 - Ideation (M25-M30)

5. Designing and planning of workshops. We anticipate these will include both remote and in-person meetings.
6. Running 4 design workshops (2 for each domain of interest) with stakeholders with expertise in the shortlisted domains and relevant policy makers and/or regulators:
    a. In these workshops, we will introduce the HACID-DSS concept and provide examples of different instantiations of this from medical diagnostics and climate services, based on the prototypes and design concepts developed during T6.1 and T7.1.
    b. We will then generate ideas for potential applications or challenges in their domains that could benefit from this type of technology. We will use a variety of design thinking tools, including those from Nesta’s collective intelligence design playbook

Phase 3 - Concept development and design (M30-M34)

7. Gathering the output from the process as a (set of) design concept(s) that will support further exploitation of the HACID approach. This might be, for example, a service blueprint illustrating how a HACID system could work in another domain or a roadmap, which details the stepwise process for HACID technology development, including the necessary technical and process infrastructure.
8. Iterating and reviewing the design concepts with HACID partners and external stakeholders.
9. Publishing and disseminating design outputs with relevant audiences.