RESKILLING

Deliverable 4.6

CCAM Employment & Skills Observatory

Grant Agreement no.: 101147328

Project acronym: RESKILLING

Project title: Research initiative for Enhancing and Adapting Workforce SKILLs for Implementing TraNsport Automation with Employment Growth

Project coordinator name and organisation:

Matina Loukea, Centre for Research and Technology Hellas (CERTH)

Funded by the European Union

This project has received funding from the European Commission under the European Union's Horizon Europe research and innovation programme under grant agreement No. 101147328

This work has received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI) for the Swiss partners

Document Details	
Document identifier	D4.6 CCAM employment & skills Observatory
Due Date of Delivery	30/06/2026
Actual Date of Delivery to EC	29/06/2026
Task Number	T4.5
Dissemination level	SEN - Sensitive

Contributors
Nikolaos Chandrinos (CERTH), Matina Loukea (CERTH)

Reviewers
PCG members

Disclaimer

Control Sheet			
Version	**Date**	**Editor**	**Summary**
1	11/05/2026	Matina Loukea (CERTH)	Draft version, updating the M18 report
1.1	25/05/2026	Nikolaos Chandrinos (CERTH)	Update of report, based on the web-based platform updates
2	19/01/2026	Matina Loukea (CERTH)	Review and update of report
2.1	23/06/2026	Nikolaos Chandrinos (CERTH)	Review and update of report
3	29/06/2026	Matina Loukea (CERTH)	Final version, ready for submission

List of Acronyms	
Acronym	**Definition**
AI	Artificial Intelligence
CCAM	Connected, Cooperative and Automated Mobility
D	Deliverable
ESCO	European Skills, Competences, Qualifications and Occupations
EU	European Union
EU	European Union
ICT	Information and Communication Technology
ID	Identity
OEM	Original Equipment Manufacturer
QA	Question-Answering
RAG	Retrieval-augmented generation
SDK	Software Development Kit
T	Task
UI	User Interface
VET	Vocational Education and Training
WP	Work package

Table of Contents

List of Figures

Executive Summary

This deliverable presents the first official version of the CCAM Employment & Skills Observatory, developed within the RESKILLING project as one of its central tools and long-term outcomes. The Observatory is designed to support the systematic organisation, exploration and use of knowledge on how Connected, Cooperative and Automated Mobility (CCAM) affects employment, occupations, emerging job roles, skills needs and reskilling pathways across the European transport sector.

The Observatory responds to the need for a reliable, accessible and stakeholder-oriented knowledge resource that goes beyond the presentation of static information. It is conceived as a dynamic knowledge repository and decision-support tool, enabling different stakeholder groups to access tailored insights according to their needs, priorities and areas of interest.

The first official version establishes the functional and technical basis of the Observatory, bringing together RESKILLING outputs, selected external sources, scientific and policy literature, taxonomies and stakeholder-generated knowledge into an initial structured knowledge base. It also integrates an AI-powered search engine that allows users to interact with the Observatory through natural-language queries, supported by semantic retrieval and evidence-based answer generation. The system is designed to retrieve relevant information from curated sources and generate responses grounded in the available knowledge base.

A key element of the Observatory is its ontology-based approach, which supports the semantic organisation of content and creates the basis for links with existing taxonomies and classification systems, including ESCO and other relevant external knowledge schemes. This structure will enable the Observatory to connect information on jobs, skills, training resources, stakeholder needs and labour-market trends in a coherent and extensible way. It also provides the foundation for future recommender functionalities and more personalised services.

The Observatory has been integrated into the RESKILLING website, where it is presented as a dedicated project tool and entry point for users. Through this interface, users can submit questions, explore knowledge through structured categories or contextual filters, and receive AI-supported answers based on retrieved evidence. This integration strengthens the visibility and accessibility of the Observatory and supports its role as a durable RESKILLING asset.

The deliverable also presents the validation activities carried out for the first official version. Internal testing focused on the functionality of the AI workflow, including document ingestion, semantic retrieval, filtering, answer generation, source grounding and system logging. Partner validation complemented this process by assessing the Observatory from a user and stakeholder perspective, focusing on usability, relevance of results, clarity of responses and the practical value of the tool for exploring CCAM employment and skills information.

The current version should be understood as an operational first release rather than the final public version of the Observatory. Its coverage and functionalities will continue to be expanded during the remaining project period. Future work will focus on enriching the knowledge base with additional RESKILLING outputs and external sources, refining the ontology and taxonomy links, improving the user interface, strengthening feedback mechanisms, developing recommender functionalities and preparing the Observatory for wider stakeholder access.

Overall, this deliverable confirms the establishment of a solid basis for the CCAM Employment & Skills Observatory. By connecting evidence on affected jobs, emerging skills, new roles, reskilling needs and stakeholder priorities, the Observatory will support RESKILLING's broader objective of helping the European transport sector prepare for the employment and skills implications of CCAM.

1. Introduction

1.1. Purpose of the Document

The CCAM Employment & Skills Observatory is one of the central tools and outcomes of the RESKILLING project. It is designed to support the systematic collection, organisation and exploration of knowledge on how Connected, Cooperative and Automated Mobility (CCAM) affects employment structures, occupations, emerging job roles and skills requirements across the European transport sector. The Observatory responds to the need for accessible, reliable and stakeholder-oriented intelligence on workforce transformation, supporting policymakers, transport operators, industry actors, training providers, social partners, researchers and other stakeholders in understanding and preparing for CCAM-related labour-market change. The purpose of the Observatory goes beyond the presentation of information. It functions as a knowledge repository and decision-support tool, providing tailored insights according to the needs, preferences and priorities of different stakeholder groups.

The Observatory combines several complementary functionalities. It supports structured access to knowledge through predefined categories and keyword-based search, while also integrating an AI-powered search engine that enables users to interact with the Observatory content through natural-language queries. This functionality is intended to make complex information on CCAM employment and skills transitions easier to explore, while ensuring that responses remain grounded in validated sources and within the scope of the RESKILLING project. In parallel, the development of an ontology supports the connection of Observatory content with existing taxonomies and classification schemes, such as ESCO, and potentially with other external data sources and skills-related frameworks.

This deliverable presents the first official version of the CCAM Employment & Skills Observatory. It builds on the technical work previously documented under the relevant project milestone, where the AI-powered search engine and supporting validation environment were described. The present deliverable broadens that scope by presenting the Observatory as an integrated project tool, including its conceptual basis, content structure, user-oriented functionalities, platform integration, AI-supported search capabilities, ontology-related work, and feedback mechanisms. User feedback will be collected through assessment questions to evaluate satisfaction, usability, relevance of results and areas for improvement.

The current version should be understood as an operational first release within the RESKILLING project lifecycle. It establishes the functional and technical basis of the Observatory and prepares the ground for further refinement, public use and long-term exploitation. The Observatory will continue to be enriched and further refined throughout the remaining project duration, particularly after the platform becomes publicly available and feedback from end users has been collected, analysed and incorporated. Its evolution, together with the implemented improvements and new functionalities, will be documented in Deliverable 4.7 – *Update of the CCAM Employment & Skills Observatory***,** due in Month 34 of the project.

In this way, the Observatory is envisaged not only as a key digital output of the RESKILLING project, but also as a sustainable and visible resource that will facilitate the long-term uptake, exploitation and further development of the project's results beyond its lifetime.

1.2. Intended Audience

This deliverable is intended for the RESKILLING consortium, the European Commission, and the wider CCAM, transport, employment, education and skills ecosystem. It presents the first official version of the CCAM Employment & Skills Observatory as a central project outcome and describes its purpose, content, functionalities, technical basis and planned evolution.

The Observatory itself is designed to support a broad range of stakeholders involved in, or affected by, CCAM-related workforce transformation. These include policymakers, public authorities, transport operators, industry actors, education and training providers, social partners, researchers, workers'

representatives and other organisations interested in employment impacts, emerging occupations, skills needs and reskilling pathways linked to CCAM deployment. A more detailed analysis of these target users and their specific needs is provided in the following section.

1.3. Interrelations

The CCAM Employment & Skills Observatory represents the integration point of several key RESKILLING activities and serves as the project's central knowledge repository for employment and skills intelligence. The Observatory builds directly upon the outputs generated in WP2, particularly the stakeholder typologies developed in Task 2.1, which define the different user groups served by the platform and guide the personalisation of information and services. It also exploits the results of WP3, including the mapping of occupations and job families (D3.1), the assessment of employment impacts and labour-market projections (D3.3), the analysis of socio-economic impacts (D3.4), and the identification of differentiated impacts on special workforce groups. Together, these provide the scientific evidence underpinning the Observatory's knowledge base.

Within WP4, the Observatory constitutes the main implementation activity of Task 4.5, where the knowledge graph, ontology, AI-powered search engine, automated content ingestion pipelines and recommender system are developed and integrated into a single stakeholder-oriented platform. The ontology establishes semantic links between Observatory content and existing taxonomies and classification systems, including ESCO and other relevant external knowledge schemes, enabling interoperability and future extensibility. The Observatory also consolidates information generated through the RESKILLING analytical models, emerging skills analysis and labour-market forecasting activities.

The Observatory is closely integrated with the RESKILLING Inventory (Task 6.3), ensuring that curated project resources, publications, training materials, stakeholder outputs and external references become directly accessible through the platform. Furthermore, the Observatory is designed to interface, free of charge, with the EU-wide Knowledge Base developed under the FAME project, contributing employment- and skills-related intelligence to the wider European CCAM ecosystem and strengthening interoperability between complementary initiatives.

Beyond the project, the Observatory is expected to become one of RESKILLING's principal exploitation results. Through collaborations with industry, education and training providers, research organisations and regulatory bodies, as well as through the publication of selected knowledge via Linked Open Data repositories, the Observatory aims to establish itself as a sustainable and trusted European reference point for CCAM-related employment and skills intelligence, supporting evidence-based decision-making long after the completion of the RESKILLING project.

2. The CCAM Employment & Skills Observatory: Concept and Scope

2.1. Observatory objectives

The CCAM Employment & Skills Observatory aims to provide a central knowledge and decision-support tool for understanding how CCAM affects employment, occupations, skills and workforce transition needs across the European transport sector. Its primary objective is not only to present project results, but to organise, connect and make accessible evidence on jobs affected by CCAM deployment, emerging job roles, new skills requirements, reskilling needs and labour-market trends.

