AI4Debunk

D6.4 FIRST REPORT ON THE PROCESS OF CONTINUOUS GRAPH ADAPTATION SEPTEMBER 2025

Funded by
the European Union

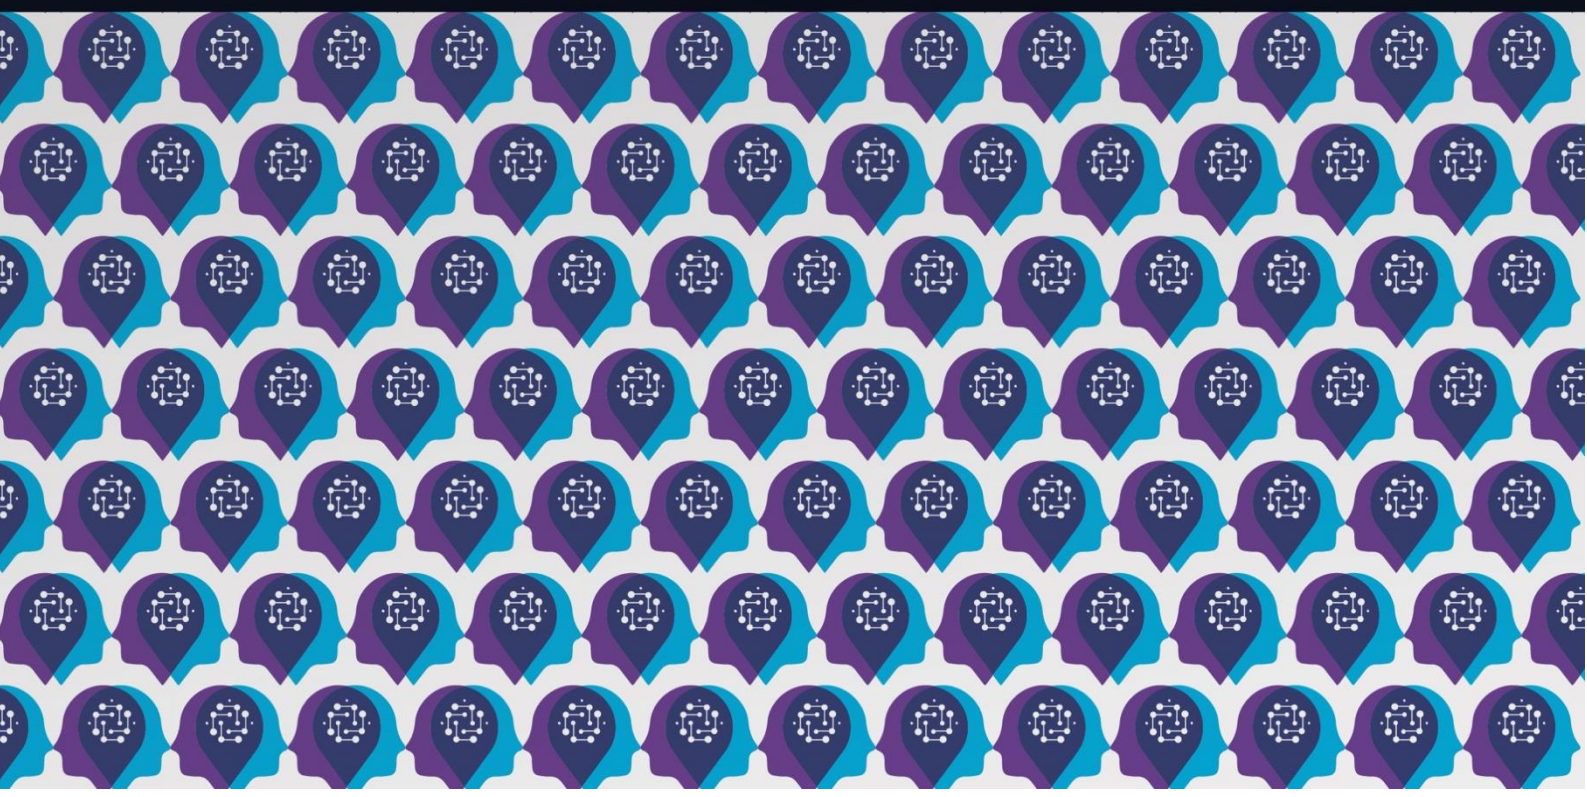

Grant Agreement No.: 101135757
Call: HORIZON-CL4-2023-HUMAN-01-CNECT
Topic: HORIZON-CL4-2023-HUMAN-01-05
Type of action: HORIZON Innovation Actions

D6.4 FIRST REPORT ON THE PROCESS OF CONTINUOUS GRAPH ADAPTATION

Project Acronym	AI4Debunk
Project Number	101135757
Project Full Title	Participative Assistive AI-powered Tools for Supporting Trustworthy Online Activity of Citizens and Debunking Disinformation
Work package	WP 6
Task	Task 4
Due date	30/09/2025
Submission date	04/09/2025
Deliverable lead	Partner UMONS
Version	1.0
Authors	Kevin El Haddad (Partner UMONS)
Contributors	Georgi Gotev (Partner EUalive), Kalina Angelova (Partner EUalive), Dzvenislava Shcherba (Partner IUI), Arianna D'Ulizia (Partner CNR), Alessia D'Andrea (Partner CNR), Roberto Caldelli (Partner CNIT), Stefano Berretti (Partner MICC-UNIFI), Jan Kragt (Partner IP), Marcel Keijzer (Partner IP), Franc Van Der Bent (HU), Jamal Nassir (Partner NUIG)
Reviewers	Žaneta Ozoliņa (University of Latvia)
Abstract	This deliverable presents a robust, scalable, and forward-looking framework for enriching the AI4Debunk knowledge graph with verified, structured, and semantically enriched data—laying the groundwork for a resilient, community-driven disinformation detection platform that can evolve and thrive beyond the initial funding period.

Funded by the European Union

Keywords	Knowledge graphs, framework, disinformation, data, platform

DOCUMENT DISSEMINATION LEVEL

Dissemination level	
X	PU - Public
	SEN - Sensitive

DOCUMENT REVISION HISTORY

Version	Date	Description of change	List of contributor(s)
0.1	01/08/2025	First draft	UMONS, EUalive, IUI
0.2	12/08/2025	Draft revised by the Task 6.4 participants	CNR, MICC-UNIFI, CNIT
0.3	18/08/2025	Final draft version revised by the internal reviewer	UL
0.4	31/08/2025	Final version	UL

STATEMENT ON MAINSTREAMING GENDER

The AI4Debunk consortium is committed to including gender and intersectionality as a transversal aspect in the project's activities. In line with EU guidelines and objectives, all partners – including the authors of this deliverable – recognise the importance of advancing gender analysis and sex-disaggregated data collection in the development of scientific research. Therefore, we commit to paying particular attention to including, monitoring, and periodically evaluating the participation of different genders in all activities developed within the project, including workshops, webinars and events but also surveys, interviews and research, in general. While applying a non-binary approach to data collection and promoting the participation of all genders in the activities, the partners will periodically reflect and inform about the limitations of their approach. Through an iterative learning process, they commit to plan and implement strategies that maximise the inclusion of more and more intersectional perspectives in their activities.

Funded by the European Union

DISCLAIMER

The AI4Debunk project has received funding from the European Union's Horizon Europe Programme under the Grant Agreement No. 101135757.

Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

COPYRIGHT NOTICE

How to cite this report: Kevin El Haddad, Georgi Gotev, Kalina Angelova, Dzvenislava Shcherba, Roberto Caldelli, Stefano Berretti, Jan Kragt, Marcel Keijzer, Franc Van Der Bent, Jamal Nassir, Alessia D'Andrea, Arianna D'Ulizia (2025). AI4Debunk D6.4: First Report on the process of continuous graph adaptation. https://ai4debunk.eu/wp-content/uploads/2025/09/AI4Debunk-Deliverable-6.4-final_300825.pdf

The AI4Debunk consortium is the following:

Participant number	Participant organisation name	Short name	Country
1	LATVIJAS UNIVERSITATE	UL	LV
2	FREE MEDIA BULGARIA SRL	EUalive	BE
3	PILOT4DEV	P4D	BE
4	INTERNEWS UKRAINE	IUA	UA
5	CONSIGLIO NAZIONALE DELLE RICERCHE	CNR-IRPPS	IT
6	UNIVERSITA DEGLI STUDI DI FIRENZE	MICC/UNIFI	IT
6.1	CONSORZIO NAZIONALE INTERUNIVERSITARIO PER LE TELECOMUNICAZIONI	CNIT	IT
7	BARCELONA SUPERCOMPUTING CENTER CENTRO NACIONAL DE SUPERCOMPUTACION	BSC	ES
8	DOTSOFT OLOKLIROMENES EFARMOGES DIADIKTIOY KAI VASEON DEDOMENON AE	DOTSOFT	EL
9	UNIVERSITE DE MONS	UMONS	BE
10	UNIVERSITY OF GALWAY	UOG	IE
11	STICHTING HOGESCHOOL UTRECHT	HU	NL
12	STICHTING INNOVATIVE POWER	IP	NL
13	F6S NETWORK IRELAND LIMITED	F6S	IE

Funded by the European Union

Funded by the European Union

ABBREVIATIONS

WP	Work Package
JSON	JavaScript Object Notation

Funded by the European Union

EXECUTIVE SUMMARY

This deliverable outlines the design and implementation of a system for adapting and enriching the AI4Debunk knowledge graphs, developed in Task 6.3, with incoming news data, ensuring that the system remains current and context-aware in the face of evolving information. The core objective is to maintain an up-to-date and trustworthy knowledge base by integrating news items that are either tagged manually by users or automatically using AI models developed in T8.1 and T8.2.

A structured validation process is established to maintain data reliability, where flagged news items must be reviewed and approved by a validation committee chaired by project partners EUalive (Free Media Bulgaria) and IUI. Once validated, news is further enhanced using information extraction algorithms from Tasks T8.1 and T8.2, as implemented in T6.2, before being inserted into the knowledge graphs.

To support consistency and ease of integration, a structured dataset—serialized in JSON format—has been deployed on the Hugging Face Hub, facilitating centralized access and supporting interoperability among project contributors. This dataset is supported by a suite of preprocessing pipelines to parse, validate, and normalize incoming news, enabling seamless integration into downstream systems and research workflows. Public accessibility of the dataset fosters transparency and strengthens the project's sustainability beyond its funded lifecycle.

Recognizing the scalability challenges of manual validation, the project proposes an automated validation system based on expert-defined rules, along with automated onboarding for new validation committee members. Periodic human review of a sample of automatically validated content is planned to ensure continued data quality.

Additionally, user-facing interfaces will be implemented to allow individuals to report and flag news items directly. These submissions will include metadata to assist in automated validation, increasing the efficiency and scalability of the overall system.

In summary, this deliverable presents a robust, scalable, and forward-looking framework for enriching the AI4Debunk knowledge graph with verified, structured, and semantically enriched data—laying the groundwork for a resilient, community-driven disinformation detection platform that can evolve and thrive beyond the initial funding period.

Funded by the European Union

1 INTRODUCTION

This deliverable presents the mechanisms and systems built for the adaptation of the knowledge graphs, developed in Task 6.3, by enriching them with new incoming news of different modalities, and updating the existing information in them. For this, the news could be tagged either by users manually or through the different tools provided by our platform (through the models developed in WP8.1).