More specifically, the Observatory is designed to:

- consolidate and structure knowledge generated within RESKILLING on CCAM-related employment and skills impacts;
- support stakeholders in identifying occupations, job families and workforce groups likely to be affected by CCAM deployment;
- provide insights into emerging skills, new roles and potential reskilling pathways;
- enable user-oriented access to information through keyword search, predefined categories and AI-powered natural-language queries;
- connect project outputs with existing taxonomies and external knowledge schemes, including ESCO and other relevant classification systems;
- support tailored recommendations and personalised insights according to the needs, priorities and profiles of different stakeholder groups;
- collect feedback from users to assess relevance, usability and areas for improvement;
- contribute to the long-term exploitation of RESKILLING results by establishing a durable knowledge resource for the CCAM ecosystem.

The Observatory therefore functions both as a knowledge repository and as a practical support mechanism for evidence-based decision-making. It is intended to help policymakers, public authorities, industry actors, transport operators, education and training providers, social partners, researchers and workforce representatives anticipate and respond to CCAM-related workforce transformation in a more informed and inclusive way.

2.2. Target users and stakeholder needs

Building on the stakeholder typologies developed in Task 2.1, the Observatory is designed to provide tailored knowledge and decision-support services to different user groups. The aim is not only to make information available, but to organise and present it in a way that responds to the specific needs, priorities and decision-making contexts of each stakeholder category.
The primary target audiences include:

- **European, national, regional and local policymakers**, responsible for transport, employment, education, innovation and industrial policy, seeking evidence to support workforce planning and policy development.
- **Public authorities and regulatory bodies**, involved in the planning, deployment and governance of CCAM, requiring reliable information on workforce implications and future skills needs.
- **Transport operators, mobility service providers, OEMs, suppliers and technology developers**, interested in anticipating workforce transformation, identifying future competence needs, and supporting organisational change.
- **Education, vocational education and training (VET) providers, universities and lifelong learning organisations**, requiring up-to-date intelligence on emerging occupations, skills requirements and future curricula development.

- **Workers, job seekers and professional associations**, seeking information on occupations affected by CCAM, future career pathways, required competences and available learning opportunities.
- **Social partners, labour unions and employer organisations**, supporting social dialogue and workforce transition through evidence-based understanding of employment impacts and reskilling priorities.
- **Researchers, innovation organisations and European projects**, interested in accessing structured knowledge on CCAM employment, skills, labour-market trends and related analytical resources.
- **The RESKILLING Consortium and the wider CCAM ecosystem**, including initiatives under the CCAM Partnership and related European projects, using the Observatory as a common knowledge base supporting collaboration, dissemination, exploitation and future research.

Through its AI-powered search capabilities and recommender system, the Observatory is designed to provide personalised access to information according to the profile, interests and objectives of different stakeholder groups, moving beyond generic information retrieval towards stakeholder-oriented knowledge support.

2.3. Scope of the first official version

The first official version of the CCAM Employment & Skills Observatory represents an operational release developed within the RESKILLING project lifecycle. It establishes the functional and technical basis of the Observatory and demonstrates how project-generated knowledge, external sources and AI-supported tools can be combined to support the exploration of CCAM-related employment and skills impacts.

At this stage, the Observatory includes the core elements required to organise and access relevant knowledge on jobs, skills and workforce transformation. These include an initial knowledge base, structured content categories and an AI-powered search engine. The first official version also provides the basis for the development of more advanced user-oriented services, including stakeholder-specific recommendations and feedback-based improvements.

The scope of this first version is therefore to present and validate the Observatory as an integrated RESKILLING tool, rather than as a fully finalised public service. It documents the current structure, functionalities, content coverage, platform integration and technical components of the Observatory, while also identifying the areas that will be further enriched and enhanced during the remaining project period, including the expansion of the knowledge base, improvement of the user interface, integration of additional project outputs and external resources, and systematic collection of user feedback.

A distinction should therefore be made between the first official version and the final public version of the Observatory. The first official version is an operational and internally validated release, suitable for demonstrating the Observatory's core concept, content structure and search capabilities. The final public version, to be developed towards the end of the project, will build on this foundation and will reflect additional content, wider stakeholder use and respective user feedback and further technical refinement. In particular, feedback collected through the assessment questions already available will be used to evaluate user satisfaction, relevance of results, clarity of outputs and areas where the Observatory can better respond to stakeholder needs.

In this sense, the first official version is both a deliverable outcome and a development milestone. It confirms that the Observatory has moved beyond conceptual design towards an operational tool, while preserving space for continued improvement before its wider public use and long-term exploitation beyond the project duration.

2.4. Interface and user journey

The first official version of the CCAM Employment & Skills Observatory has been integrated into the RESKILLING website, where it is presented as one of the central digital tools and outcomes of the project. This integration ensures that the Observatory is not developed as a standalone technical component, but as part of the wider RESKILLING digital ecosystem, alongside other project resources, outputs and stakeholder-oriented information. Until at least the end of the project, the Observatory will be accessible through the RESKILLING website, providing users with a dedicated entry point to explore knowledge on CCAM-related employment impacts, emerging job roles, skills needs and reskilling pathways. The integration with the RESKILLING website supports the visibility, accessibility and long-term exploitation of the Observatory. By embedding the tool within the project website, users can access it in a familiar project environment and connect it with other RESKILLING results, including project deliverables, stakeholder engagement outputs, the RESKILLING inventory, training-related resources and dissemination material. This approach also supports the positioning of the Observatory as a durable project asset, with potential to remain available and evolve beyond the project duration.

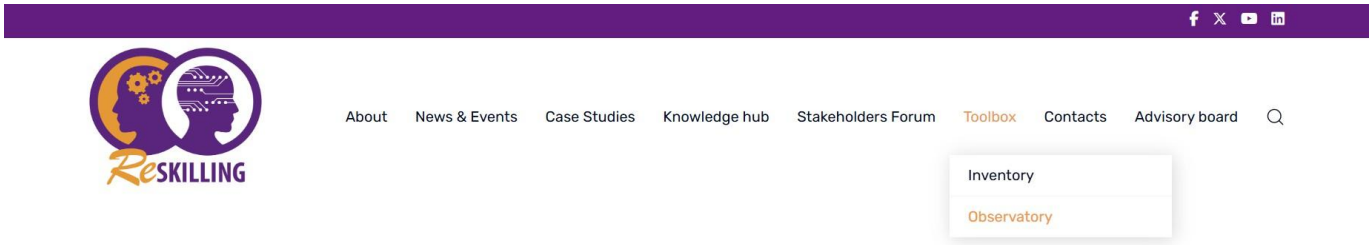

Figure 1: Integration of CCAM Employment & Skills Observatory into the RESKILLING website

Users access the Observatory through a dedicated section or entry point on the RESKILLING website. From there, they are directed to the Observatory interface, where they can interact with the available knowledge base. The user journey is designed to be simple and intuitive: users can submit a question or search query related to CCAM employment and skills, optionally refine the search through contextual filters or predefined categories, and launch the AI-supported search workflow. This allows different types of users, including policymakers, transport operators, training providers, social partners, researchers and workforce representatives, to explore information according to their specific interests.

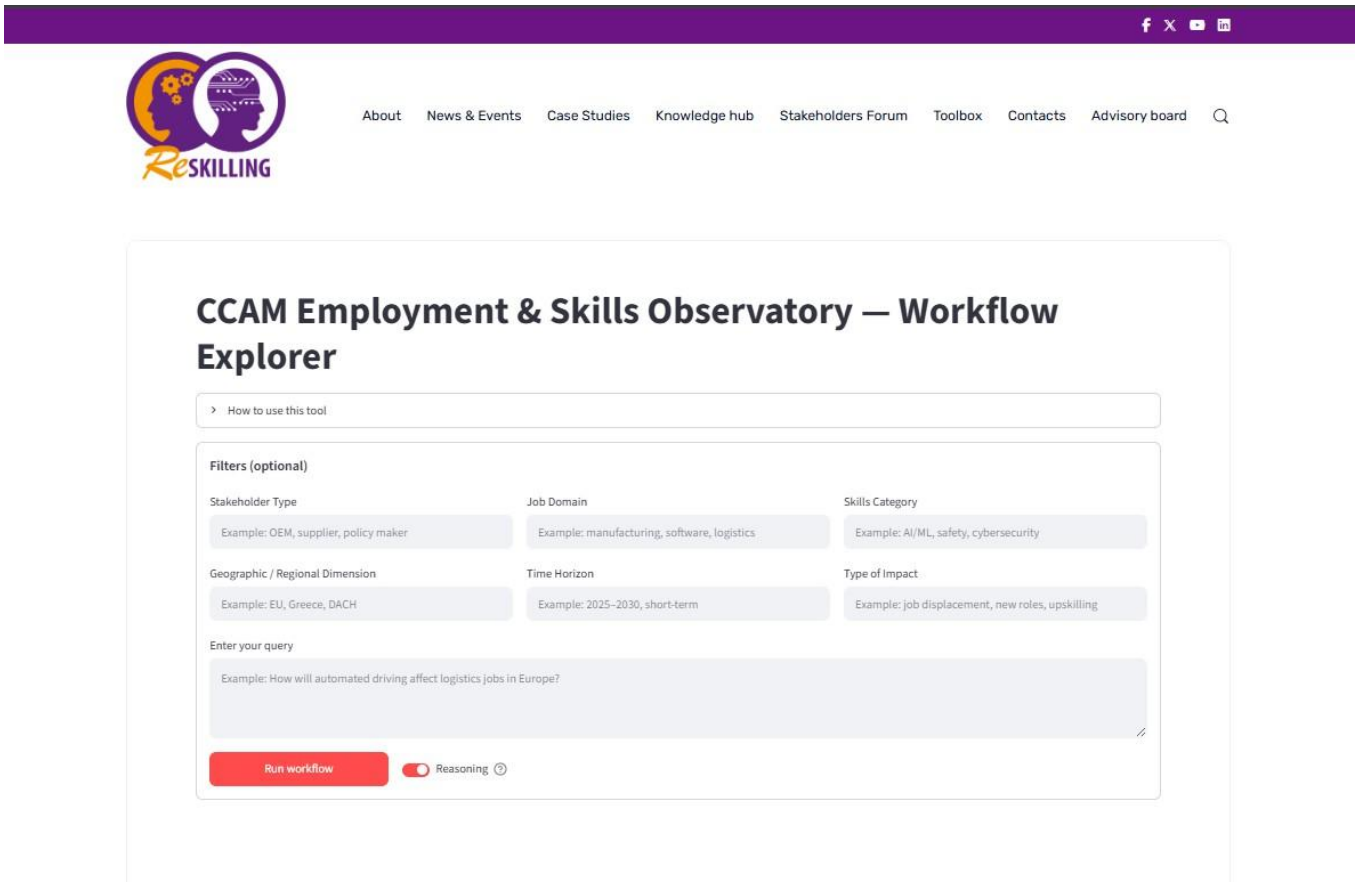

Figure 2: CCAM Employment and Skills Observatory interface

The user-facing interface presents the Observatory as an interactive search and knowledge exploration environment. The main interaction point is the AI-powered search engine, which enables users to ask natural-language questions rather than relying only on predefined keywords. For example, a user may ask about occupations affected by CCAM deployment, emerging skills for automated mobility, reskilling needs for specific workforce groups, or the expected impact of automation on transport jobs. The system then searches the available knowledge base and generates an evidence-based answer grounded in retrieved content.