Adapting the knowledge graphs with new data allows our knowledge base and therefore our entire system to stay up to date with the evolution of contextual information for better performances.

The tagged news must first be validated by a validation committee chaired by the partners EUalive (Free Media Bulgaria) and IUA before being inserted into the graph. Once validated the data are then enhanced by the information extraction algorithms developed in T8.1 and T8.2 and implemented in T6.2, before being inserted into the graph.

For the system to outlive the funding period of the AI4Debunk project on one side and to scale the processes, we attempt to go further by attempting to automate the validation process of the flagged news as well as the inclusion of a new validation committee member.

1.1 OBJECTIVES

The objective of this deliverable is to build an update mechanism for the adaptation of the knowledge graphs. The path of the data will be as follows:

1- Some news is reported/flagged by a user or by our models

2- The flag is checked by the validation committee

3- If validated, the news will be processed for information extraction, then inserted in our database, which is used to update the knowledge graphs.

We break down this task in the following objectives:

- Develop mechanisms and systems for adapting the knowledge graphs: Integrate incoming news data (multi-modal) into the existing knowledge graphs;
- Implement a structured validation workflow: Establish a validation committee (EUalive and IUA) for approving tagged news before insertion into the graph;
- Deploy a centralized and extensible data repository: Use Hugging Face Hub to store and access the AI4Debunk news corpus in a structured, extensible JSON format;
- Develop standardized preprocessing and ingestion pipelines: Ensure consistency and interoperability through an automated normalization workflow;

Funded by the European Union

- Promote sustainability and openness: Ensure the system's longevity and encourage community involvement by making datasets publicly accessible and by providing tools and resources to process them, making it easier to use the data in workflows;
- Build an interface for user-based news flagging: develop user-facing tabs/windows for disinformation flagging with metadata support for automatic validation.

Taking steps further than our initial objectives, we aim to reach the following objectives:

- Automate validation and committee member registration: design and deploy a rule-based system to scale validation and allow automatic onboarding of new committee members.
- Maintain quality assurance in automation: introduce periodic human sampling and validation of auto-validated news to maintain trust and accuracy.

1.2 EXPECTED OUTCOME

The above-mentioned objectives should achieve the following outcome:
- Continuously updated knowledge graph: The graph stays relevant with current events, improving system performance and contextual accuracy;
- User engagement through intuitive reporting interfaces: Empowered users contribute to system robustness by flagging content directly with meaningful metadata;
- Reliable data validation mechanism: Enhanced data integrity through oversight by a human validation committee;
- Richer and context-aware news data: Validated news is semantically enhanced before integration, supporting deeper graph reasoning;
- Open, accessible, and reusable dataset: Centralized, extensible repository fosters collaboration and long-term utility for the AI research community;
-Streamlined and scalable ingestion process: Unified pipelines ensure quality control and reduce friction in data onboarding across partners;
- Sustainable project infrastructure: Reduced dependency on human validators and funding cycles through automation and open platforms;
- Automated validation with expert oversight: Scalable system for news validation that reduces bottlenecks without compromising accuracy;
- Ongoing quality assurance: Periodic checks maintain high-quality standards for auto-validated content.

Funded by the European Union

2 DATA INGESTION PIPELINE

A structured dataset has been deployed on the Hugging Face Hub to store the AI4Debunk news corpus collected in Task 6.1. The dataset is serialized in JSON format, allowing for flexible schema design and extensibility at the level of individual data entries. To standardize data ingestion, a set of preprocessing pipelines and utility functions has been developed to parse, validate, and normalize incoming news items prior to insertion. This ensures consistency and interoperability across contributions from all project partners. Leveraging Hugging Face as the hosting platform facilitates programmatic access via the datasets and transformers libraries, enabling seamless integration into downstream workflows. Furthermore, the public availability of the dataset supports external stakeholder engagement and contributes to building an open research community—key components of the project's long-term sustainability strategy beyond the AI4Debunk timeline.