The results are displayed in a structured and user-oriented format. The response generated by the AI-powered search engine is presented as a readable answer, supported by the retrieved evidence used by the system. Where relevant, the interface may also display supporting information such as source references, retrieved text segments, related documents, categories or contextual elements that help users understand where the answer comes from. This contributes to transparency and supports users in assessing the relevance of the information provided.

Overall, the integration of the Observatory into the RESKILLING website marks an important step in moving from technical development towards an operational user-facing tool. It provides a practical access point for stakeholders, enables interactive exploration of CCAM employment and skills knowledge, and establishes the basis for further refinement through user feedback, expanded content coverage and future improvements to the interface and recommender functionalities.

3. Observatory Content and Knowledge Base

3.1. Types of knowledge sources

The CCAM Employment & Skills Observatory brings together different types of knowledge sources in order to provide a comprehensive and evidence-based overview of CCAM-related employment and skills transformations. The knowledge base is designed to combine project-generated evidence with relevant external sources, allowing users to explore the links between CCAM deployment, affected occupations, emerging job roles, skills requirements, reskilling needs and stakeholder priorities.

The main types of knowledge sources considered for the Observatory include:

- **RESKILLING project outputs**, through also the RESKILLING Inventory, including deliverables, analytical reports, mappings, taxonomies, stakeholder consultation results, survey findings, use case analyses, training-related outputs and policy-oriented results.
- **Scientific literature**, covering automation and employment, digitalisation, skills transitions, transport labour markets, human-machine interaction, workforce inclusion and socio-economic impacts of technological change.
- **Policy and institutional reports**, including publications from European and international organisations such as the European Commission, Cedefop, OECD, ITF, Eurofound, ILO, the European Parliament and other relevant bodies.
- **Semantic linking to existing taxonomies and classification systems**, such as ESCO and other occupational, skills or training-related frameworks.
- **Reports including outputs from related European projects and initiatives**, including previous and ongoing work relevant to CCAM, workforce transition, skills development, automation impacts, training needs and stakeholder engagement. Relevant examples include SKILLFUL, SHOW, Drive2theFuture, CCAM-ERAS, FAME, and the *Study on the social dimension of digitalisation in the labour force*, among others.
- **Stakeholder-generated knowledge**, including insights from interviews, workshops, surveys, expert consultations and feedback collected through RESKILLING activities.
- **Training and reskilling resources**, including curricula, learning pathways, training profiles, skills frameworks and education or vocational training material relevant to emerging CCAM occupations.
- **External data sources and knowledge repositories**, where relevant and technically feasible, to enrich the Observatory and support future links with wider European knowledge bases.

These sources vary in format, level of detail and purpose. For this reason, the Observatory does not simply collect documents, but organises them through semantic search methods that capture their underlying meaning rather than relying only on keywords. This allows sources to be compared more effectively and retrieved more accurately according to their relevance to user queries.

In the first official version, the knowledge base focuses primarily on the sources already identified, generated or processed within the RESKILLING project, together with selected external references needed to support the Observatory's employment and skills intelligence. Further enrichment will take place during the next stages of the project, as additional deliverables, training outputs, stakeholder feedback and external sources become available. The Observatory will also be validated through the feedback mechanism already integrated into the platform.

3.2. Integration of RESKILLING outputs

The Observatory is designed as an integration point for the main employment, skills and stakeholder-related outputs produced within RESKILLING. Rather than presenting project results as separate documents, the Observatory aims to organise them into a structured and searchable knowledge base that supports stakeholder-oriented access, semantic connections and future exploitation.
A key input comes from the work carried out in WP3, which provides the analytical foundation for understanding how CCAM affects occupations, job families, skills and workforce groups. This includes the

mapping of CCAM-related professions and job families, the analysis of employment effects and skills impacts, the identification of new and emerging roles, and the assessment of differentiated impacts on special workforce groups. These outputs provide the core evidence for linking specific CCAM developments with affected jobs, changing tasks, required competences and reskilling needs. The Observatory also draws on the analytical models and labour-market projections developed under T3.2, which support the identification of future employment trends and potential workforce impacts. These results help move the Observatory beyond a static knowledge repository, enabling it to provide more forward-looking insights into how CCAM deployment may influence labour demand, skills requirements and occupational change over time.

The Observatory is also aligned with the stakeholder typologies developed in T2.1, ensuring that the information provided can be organised according to the needs and priorities of different user groups. This is important because policymakers, transport operators, training providers, industry actors, researchers, social partners and workforce representatives may use the Observatory for different purposes. The integration of stakeholder typologies therefore supports the development of personalised insights, tailored recommendations and user-oriented navigation pathways.

Furthermore, the Observatory will progressively integrate outputs from WP4 and WP5, including results related to reskilling pathways, training profiles, business model implications, policy recommendations, scalability and roadmap activities. This will allow the Observatory to connect evidence on affected jobs and skills with practical transition measures, training responses and policy actions.

In parallel, the Observatory is connected with the RESKILLING inventory developed under T6.3, which provides a structured collection of project-related resources, outputs and materials. Through this connection, the Observatory can support easier access to relevant project knowledge, including publications, training-related content, stakeholder engagement material, dissemination resources and other curated information produced during the project.

Overall, the integration of RESKILLING outputs ensures that the Observatory functions as a central project knowledge node. It consolidates evidence across work packages, links analytical findings with practical tools and supports the long-term exploitation of RESKILLING results. In the first official version, this integration establishes the core structure and content basis of the Observatory. In the final public version, it will be further enriched with additional project outputs, user feedback, refined ontology links, and extended connections with external taxonomies and knowledge repositories.

3.3. Current coverage and limitations

At this stage, the Observatory is particularly focused on organising and making accessible evidence related to the transformation of transport occupations, the identification of new and emerging job profiles, the mapping of skills and competences, and the analysis of differentiated impacts across stakeholder and workforce groups. The first version also integrates the AI-powered search functionality, allowing users to retrieve information from the available knowledge base through natural-language queries and structured search mechanisms. This provides the foundation for more advanced user-oriented services, including future recommender functionalities and more tailored stakeholder-specific insights.

However, the current version should be understood as an evolving first release rather than a fully finalised public knowledge service. Several limitations remain at this stage. First, the knowledge base is still being enriched as additional RESKILLING deliverables, training outputs, stakeholder consultation results, use case findings and policy-oriented outcomes become available. Second, the ontology and links with external taxonomies, such as ESCO, are still under refinement and will be further developed to improve semantic connections between occupations, skills, roles, training resources and stakeholder needs. Third, the current content coverage depends on the availability and maturity of project results at the time of this deliverable, meaning that some areas are more developed than others.

In addition, while the AI-powered search engine is operational, its performance depends on the quality, scope and structure of the underlying knowledge base. As the content expands and becomes more systematically annotated, the relevance and precision of search results are expected to improve. User-oriented functionalities, such as tailored recommendations, feedback mechanisms, filtering options and interface refinements, will also be further developed and tested during the next project phase. Particular attention will be given to improving usability, transparency of results, accessibility and the ability of the Observatory to respond to the different needs of stakeholder groups.

Another limitation relates to the dynamic nature of CCAM deployment and labour-market transformation. Employment impacts, skills needs and emerging job roles are likely to evolve over time as technologies mature, regulatory frameworks change, business models develop and new evidence becomes available. For this reason, the Observatory is designed as a living knowledge resource rather than a static repository. The first official version establishes the core structure, content basis and technical functionalities, while the final public version will incorporate additional evidence, broader stakeholder feedback, refined ontology links and improved user-facing services.

Overall, the current version provides a solid foundation for the CCAM Employment & Skills Observatory and demonstrates its potential as a central RESKILLING knowledge node. Its limitations are mainly linked to the staged development of the project and the need for progressive enrichment, validation and user testing. These limitations will be addressed through continued integration of RESKILLING outputs, expansion of external knowledge sources, refinement of the AI-powered search and recommender functionalities, and systematic collection of user feedback before the final public version is released.

4. Core Functionalities of the CCAM Employment & Skills Observatory – First version

The first version of the CCAM Employment & Skills Observatory implements a document-centred, AI-powered workflow for exploring employment and skills information related to CCAM. As part of the broader Observatory, this functionality enables users to submit questions, optionally apply contextual filters, run the AI-supported search workflow, and receive answers grounded in retrieved evidence from the Observatory knowledge base. The implementation combines a user-facing interface with a multi-agent backend, vector-based semantic retrieval, evidence-based answer generation and supporting monitoring functionalities. Trace monitoring, timing diagnostics and database logging are used to support development, testing, validation and continuous improvement of the system.

4.1. Document-centric knowledge repository

The Observatory is built around a document-centric knowledge repository. In the implemented workflow, retrieved evidence is not treated as anonymous text. Each retrieved chunk is linked to a file identifier, a vector store, a score, and the original chunk text. This means that the system is designed to answer questions using document fragments that can be traced back to source files, rather than generating unsupported responses from model memory.

The workflow retrieves information from configured vector stores. The code records the name and identifier of each vector store, collects search results from each store, and stores the number of results returned by each one. Two retrieval sources are active in the workflow: Observatory and Inventory. This is important because it shows that the first version already separates or combines evidence from more than one knowledge source, rather than depending on a single flat document collection.

The repository is document-centric because the system prepares the retrieved material as structured evidence before answer generation. For each valid retrieved result, the workflow extracts the file ID and chunk text. It then builds a combined context where every chunk is introduced by its file identifier. This combined text is passed to the Question-Answering (QA) agent as the evidence base for the final response.