The structured dataset can be found at the following link:

War in Ukraine use-case: https://huggingface.co/datasets/AI4Debunk/war-in-ukraine-disinformation-detection_3rdrelease

Climate change use-case: https://huggingface.co/datasets/AI4Debunk/climate-change-disinformation-detection-3rdrelease

3 INTERFACE FOR DATA FLAGGING AND USER ENGAGEMENT

Input windows or tabs will be implemented in the context of WP10 and WP11, allowing for users to report news and flag them as disinformation or not with metadata allowing for automatic validation. The interface will be intuitive and easy to access to reduce barriers or adoption.

4 GOING BEYOND: AUTOMATING THE VALIDATION AND THE COMMITTEE MEMBER REGISTRATION PROCESSES FOR SCALING

Manual validation of flagged news can become a bottleneck as the volume of flagged items grows, particularly when the number of available human validators is limited. To address this challenge and ensure the system's sustainability beyond the AI4Debunk funding period, as well as to support scalability, we propose the automation of two key processes:

(i) the validation of flagged news items and

(ii) the registration of new validation committee members.

These will be achieved through a rule-based system built on expert-defined criteria. The work to define and implement these rules is in progress and will be completed in the context of WP7.

To maintain data integrity, periodic human reviews will be conducted on samples of automatically validated news, ensuring consistent quality and reliability across the dataset.

A specific policy will be adopted for the correct management of the personal data of the new committee members; in addition to this, some reward mechanisms could be foreseen in order to stimulate the affiliation of new members to the committee.

Funded by the European Union

5 CONCLUSION

This deliverable makes a crucial contribution to the AI4Debunk project by establishing the core framework for its knowledge graph's evolution and sustainability. It directly addresses the project's central mission of fighting disinformation by ensuring the AI4Debunk system remains a living, trusted resource.

Indeed, the development of adaptive mechanisms for updating the knowledge graph with new, validated, and enriched news data represents a critical advancement for the AI4Debunk project. By combining user-driven inputs, automated tagging tools, structured validation workflows, and robust information extraction methods, the system ensures both the accuracy and relevance of its knowledge base over time. The deployment of a centralized, extensible dataset on Hugging Face Hub enhances data accessibility, interoperability, and long-term sustainability.

Moreover, by addressing scalability through the automation of validation processes and committee member registration, the project moves toward a self-sustaining and future-proof architecture. Periodic human oversight safeguards data quality, while the development of user-friendly interfaces promotes community participation in disinformation detection. Together, these components lay a strong foundation for a resilient, efficient, and collaborative ecosystem capable of enduring beyond the funding horizon of AI4Debunk and adapting to the evolving challenges of misinformation in the digital age.

Funded by the European Union

Review Sheet of Deliverable/ Milestone Report
D6.2 Deliverable Title

Editor(s):	Kevin El Haddad
Responsible Partner:	UMONS (UNIVERSITE DE MONS)
Status-Version:	Draft 0.1
Date:	01/08/2025
Distribution level (CO, PU):	Public
Reviewer (Name/Organization)	Žaneta Ozoliņa (University of Latvia)
Review date	14/08/2025 and 30/08/2025

Disclaimer: This assessment reflects only the author's views and the European Commission is not responsible for any use that may be made of the information contained therein"

Funded by the European Union

Mark with X the corresponding column:

Y= yes	N= no	N = not applicable

ELEMENT TO REVIEW	Y	N	NA	COMMENTS
FORMAT: Does the document … ?				
…include editors, deliverable name, version number, dissemination level, date, and status?	Y			
…contain a license (in case of public deliverables)?			NA	
…include the names of contributors and reviewers?	Y			
….has a version table consistent with the document's revision?	Y			
… contain an updated table of contents?	Y			
… contain a list of figures consistent with the document's content?			NA	
… contain a list of tables consistent with the document's content?	Y			
… contain a list of terms and abbreviations?	Y			Only one term is mentioned but more would be needed, i.e. JSON.
… contain an Executive Summary?	Y			
… contain a Conclusions section?	Y			
… contain a List of References (Bibliography) in the adequate format, if relevant?			NA	
… use the fonts and sections defined in the official template?	Y			
… use correct spelling and grammar?	Y			
… conform to length guidelines (50 pages maximum (plus Executive Summary and annexes)	Y			
… conform to guidelines regarding Annexes (inclusion of complementary information)			NA	
… present consistency along the whole document in terms of English	Y			

Funded by the European Union

ELEMENT TO REVIEW	Y	N	N A	COMMENTS
quality/style? (to avoid accidental usage of copy&paste text)				
About the content…				
ELEMENT TO REVIEW	**Y**	**N**	**N A**	**COMMENTS**
Is the overall style of the deliverable correctly organized and presented in a logical order?	Y			
Is the Executive Summary self-contained, following the guidelines and does it include the main conclusions of the document?	Y			
Is the body of the deliverable (technique, methodology results, discussion) well enough explained?	Y			
Are the contents of the document treated with the required depth?	Y			
Does the document need additional sections to be considered complete?		N		
Are there any sections in the document that should be removed?		N		
Are all references in the document included in the references list?			N A	
Have you noticed any text in the document not well referenced? (copy and paste of text/picture without including the reference in the reference list)		N		
SOCIAL and TECHNICAL RESEARCH WPs (WP4, 5, 12, 13, 14)				
ELEMENT TO REVIEW	**Y**	**N**	**N A**	**COMMENTS**
Is the deliverable sufficiently innovative?			N A	
Does the document present technical soundness and its methods are correctly explained?			N A	
What do you think is the strongest aspect of the deliverable?			N A	

Funded by the European Union

ELEMENT TO REVIEW	Y	N	NA	COMMENTS
What do you think is the weakest aspect of the deliverable?			NA	
Please perform a brief evaluation and/or validation of the results, if applicable.			NA	
AI AND TECNOLOGICAL WPS (WP6 – WP11)				
ELEMENT TO REVIEW	**Y**	**N**	**NA**	**COMMENTS**
Does the document present technical soundness and the methods are correctly explained?	Y			
What do you think is the strongest aspect of the deliverable?				Clear structure and soundness of the approach.
What do you think is the weakest aspect of the deliverable?				
Please perform a brief evaluation and/or validation of the results, if applicable.				
DISSEMINATION AND EXPLOITATION WPs (WP15 – WP17)				
ELEMENT TO REVIEW	**Y**	**N**	**NA**	**COMMENTS**
Does the document present a consistent outreach and exploitation strategy?			NA	
Are the methods and means correctly explained?			NA	
What do you think is the strongest aspect of the deliverable?			NA	
What do you think is the weakest aspect of the deliverable?			NA	
Please perform a brief evaluation and/or validation of the results, if applicable.			NA	
DISSEMINATION AND EXPLOITATION WPs (WP18)				
ELEMENT TO REVIEW	**Y**	**N**	**NA**	**COMMENTS**
Does the document present the main ethical aspects regarding the use of methods and human involvement?			NA	
What do you think is the strongest aspect of the deliverable?			NA	

Funded by the European Union

ELEMENT TO REVIEW	Y	N	N A	COMMENTS
What do you think is the weakest aspect of the deliverable?			N A	
Please perform a brief evaluation and/or validation of the results, if applicable.			N A	

CONCLUSION

Mark with X the corresponding line.

X	Document accepted, no changes required.
	Document accepted, changes required.
	Document not accepted, it must be reviewed after changes are implemented.

Please rank this document globally on a scale of 1-5 (1 = poor, 5= excellent) – using a half point scale. Mark with X the corresponding grade.

Document grade	1	1.5	2	2.5	3	3.5	4	4.5	5
								X	

Funded by the European Union