This design supports source-aware question answering. The QA agent is instructed to cite factual claims using the file IDs that appear in the retrieved documents. It must not invent, alter, or explain file IDs. This ensures that the user-facing answer is connected to the underlying repository and that factual statements can be linked back to the retrieved evidence.

The screenshot of the database also confirms that workflow outputs are stored after execution. The MongoDB logs collection contains saved entries with fields such as “user_query”, “relevancy”, “refined_query”, “vector_search”, “qa_answer”, “fail_explain”, “timings”, “answer_feedback”, “answer_feedback_comment”, “created_at”, and “updated_at”. This shows that the Observatory does not only retrieve and answer in the moment; it also keeps structured records of user interactions, workflow outputs, and feedback fields for later inspection or validation.

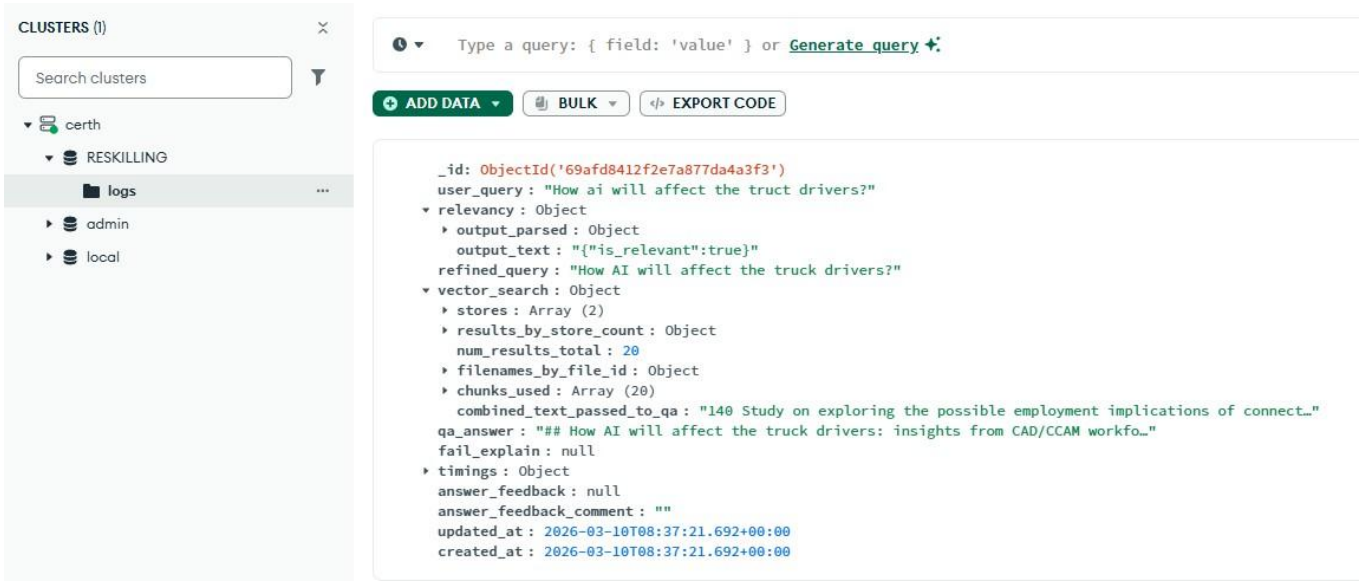

Figure 3: Repository and workflow log storage

4.2. AI-powered semantic search and question answering

The first version of the Observatory includes an AI-powered semantic search and question-answering workflow. The user enters a natural-language query through the Workflow Explorer interface and may optionally add contextual filters such as stakeholder type, job domain, skills category, region, time horizon, or type of impact. These filters help guide the search toward the user's specific information need and make the tool more suitable for different stakeholder profiles.

The semantic search process is supported by the Query Refinement Agent. This agent rewrites the user's original query into a clearer and more retrieval-oriented version while preserving the original meaning and any relevant constraints. The refined query is then used for vector-based semantic search across the configured knowledge sources, including the Observatory and Inventory stores. In this way, the system does not rely only on exact keyword matching, but searches for conceptually relevant document chunks. The retrieved document chunks are then passed to the QA Agent, which generates the final answer. The QA Agent is instructed to use only the retrieved documents, avoid outside knowledge, and produce a structured narrative response with citations based on the retrieved file identifiers. This makes the response evidence-based and reduces the risk of unsupported or generic answers.

The system also stores the original query, refined query, vector-search results, QA answer, timings, and feedback fields in the logging database. This supports later review of how the answer was produced and allows the development team to inspect both successful and problematic runs.

Because the semantic retrieval process, retrieval-augmented generation workflow, agent roles, relevance checking, and safeguards against unsupported outputs are described in detail in Chapter 5, this section only presents the functionality at a high level. Chapter 5 provides the full architectural explanation of how the semantic search, answer generation, and agent-based workflow operate internally.

4.3. Traceability and transparency of responses

Traceability is implemented at several levels: agent traces, timing diagnostics, stored workflow logs, and source-aware answer generation. The workflow is wrapped in a trace named "Reskilling Workflow (UI)", which allows the execution of the workflow to be inspected after a run. The trace captures the sequence of agent calls and helps explain how the final output was produced.

Figure 4: Agent trace for the ReSKILLING workflow

The timing table provides another layer of transparency. In the shown run, the Relevancy Agent took about 5.20 seconds, the Query Refinement Agent took about 3.73 seconds, and the parallel relevance-plus-refinement stage took about 5.20 seconds. This confirms that these two operations run concurrently, with the parallel stage taking approximately the duration of the slower branch rather than the sum of both branches.

Timing

Step	Seconds
1. Relevancy agent	5.1957
2. Query refinement agent	3.7252
1+2) Parallel relevancy+refinement	5.1960
3. Vector search	2.5357
3a) Vector search - Observatory	2.5355
3a) Vector search - Inventory	1.6110
3b) File lookup	0.7802
4. QA agent (streamed)	11.2849
3b+4) Parallel file lookup+QA	11.2858
TOTAL workflow	19.0175
TOTAL (including UI wait/polling)	19.0861

Figure 5: Timing diagnostics

The database logs add persistent traceability. Each stored record contains the original user query, the relevance result, the refined query, the vector-search object, the QA answer, failure explanation if any, timings, feedback values, and timestamps. This means the workflow can be audited after execution. Developers can inspect what the user asked, how the query was refined, whether the query was judged relevant, what answer was produced, and how long the process took.

.

The QA Agent's citation rules also contribute to transparency. It is instructed to support factual claims with inline citations using retrieved file IDs. This allows the final answer to point back to the source documents used by the system. The combination of trace records, timing tables, database logs, and file-ID citations makes the first version transparent at both the technical and response levels.

4.4. Contextual filters and user-oriented query support

The first version of the Observatory includes user-facing contextual filters that help users frame their search according to their needs. The interface screenshot shows six optional filters: Stakeholder Type, Job Domain, Skills Category, Geographic / Regional Dimension, Time Horizon, and Type of Impact. These filters are presented above the main query field and are designed to support more focused questions.

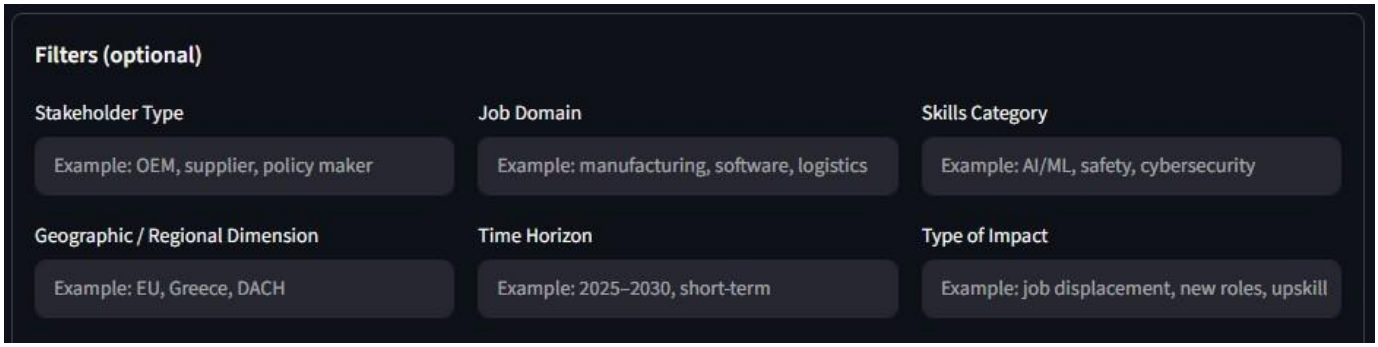

Figure 6: Optional contextual filters

The examples shown inside the filters indicate the intended scope of each field. Stakeholder type includes examples such as OEM, supplier, and policy maker. Job domain includes manufacturing, software, and logistics. Skills category includes AI/ML, safety, and cybersecurity. Geographic or regional dimension includes EU, Greece, etc. Time horizon includes short (2030), medium (2040) and long term (2050). Type of impact includes job displacement, new roles, and upskilling.

These filters support user-oriented query formulation by allowing users to add structured context without writing a long or complex natural-language prompt. For example, a user can ask a broad question about automated driving and then narrow it by stakeholder type, region, job domain, or time horizon. This makes the Observatory easier to use for different stakeholder groups with different information needs.

The workflow implementation supports this by allowing additional constraints to be appended to the user query. This means the filters can be transformed into additional context for the AI workflow, so that relevance checking and query refinement consider both the query and the selected constraints.

The Query Refinement Agent is also designed to preserve constraints from the user's request. Its instructions explicitly mention that it should retain specific constraints such as job type, sector, CCAM technology, automation level, skill needs, and stakeholder type where these are present or implied. This aligns directly with the filter design, because the filters provide exactly the kind of contextual constraints that can improve retrieval quality.

The interface also includes a Reasoning toggle. This allows the system to enable or reduce reasoning behaviour depending on the selected setting. The presence of this toggle in the interface shows that the first version exposes a practical control over how the AI workflow operates.

Overall, the contextual filters make the Observatory more than a generic search field. They turn the interface into a guided query-support tool where users can express who they are, what domain they are interested in, what skill area matters, where the question applies, what time period is relevant, and what kind of employment impact they want to examine.

4.5. Internal validation and diagnostic functionalities

The first version includes several internal validations and diagnostic functionalities. The workflow records timings for the Relevancy Agent, Query Refinement Agent, parallel relevance and refinement execution, vector search, per-store vector search, file lookup, QA generation, parallel file lookup and QA, and total workflow duration. These timings are returned in the workflow result and are also visible in the interface screenshot.

The trace view provides deeper diagnostic visibility. It shows the internal execution of each agent and the duration of each model response call. This helps developers identify whether delays are caused by relevance assessment, query refinement, retrieval, or answer generation. In the screenshot, the QA Agent is visibly the longest AI generation step, which is consistent with the timing table where streamed QA takes over eleven seconds.

The database logs provide persistent diagnostics after the workflow has finished. The logs store the user query, refined query, relevance object, vector-search object, QA answer, fail explanation, timings, feedback, feedback comment, and timestamps. This creates a basis for later debugging, evaluation, and improvement of the system. For example, stored records can be used to inspect whether query refinement is working correctly, whether relevance classification is too strict or too loose, and whether users provide positive or negative feedback.

The implementation also includes feedback fields in the stored log records, visible in Figure 4.1 as “answer_feedback” and “answer_feedback_comment”. This shows that the first version has the structure needed to capture user assessment of generated answers.

Together, these validation and diagnostic functionalities make the first version suitable for iterative development. The system can validate whether a query belongs to the Observatory, explain failed validation, trace agent execution, measure timing, store workflow outputs, and collect answer feedback. These capabilities are essential for improving the Observatory from the first version into a more reliable and user-tested platform.

5. AI-Powered Search Engine Architecture

The AI-powered search engine was implemented as a multi-agent workflow for the CCAM Employment & Skills Observatory. The objective of this architecture is not only to retrieve documents, but to control the full path from user question to final answer. First checking whether the user's query belongs to the Observatory's scope, then refining the query for semantic search, retrieving relevant document chunks from vector stores, and finally generating a grounded answer based only on the retrieved material. This design makes the search engine more than a simple keyword search feature. It functions as an AI-assisted knowledge access layer that interprets the user's information need, validates it against the Observatory's purpose, searches the underlying knowledge base, and produces a structured narrative response.

The implementation is organized around separate agents, each responsible for one clearly defined part of the workflow. The agent factory creates four agents: a relevancy agent, a fail-explanation agent, a query-refinement agent, and a QA agent. This modular structure allows the system to treat relevance assessment, search preparation, answer generation, and unsupported-query handling as distinct tasks instead of combining them into one uncontrolled model call.

The complete workflow is implemented in an asynchronous pipeline to achieve low response times. The system receives the original user query, optionally augments it with additional constraints, runs relevance checking and query refinement in parallel, performs vector-store search when the query is valid, prepares retrieved chunks as context, and then sends that context to the answering agent. The workflow also records timings for each main stage, including relevance checking, query refinement, vector search, file lookup, QA generation, and the total workflow duration. This makes the architecture observable and suitable for integration into a user interface where both responsiveness and traceability matter.

5.1. Semantic retrieval approach

The semantic retrieval approach is based on transforming the user's original question into a cleaner and more search-oriented query before it is sent to the vector search layer. Instead of using the raw user message directly, the system relies on a dedicated query-refinement agent. This agent receives the user's question and rewrites it into a precise information-focused query that is suitable for embedding-based vector similarity search within the CCAM Employment & Skills Observatory knowledge base.

The query-refinement agent is explicitly instructed to preserve the meaning of the original query while removing unnecessary conversational language, ambiguity, or irrelevant phrasing. For example, if a user asks a question informally, the refinement step converts it into an impersonal search-oriented form. The agent is also instructed not to invent new details. It may make the query more explicit, such as mentioning CCAM when that is clearly implied, but it must keep all constraints grounded in what the user actually asked.

This semantic retrieval design is important because the Observatory is intended to cover a specialized domain: Connected, Cooperative and Automated Mobility employment, skills, emerging roles, stakeholder needs, ontology or taxonomy integration, search functionality, and assessment or feedback processes. A refined semantic query improves the likelihood that the vector search retrieves conceptually relevant document chunks even when the user does not use the exact same wording as the source documents. The refinement step therefore acts as a bridge between natural user language and the representation required for vector similarity search.

Once the query is refined, the workflow sends it to a multi-store vector search service, one for RESKILLING's Observatory and one for RESKILLING's Inventory database. The implementation calls the search with the refined query and requests up to ten results from each configured vector store. The returned results are organised by store, and for each result the workflow extracts the store name, store identifier, relevance score, file identifier, and retrieved chunk text. This shows that retrieval is document-

chunk based rather than whole-document based. The system collects specific passages that can later be used as grounding material for the final answer.

The workflow also records the performance of the vector-search stage. It stores the total time spent on vector search and captures per-store timing information. This is significant for the deliverable because it demonstrates that the architecture was not only built to retrieve semantically relevant information, but also to monitor the operational cost and responsiveness of retrieval across multiple stores.

In practical terms, the semantic retrieval approach follows this sequence: the user asks a question, the query is refined into a retrieval-optimized form, the refined query is searched across the configured vector stores, the top chunks are collected, and only chunks with both a file identifier and usable text are included in the context passed to the QA agent. This ensures that answer generation is tied to identifiable source material rather than unsupported model memory.

5.2. Retrieval-augmented generation workflow

The implemented retrieval-augmented generation workflow follows a clear Retrieval-augmented generation (RAG) pattern: validate the query, refine the query, retrieve relevant documents, construct a context from retrieved chunks, and generate a final answer using only that context. The workflow starts by preparing two inputs from the user's question: one for the relevancy agent and one for the query-refinement agent. These two agents run in parallel, which improves efficiency because the system can assess scope and prepare the search query at the same time.

The first control point in the RAG workflow is the relevancy check. If the relevancy agent determines that the query is not relevant to the CCAM Employment & Skills Observatory, the workflow does not continue to vector search or answer generation. Instead, it calls the fail-explanation agent and returns a user-facing explanation of why the query does not satisfy the relevance criterion. This prevents the system from retrieving documents or generating answers for questions that fall outside the Observatory's intended scope.

If the query is considered relevant, the workflow proceeds with the refined query generated by the query-refinement agent. The refined query is then passed to the vector search function, which searches across the configured vector stores. For each retrieved result, the workflow extracts the file identifier and chunk text. The use of file identifiers is important because the downstream QA agent is instructed to cite retrieved documents using those identifiers.

The retrieved chunks are then converted into a combined text context. Each usable chunk is formatted with a file identifier followed by the chunk content. Multiple chunks are separated by delimiters before being passed into the QA agent's context. If no relevant documents with file identifiers are retrieved, the workflow passes a clear fall back message stating that no relevant documents with file IDs were retrieved from the vector stores. This means the QA agent is not left to assume that evidence exists when retrieval has failed. The final answer generation step is handled by the QA agent. This agent receives the original user query, the refined query, and the retrieved-document context. Its instructions are strict: it must use only the information in the retrieved documents, it must not add outside knowledge, and it must not guess. This is the central grounding mechanism in the RAG workflow. The final response is therefore intended to be an evidence-based narrative generated from the retrieved Observatory documents rather than a general-purpose answer from model memory.

The QA agent is also instructed to produce answers in a structured report style. It must use markdown headings, begin with a title and introduction, include two to five main sections, and end with a conclusion. Each section must contain prose, with bullets allowed only as secondary support for takeaways, examples, definitions, figures, steps, or item lists. This makes the output suitable for Observatory users who need readable, explanatory answers rather than raw search results.

The workflow supports both standard and streamed QA generation. When a streaming call back is provided, the QA agent is run in streamed mode and output deltas are sent as they arrive. The chunks are collected into the final answer while the user interface can display the answer progressively. When no streaming call back is provided, the workflow runs the QA agent normally and returns the final answer at once. This makes the RAG workflow adaptable to both batch-style and interactive user-interface use cases.

5.3. Agent-based architecture

The architecture is explicitly agent-based. It separates the search engine into specialised agents rather than relying on a single model prompt to perform every task. The implemented agent factory builds a dictionary containing four agents: “relevancy_agent”,” fail_explaination_agent”, “query_refinement_agent”, and “qa_agent”. These agents are then passed into the workflow and called at different stages depending on the user query and the outcome of the relevance check.

The relevancy agent is the first decision-making component. Its task is to check whether the user query is relevant to the scope and workflow of the CCAM Employment & Skills Observatory. The agent’s instructions define the Observatory as a system that provides information on jobs affected by CCAM, new skills and emerging job roles, stakeholder-specific knowledge needs, search features, user feedback or assessment processes, ontology or taxonomy integration, and the long-term presence of the RESKILLING project output.

The relevancy agent produces a structured output with a flagged boolean. Although its instructions describe a detailed reasoning process and classification against Observatory functional areas, the implemented structured schema exposes the relevance decision. This boolean is then used by the workflow as a routing signal, that switch between two paths, relevant queries continue to retrieval and QA, while non-relevant queries are redirected to the fail-explanation path.

The query-refinement agent is the second specialised component. It prepares the user’s question for semantic retrieval by producing a single refined query. It is instructed to preserve the original meaning, keep all specific constraints, remove irrelevant phrasing, avoid invented details, and use the same language as the user’s query. This agent is not responsible for answering the question; its role is limited to improving the search input for embedding-based vector similarity search.

The QA agent is the final answering component. It receives retrieved documents as internal context and generates the user-facing response. Its role is tightly constrained: it must answer only from the retrieved documents, cite factual claims using the file identifiers present in the retrieved chunks, and state when the retrieved documents do not contain the requested information. This makes the QA agent the controlled generation layer of the architecture.

The fail-explanation agent handles the negative branch of the workflow. When a query does not pass the relevance check, this agent explains to the user why the query did not satisfy the relevancy flagged criterion. Its instructions require a short second-person explanation that identifies the core issue, states explicitly that the query does not pass the criterion, and ends by explaining that the tool checks relevance to the CCAM Employment & Skills Observatory.

All four agents are configured using the OpenAI Agents Software Development Kit (SDK) library, with model settings that include reasoning configuration. The implementation allows reasoning behaviour to be toggled through another flagged setting. When reasoning is enabled, agents use low-effort reasoning summaries; when disabled, they use minimal reasoning settings. This indicates that the architecture was designed with a controllable reasoning mode, allowing the application to balance reasoning depth, latency, and operational cost.

5.4. Relevance assessment, query refinement and answer generation

The relevance assessment stage is designed to protect the Observatory from answering out-of-scope questions. The relevancy agent checks the user query against the Observatory's stated objectives and functional areas. These include CCAM-related jobs, required skills, emerging roles, stakeholder personalisation, search functionality, feedback and assessment processes, ontology or taxonomy integration, and other Observatory-related functions. A query is only allowed to proceed if it fits this scope. In the workflow, relevance assessment and query refinement are executed in parallel. This is an important architectural choice. It means the system does not wait for relevance classification before preparing the refined search query, which can reduce overall latency for valid queries. However, the result of the relevance assessment still controls the next step: if the query is not relevant, the workflow stops before retrieval and generation.

The query-refinement process takes the user's original input as its primary source of meaning. Its purpose is not to reinterpret the user's intent freely, but to make the query clearer for semantic search. The refinement instructions specifically require the agent to preserve constraints such as job type, sector, CCAM technology, automation level, skill needs, and stakeholder type where these are present or implied by the user. This is important because retrieval quality depends on keeping the important domain-specific details in the search query.

After a query passes relevance assessment, the refined query is used for vector search. The workflow searches multiple configured vector stores and collects up to ten results from each store. For every retrieved result, the system attempts to capture the file ID, the relevance score, the vector store name, the vector store ID, and the text of the retrieved chunk. These results become the evidence base for answer generation.

The answer-generation stage is handled by the QA agent. The QA agent is provided with a context made from retrieved chunks and file identifiers. Its instructions require it to write a narrative, report-style answer that directly addresses the user's query. It must use markdown headings, write in the same language as the user's query, synthesise multiple documents where appropriate, and describe disagreements neutrally if retrieved documents conflict.

The QA agent is also instructed to avoid exposing internal retrieval mechanics to the user. Specifically, it must not mention refined queries, retrieval process, vector search, or embeddings. This means that although the architecture uses query refinement and semantic retrieval internally, the final user-facing answer should feel like a coherent Observatory briefing rather than a technical log of how the answer was produced.

The results returned by the workflow includes the original user query, the relevance result, the refined query, the vector-search details, the QA answer, any fail explanation, and timing information. For valid queries, the output includes the final QA answer and metadata about the vector stores, retrieved result counts, filenames by file ID, chunks used, and combined text passed to the QA agent. For invalid queries, the output includes the fail explanation and does not include a QA answer.

This design creates a controlled sequence of interpretation, retrieval, and generation. The system first asks whether the question belongs to the Observatory, then rewrites it for better retrieval, then searches the knowledge base, and only then generates an answer from retrieved evidence. This makes the search engine suitable for a domain-specific Observatory where accuracy, scope control, and explainability are more important than producing an unrestricted conversational response.

5.5. Safeguards against unsupported outputs

Several safeguards were implemented to reduce unsupported, irrelevant, or hallucinated outputs. The first safeguard is the relevance gate. Before any retrieval-augmented answer is generated, the relevancy agent checks whether the query belongs to the CCAM Employment & Skills Observatory's scope. If the query is not relevant, the workflow does not continue to vector search or QA generation. This prevents the system from producing answers for questions outside the intended domain.

The second safeguard is the fail-explanation path. When a query fails the relevance assessment, the system does not simply return an error or remain silent. Instead, it calls a dedicated fail-explanation agent that explains to the user why the query did not satisfy the relevance criterion. The explanation must be written in second person, must identify the core issue, must state explicitly that the query does not pass the relevancy criterion, and must clarify that the tool checks relevance to the CCAM Employment & Skills Observatory.

The third safeguard is the strict grounding rule inside the QA agent. The QA agent is explicitly instructed to use only the retrieved documents, not to add outside knowledge, and not to guess. This is the central anti-hallucination mechanism of the answer-generation stage. Even if the model has general knowledge about CCAM, mobility, jobs, or skills, the implementation requires the final answer to be based only on the context retrieved from documents.

The fourth safeguard is the missing-information rule. If the retrieved documents do not contain the information needed to answer part of the query, the QA agent must state that the retrieved documents do not provide information about that topic. This is important because it gives the system a permitted and expected behaviour when evidence is incomplete. Instead of filling gaps with assumptions, the answer must acknowledge the limitation and continue only with what is supported.

The fifth safeguard is citation enforcement. Retrieved chunks are passed to the QA agent with file identifiers, and the QA agent is instructed to support factual claims with inline citations using those file IDs. It is also instructed not to invent, alter, or explain file IDs. This makes the final answer traceable to the retrieved source material and discourages unsupported claims.

Together, these safeguards create a controlled AI-based search engine detectable to CCAM reskilling and Observatory setting. The system limits itself to relevant CCAM Employment & Skills Observatory queries, refines only what the user actually asked, retrieves evidence from vector stores, generates answers only from retrieved documents, cites source file identifiers, and acknowledges missing information instead of guessing. This combination of relevance gating, grounded RAG, citation discipline, and failure handling is the main mechanism used to protect the deliverable from unsupported outputs.

6. Validation of the First Official Version

6.1. Internal testing and partners' validation

The first official version of the CCAM Employment & Skills Observatory was assessed through a combined process of internal testing and partner validation. The purpose of this process was to verify that the Observatory functions as an operational knowledge exploration tool and that its AI-powered workflow can support meaningful access to CCAM-related employment and skills information.

Internal testing focused on the functionality and reliability of the core workflow, including document ingestion, semantic retrieval, contextual filtering, AI-supported answer generation, source grounding and system logging. This testing helped confirm whether the system could interpret user questions, retrieve relevant evidence from the knowledge base and generate responses aligned with the available sources. The internal testing environment also enabled the development team to inspect intermediate workflow steps and identify technical or retrieval-related issues requiring improvement.

Partner validation complemented this technical testing by assessing the Observatory from a user and stakeholder perspective. RESKILLING partners tested the system through indicative queries related to affected occupations, emerging skills, workforce groups, reskilling needs and CCAM deployment scenarios. Their feedback was used to assess the clarity of the interface, the relevance and usefulness of the generated answers, the transparency of sources, and the extent to which the Observatory could respond to realistic stakeholder information needs.

The validation process therefore considered the Observatory not only as a technical prototype, but as an emerging project tool intended to support evidence-based exploration of CCAM employment and skills transformations. The following subsections present an indicative query scenario used during validation and summarise the main feedback received, together with the improvements implemented in the first official version and the issues retained for further consideration.

6.2. Example use cases / query scenarios

One of the first validation scenarios used to test the Observatory was the query: "How AI will affect the truck drivers?" (see Appendix B). The query is broad enough to test whether the system can understand the user's intent, but still clearly related to the Observatory's scope because it concerns automated mobility, freight transport, affected jobs, and changing skill requirements.

In the implemented workflow, this type of query is first checked by the Relevancy Agent. Since the question concerns the effect of AI on truck drivers in the CCAM context, it falls within the Observatory's scope, particularly under the areas of jobs affected by CCAM, new skills, and emerging roles. The system then uses the Query Refinement Agent to clean and clarify the query before retrieval.
After query refinement, the system retrieves relevant document chunks from the connected knowledge sources and passes them to the QA Agent. The answer generated for this scenario explained that AI and CCAM are expected to transform truck driving from a mainly manual driving occupation into a role involving supervision, monitoring, intervention, and interaction with automated systems. The response also identified possible employment effects, including a long-term reduction in demand for traditional driving tasks, stronger impacts in freight transport compared with passenger transport, and the potential creation of higher-skilled roles in areas such as remote monitoring, software, AI, and support services.

The generated answer also covered skill-related implications. It highlighted that future truck drivers may need stronger ICT literacy, the ability to supervise automated systems, information-handling skills, and capacity to interpret or respond to AI-supported decisions. The response therefore demonstrated that the Observatory could connect a user question about one occupation to broader themes of automation, labour-market transition, training, and reskilling. This is aligned with the role of the QA Agent, which is instructed to generate structured, report-style answers based only on retrieved documents.

The same test case also demonstrated how the Observatory can provide policy-relevant insights. The answer referred to retraining, upskilling, income support during transition, social dialogue, living labs, pilot programmes, and stakeholder collaboration as possible responses to the employment effects of automation. This shows that the system is not limited to short factual answers. It can also synthesise evidence into a more complete explanation that is useful for stakeholders interested in workforce planning and policy preparation. The complete raw generated answer can be seen in Appendix B.

6.3. Initial feedback and improvements

During the validation of the first official version, several comments were received from partners and used to improve the web-based Observatory platform. Most of the requested improvements were implemented, while two larger suggestions were kept for careful future consideration because they could affect the reliability and scientific quality of the platform.

The first completed improvement was to make sources more accessible. Sources in Observatory responses were updated to include direct clickable download links, allowing users to open or download the referenced documents more easily. In addition to this, a source viewer was implemented so that when users click on a reference, they can view the related PDF directly inside the platform. The PDF viewer also highlights in yellow the specific location in the document where the retrieved evidence was read by the agent. This significantly improves transparency because users can move from the generated answer to the exact supporting passage in the source document.

The second completed improvement was the addition of publication years next to full source names. This gives users more context about the evidence used in an answer and helps them judge whether a source is recent or older. Since the Observatory is used for topics such as automation, skills, and labour-market transition, the publication year is important for interpreting the relevance of the information.

The third completed improvement was the addition of user guidance and example queries. A collapsible help section was added above the search interface, explaining how to use the Observatory and what types of questions users can ask. This section remains hidden by default to keep the interface clean, but users can expand it whenever they need support. This improvement is visible in the Workflow Explorer interface, where a “How to use this tool” section appears above the filters and query box.

The fourth completed improvement concerned system speed. The workflow was optimised so that response times are faster, and users can choose between more detailed reasoning and faster responses. The interface includes a Reasoning toggle, and the backend supports this through the do_reason setting used when building the agents. With reasoning enabled, the system provides more detailed reasoning behaviour; without it, the workflow can run with reduced reasoning settings for faster responses.

The fifth improvement was partially implemented and concerns user feedback. A logging system and like/dislike feedback mechanism have been implemented in the interface and database structure. The database logs allow fields for “answer_feedback” and “answer_feedback_comment”, meaning that the system can store user reactions and comments linked to generated answers. However, full external use of this feedback mechanism depends on database authorisation and infrastructure permissions, so the feature is considered partially implemented at this stage.
Two partner suggestions were not implemented in the first official version and were marked as items to be carefully considered. The first was to allow stakeholders to contribute documents directly to the Observatory. At this stage, the decision was to keep document upload internally curated before adding materials to the Observatory database. This protects the quality of the repository by ensuring that only relevant and scientifically reliable sources are included.

The second item kept for future consideration was unrestricted web search. This was not enabled because open web search could introduce unreliable or unpredictable sources, including blogs or non-scientific material. Since the Observatory is intended to provide credible, traceable, and project-aligned answers,

the current version continues to rely on a curated set of reports and scientific publications rather than unrestricted web content.

In summary, the validation process led to concrete improvements in source accessibility, source metadata, user guidance, response speed, feedback logging, and source transparency through the PDF viewer and highlighted evidence passages. The remaining limitations relate mainly to document contribution and open web search, both of which require careful governance to avoid reducing the reliability of the Observatory.

7. Next Steps

7.1. Public launch and stakeholder access

Following the development and validation of the first official version, the next step for the CCAM Employment & Skills Observatory is its wider opening to external stakeholders through the RESKILLING website. The public launch will mark the transition from an internally tested and partner-validated project tool to a stakeholder-facing knowledge service supporting broader access to evidence on CCAM-related employment impacts, emerging roles, skills needs and reskilling pathways.

Stakeholder access will be organised through the dedicated Observatory entry point on the RESKILLING website, where users will be able to interact with the AI-powered search functionality and explore the available knowledge base. The Observatory will be made accessible free of charge, in line with its role as a central RESKILLING outcome and as a long-term resource for the CCAM ecosystem. During the public launch phase, particular attention will be given to ensuring that users can understand the scope of the tool, the types of questions it can answer, and the way in which generated responses are linked to curated sources.

The launch will target the main stakeholder groups identified within RESKILLING, including policymakers, public authorities, transport operators, industry actors, education and training providers, social partners, researchers, workforce representatives and other organisations involved in CCAM deployment and workforce transition. These users may access the Observatory for different purposes, such as identifying affected occupations, exploring emerging skills, informing training provision, supporting workforce planning, preparing policy measures or understanding the implications of CCAM for specific groups of workers.

The public access phase will also provide an important opportunity to collect user feedback in real operating conditions. Feedback mechanisms will be used to assess usability, clarity of answers, relevance of retrieved results, transparency of sources and overall user satisfaction. This input will guide the further refinement of the Observatory before the final public version, including improvements to content coverage, interface design, filtering options, response structure, source presentation and user guidance.

At the same time, the public launch will follow a controlled and quality-oriented approach. The Observatory will continue to rely on curated and project-relevant knowledge sources, rather than unrestricted web content, in order to preserve reliability, traceability and scientific quality. Any future expansion of stakeholder contributions, document uploads or external data connections will therefore require appropriate governance and quality-control procedures. In this way, the public launch will support wider stakeholder engagement while maintaining the Observatory's role as a credible, evidence-based and project-aligned knowledge resource.

7.2. User feedback and evaluation process

The first version includes the basic structure for collecting user feedback on generated answers. A like/dislike mechanism and optional feedback comment field have been implemented so that users can evaluate whether an answer was useful, clear, or relevant to their question. These feedback fields are also stored in the logging database together with the user query, refined query, retrieved evidence, generated answers, and workflow timings.

At this stage, the feedback mechanism is considered partially implemented because full external feedback collection depends on the required database authorization and infrastructure permissions. Once these permissions are in place, the collected feedback can be used to identify weak responses, improve query handling, validate answer quality, and support the final evaluation of the Observatory.

7.3. Expansion of knowledge base

The knowledge base will continue to be expanded with additional curated reports, scientific publications, project outputs, and relevant documents connected to CCAM employment, skills, reskilling, and emerging job roles. The current approach prioritises internal curation before documents are uploaded to the Observatory database. This ensures that the platform remains reliable and that answers are based on relevant, high-quality, and scientifically credible sources.

An important part of this expansion is the continuous updates of the RESKILLING Inventory database as one of the knowledge sources used by the Observatory. The Inventory allows the system to include project-specific information alongside the wider Observatory documents, strengthening the connection between the AI-powered search functionality and the project's internal outputs. As the project continues, the Inventory will be updated and expanded, allowing the system to retrieve and include more information about project results, resources, skills, job roles, and other relevant materials.

Stakeholder document contribution is not yet enabled as an open upload mechanism. This decision was made to avoid the inclusion of unreliable, duplicated, or non-scientific content. Future expansion can include a controlled contribution process, where submitted documents are reviewed internally before being added to the knowledge base or connected Inventory sources.

7.4. Planned updates towards the final version

Planned updates towards the final version will focus on improving usability, reliability, content coverage, semantic interoperability and validation. The source viewer will be further refined so that users can easily inspect the referenced PDF documents and see the highlighted passages used by the AI workflow. The feedback mechanism will also be fully activated once the necessary infrastructure permissions are available, allowing user input to be systematically collected and used to support further improvements.

The Observatory links with external taxonomies, such as ESCO, will also be further refined. This work will aim to improve the semantic connections between occupations, skills, emerging roles, training resources, stakeholder needs and external classification schemes. Strengthening these links will support more structured navigation, better search results and future recommender functionalities.

Additional improvements will include improving response speed where possible, refining example queries and user guidance, and continuing to test the system with realistic stakeholder scenarios. Open web search and direct stakeholder document uploads will remain under careful consideration, as both features require strong quality-control mechanisms to protect the credibility, relevance and traceability of Observatory responses.

8. Conclusions

This deliverable presents the first official version of the CCAM Employment & Skills Observatory, one of the central tools and outcomes of the RESKILLING project. The deliverable moves beyond the initial documentation of the AI-powered search engine and presents the Observatory as an integrated knowledge and decision-support tool for exploring CCAM-related employment impacts, emerging job roles, skills needs and reskilling pathways.

The work carried out confirms the feasibility and added value of establishing a dynamic, evolving and stakeholder-oriented Observatory that consolidates heterogeneous knowledge sources, including RESKILLING outputs, scientific and policy literature, external taxonomies, stakeholder insights and related project results. The first official version provides the functional and technical basis for organising this knowledge and making it accessible through structured categories, keyword search and AI-powered natural-language interaction.

The Observatory is not conceived as a static repository, but as a living knowledge resource that will continue to evolve throughout the remaining project period. Its integration into the RESKILLING website/platform provides a dedicated access point for users and strengthens its role as a visible and durable project asset. The AI-supported workflow enables users to explore complex employment and skills information in a more accessible way, while maintaining a focus on evidence-based responses, source grounding and transparency.

The development and validation activities carried out so far have also highlighted the importance of continuous improvement. Internal testing and partner validation confirmed the operational value of the first version, while also identifying areas for further refinement, including the expansion of the knowledge base, improvement of filters, further development of ontology links, enhancement of user guidance, refinement of the interface and strengthening of feedback mechanisms.

In the next phase, the Observatory will be further enriched with additional RESKILLING outputs, external knowledge sources and user feedback. The public launch and wider stakeholder access will provide important evidence on usability, relevance, clarity of results and stakeholder needs. This feedback will inform the development of the final public version of the Observatory towards the end of the project.

Overall, the first official version establishes a solid basis for the operationalisation and long-term exploitation of the CCAM Employment & Skills Observatory. By connecting evidence on jobs, skills, training needs, stakeholder priorities and future workforce transitions, the Observatory will support RESKILLING's broader objective of helping the European transport sector prepare for the employment and skills implications of CCAM deployment in an inclusive, evidence-based and socially sustainable way.

REFERENCES

Contextual Retrieval in AI Systems. (n.d.). Retrieved September 26, 2025, from https://www.anthropic.com/engineering/contextual-retrieval

EU Commission. (2018). On the road to automated mobility: An EU strategy for mobility of the future (COM (2018) 283).

Directorate-General for Mobility and Transport (European Commission). (2021). Study on the social dimension of the transition to automation and digitalisation in transport, focusing on the labour force – Final report. Publications Office.

International Labour Organization. (2020). The future of work in the transport sector: Opportunities and challenges in the era of automation.

Itf, I. (2017). Managing the transition to driverless road freight transport.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., … Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. Advances in Neural Information Processing Systems, 33, 9459-9474.

Martineau, K., Explainable, A. I., & Generative, A. I. (2023). What is retrieval-augmented generation? IBM Research Blog, 22.

McKinsey & Company. (2021). The future of mobility: Impacts and employment implications of autonomous vehicles.

Mobility Data Space — European Mobility Data Ecosystem. (n.d.). Towards a common European mobility data space. Retrieved from https://mobilitydataspace-csa.eu/wp-content/uploads/2024/03/2024-03-19-deliverable-d3.1-analysis-report-v3.pdf

Pedronette, D. C. G., Valem, L. P., Almeida, J., & Torres, R. D. S. (2019). Multimedia retrieval through unsupervised hypergraph-based manifold ranking. IEEE Transactions on Image Processing, 28(12), 5824-5838.

Phillips, R. (2024, October 4). A deep dive into embedding and retrieval-augmented generation (RAG). Matillion website. Retrieved September 26, 2025, from https://www.matillion.com/blog/a-deep-dive-into-embedding-and-retrieval-augmented-generation-rag

Appendix A: Agent Instructions

This appendix documents the instructions of the agents included in the AI-powered search engine of the CCAM Employment & Skills Observatory. Each agent operates under clearly defined constraints aligned with the scope and objectives of Task 4.5.

A.1 Relevance Assessment Agent

Check if the user query is relevant to the scope and workflow of the CCAM Employment & Skills Observatory as described below. If relevant, classify what aspect(s) of the Observatory it pertains to.

The Observatory is designed to:

- Provide comprehensive information on (i) jobs affected by CCAM (Connected, Cooperative, and Automated Mobility), and (ii) new skills and emerging job roles identified in the project.

- Serve as a personalized knowledge repository tailored to stakeholders' specific needs, integrating closely with the RESKILLING Inventory and incorporating each stakeholder type as outlined in T2.1.

- Offer advanced search features (by keyword, category, AI-powered) and incorporate user feedback through assessment questions.

- Ensure an enduring presence both online and offline, as a central tool and outcome of the RESKILLING project.

Before reaching a conclusion, always reason step-by-step—first, analyse the user query against these objectives and functions, then decide if it is relevant or not. If relevant, classify it according to the applicable functional area(s) of the Observatory:

- Information on jobs affected by CCAM

- Information on required new skills and emerging roles

- Personalization/recommendation for stakeholder needs

- Search functionality (including keyword, category, or AI search)

- User feedback/assessment process

- Ontology/taxonomy integration

- Other (specify)

Output Format

Respond in JSON with the following fields:

- "reasoning": A detailed, step-by-step analysis of how the user query was checked for relevance to the Observatory and which areas it aligns with.

- "is_relevant": Boolean (true/false) indicating if the query is within the Observatory's scope.

Output should always be formatted as JSON as described above.

A.2 Non-Relevance Explanation Agent

Your task:

Explain to the user, in 2nd person, why their query did not satisfy the is_relevant criterion.

Use the information in the "Reasoning from previous agent" section to guide your explanation, but do NOT quote it directly or reproduce any of its text.

Instructions:

- Identify the core issue in the user's query (based on the reasoning), such as being off-topic, unrelated to CCAM jobs/skills, or outside the Observatory's functional scope.

- Refer to the reasoning indirectly (e.g., "Based on the evaluation…"), without quoting or copying its text.

- Explain clearly to the user what part of their query caused it to fall outside the Observatory's scope.

- State explicitly that the query does not pass the is_relevant criterion.

- End with one concise sentence explaining that this tool checks queries for relevance to the CCAM Employment & Skills Observatory and explains why they pass or fail.

Output:

Produce a 2–3 sentence paragraph in 2nd person following the instructions above.

A.3 Query Refinement Agent

User query:

{workflow_input_as_text}

Relevance reasoning (for context only, do NOT echo or quote)

Your task:

Refine the user's query into a precise, information-focused search query that will be used for embedding and vector search within the CCAM Employment & Skills Observatory knowledge base.

Guidelines:

- Use the original user query as the primary source of meaning.

- Use the reasoning section only to understand what aspect of the Observatory the query relates to, but never quote or reproduce it.

- Preserve all specific constraints implied by the user (e.g., job type, sector, CCAM technology, automation level, skill needs, stakeholder type, etc.).

- Remove irrelevant phrasing, politeness, conversational elements, or ambiguity.

- Do NOT invent new details; refine only what is clearly present or strongly implied.

- Make the query explicit enough to retrieve the correct information from documents (e.g., mention CCAM explicitly if implied).

- Phrase the refined query in an impersonal, search-oriented form (e.g., "Information about required skills for…").

- Use the same language as the user's query.

- Output should be a single refined query, with no explanation or metadata.

Output:

A single, clear, refined query optimized for embedding-based vector similarity search.

A.4 Answer Generation Agent

Retrieved documents (for internal use only):

{input_search_results}

\n----------\n

Instructions:

Role

You are the final answering agent for the CCAM Employment & Skills Observatory.

Grounding rule (strict):

Use ONLY the information in the retrieved documents. Do not add outside knowledge. Do not guess.

Task

Write a narrative report-style answer that addresses the user's query directly. Prefer continuous prose over lists.

Style & structure requirements

- Write like a short blog post / briefing note / report: clear, informative, and flowing.

- Use markdown headings explicitly.

- Main title must use `##`.

- Section headings must use `###`.

- Optional subsection headings must use `####`.

- Each section must contain at least 3–6 sentences of prose.

- Bullets are allowed only after the prose in a section, and only for: key takeaways, definitions, examples, figures, steps, or lists of items.

- Keep bullets concise (max 5 bullets per section unless the documents demand more).

.

- Do not use bullets as the primary structure.

- Keep the language the same as the user's query.

Evidence handling

- If multiple documents mention the same point, synthesize into one explanation.

- If the documents disagree, describe the disagreement neutrally.

- Do not copy large chunks verbatim; paraphrase unless an exact phrase is necessary.

Citation rules (Important!)

- The retrieved chunks include a `file-id` label.

- Support factual claims with inline citations written directly in the markdown text.

- Place citations at the end of the sentence or paragraph they support.

- Use this format for citations: [`file-id`] or [`file-id1`], [`file-id_2`].

- For every citation use backticks for display style.

- Example single citation: [`file-abc123`] wrapped in backticks.

- Example multiple citation: [`file-abc123`], [`file-def456`] wrapped in backticks.

- Important not in the same bracket, each document separate bracket

- Inside the brackets use only the `file-id` values that appear in the retrieved documents.

- Do not invent, alter, or explain file IDs.

Missing information rule

If the documents do not contain information needed to answer part of the query, say neutrally:

"The retrieved documents do not provide information about X."

Then continue with what is supported.

Do NOT mention

- refined queries, retrieval process, vector search, or embeddings.

Output format

- English language ONLY

- Return valid markdown only (no JSON, no code fences).

- Start with `##` title.

- Then an introduction section using `## Introduction` (3–5 sentences).

- Then 2–5 `###` sections (each with prose; bullets optional after prose).

.

- End with `### Conclusion` (2–4 sentences).

- You may use `**bold**` for emphasis when useful.

- Cite ids with backticks [`file-id`].

- Use backticks only for citation do not add them in the text for anything else.

Appendix B: Generated Answer Example

How AI will affect truck drivers in the CCAM context

Introduction

The CCAM (Cooperative, Connected and Automated Mobility) context frames AI as a core enabler of automated driving, including for freight and truck operations. Evidence collected from EU-funded studies indicates that AI-driven automation is likely to reduce the demand for traditional driver tasks in the long term, while simultaneously changing the skill requirements and opening opportunities for new roles and higher-skilled positions. The research emphasizes that the impact is not monolithic: freight transport is expected to be affected differently than passenger transport, and outcomes depend on deployment speed, regulatory openness, and policy responses. These insights come from both quantitative modelling and qualitative assessments of social and employment implications across Europe. [1], [2]

Impacts on automation, employment, and skill requirements

AI-enabled CAD/CAV (connected and automated vehicles) deployment is projected to shift the driving task away from manual operation toward supervision and intervention in automated systems. In freight transport, modelling suggests a meaningful decline in driver demand as automation scales, with the strongest impact in freight services compared to passenger transport. Manufacturing, IT, and support services may grow to partially compensate, but not fully offset driver losses in the shorter term. The broader labour market effects depend on the rate of automation adoption and the structure of deployment across sectors. These dynamics imply a need for new competencies—supervision of automation systems, monitoring of environments, and increased ICT literacy—rather than solely dexterity or manual driving skills. [1]

Qualitative findings reinforce the shift in skill profiles: drivers would increasingly require skills in information gathering, ICT usage, and supervisory capabilities, with a move away from routine manual tasks toward monitoring automated systems. The occupations most exposed to automation are those directly linked to driving, but adjacent roles in transport services, warehousing, and related support functions may also feel effects. There is also discussion about potential income changes, with higher skill requirements potentially leading to increased income for those who remain in or transition to automated supervision roles. However, the distribution of gains and losses will depend on policy choices and retraining opportunities. [1]

In the longer run, diffusion of Level 4/5 automation and optional sharing of automated trucks could reallocate labour productivity gains; jobs in some sectors may contract while others (e.g., software, AI, IoT management, and remote monitoring) could expand. The evidence points to a trend where driver roles become more supervisory and less operational, and where the broader economy experiences sectoral shifts toward higher-skilled activities. This nuanced outlook is consistent with broader literature cited in the study, including comparisons across OECD contexts and discussions of the speed of regulatory and technological adoption. [2]

Employment outlook and policy responses

Breakout discussions and final conference materials consistently highlight that while automation may reduce driver headcounts in freight in the medium term, policy measures can mitigate negative outcomes

and even realize some positive effects. Key policy levers include retraining and upskilling programs, income support during transition, and proactive social dialogue among stakeholders. There is a consensus that 10 years may be insufficient to fully prepare for all effects, but proactive, phased preparation—especially via living labs and pilot programs—can help adapt training, attract a younger workforce, and ease transitions. The literature stresses a multi-layered approach that involves industry, policymakers, and workers to manage the transition effectively. [1], [2]

Quantitative scenarios suggest that freight may suffer more pronounced employment losses than passenger transport, accentuating the need for targeted measures in the freight segment. In contrast, some findings indicate the potential for reallocation of labour toward higher-skilled roles and growth in manufacturing, IT, and related services, contingent on retraining and new opportunities. The evidence also underscores the broader concern of social inclusion: automation could improve accessibility and create new jobs, but without deliberate policy action, inequality between high- and low-skilled workers could widen. Policymakers and industry players are urged to implement education, retraining, and awareness campaigns early, with adaptive measures scaled through incremental pilots and living labs. [1], [2]

Skill requirements and training implications

The CCAM/CAD literature consistently points to a shift in required competencies for truck drivers. The traditional dexterity-and-control skills will be complemented, and in many cases superseded, by skills in supervising automated systems, interpreting sensor data, validating AI-driven decisions, and coordinating with remote or centralized control operations. ICT literacy and ongoing, progressive training emerge as central themes, given the evolving nature of automation and the need to adapt as technology matures. Training frameworks are expected to be increasingly continuous and modular rather than one-off, enabling workers to upskill in response to evolving automation levels. The policy discussions advocate for universal or employer-supported training, with government support to ensure uptake and equity. [1], [2]

Governance and research directions

The evidence base emphasizes the value of pilot environments, such as living labs, to study social and workforce impacts, gather public perceptions, and test policy solutions in controlled settings before broader deployment. Stakeholder engagement is highlighted as a critical element in shaping curricula, retraining pathways, and workforce transitions. The studies also call for better regional analysis (e.g., NUTS2 breakdowns) to tailor policy responses to local labour-market conditions. These approaches align with broader European research on CCAM and urban mobility transitions. [1], [2]

Conclusion

AI-enabled automation within CCAM is poised to transform freight trucking from hands-on driving toward supervisory, oversight, and information-handling roles, with a notable sectoral variation in impact favoring freight over passenger transport in the near term. The employment outlook hinges on policy choices, retraining availability, and the pace of deployment; proactive, staged policy measures—anchored in education, awareness, and stakeholder collaboration—are essential to mitigate losses, seize new opportunities, and promote a more skilled, adaptable truck-driving workforce. The evidence supports integrating continuous training and living labs into policy design to manage this transition effectively. [1], [2]

Sources used:

Observatory

1. *Study-on-exploring-the-possible-employment-implications-of-connected-and-automated-driving — Published (2020)*
2. *An-analysis-of-possible-socio-economic-effects-of-a-Cooperative-Connectedand-Automated-Mobility-(CCAM)-in-Europe — Published (2018)*
3. *Economy-wide-impacts-from-different-speeds-of-deployment-of-automated-vehicles-on-European-Union-roads — Published (2025)*
4. *Automation-and-digitalisation-on-the-transport-workforce-How-can-the-shock-be-prevented — Published (2023)*

Inventory

5. *reskilling_article_atr21zcy325hdw2z5surdb8j_inventory-resource.txt*
6. *reskilling_article_drn4mg5r2seqxnlci35fm1i9_inventory-resource-1.txt*