Robust Learning and Reasoning
for Complex Event Forecasting

Project Acronym:	EVENFLOW
Grant Agreement number:	101070430 (HORIZON-CL4-2021-HUMAN-01-01 – Research and Innovation Action)
Project Full Title:	Robust Learning and Reasoning for Complex Event Forecasting

DELIVERABLE

D2.3 – Final Report on the Dissemination and Exploitation Plan

Dissemination level:	PU - Public, fully open
Type of deliverable:	R - Document, report
Contractual date of delivery:	31 December 2025
Deliverable leader:	NCSR
Status - version, date:	Final, v1.0, 2025-12-17
Keywords:	Dissemination, communication, exploitation, market research, open source, IPR management, communication activities, outreach, neurosymbolic AI

Funded by the European Union

This document is part of a project that is funded by the European Union under the Horizon Europe agreement No 101070430. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the Commission. The document is the property of the EVENFLOW project and should not be distributed or reproduced without prior approval. Find us at www.evenflow-project.eu.

Executive Summary

The EVENFLOW project has successfully completed its dissemination, communication, and exploitation activities over its 39-month duration. This deliverable (D2.3) provides the final overview of all actions implemented under Work Package 2, summarising how the consortium promoted the project's visibility, strengthened stakeholder engagement, and maximised the exploitation potential of its results.

The dissemination and communication strategy, initially defined in Deliverable D2.2, was fully implemented and continuously monitored throughout the project's lifetime. It aimed to reach diverse audiences—scientific, industrial and general public—through a balanced combination of digital communication, targeted campaigns, and participation in international events. The project's website served as the central communication hub, complemented by social media channels, newsletters, and a series of specific campaigns such as *Meet the Partner*, *Use Cases in a Nutshell*, *EVENFLOW AI Tools*, and *Subscribe to Our Newsletter and Meet the TrustWorthy AI Cluster*. These coordinated actions enhanced the project's digital presence, community building, and outreach impact.

EVENFLOW partners actively contributed to international and European conferences, workshops and webinars, promoting scientific results and engaging with stakeholders. In total, more than 20 publications were produced, openly accessible through a dedicated Zenodo community and the project website, ensuring compliance with EU open access. A key achievement of EVENFLOW has been the coordination of the TrustWorthy AI Cluster, which brought together nine Horizon Europe projects focused on reliable AI. The cluster fostered knowledge exchange and joint outreach through regular coordination meetings, shared visual identity, collective participation in major European events, and the joint webinar *Trustworthy AI: Landscaping Verifiable Robustness and Transparency*. These efforts amplified the collective visibility and impact of EU-funded AI research.

The project's exploitation work focused on making EVENFLOW's technical results sustained beyond the end of the project. Building on the strategy defined in D2.2, partners identified and tracked all exploitable results via an innovation management log, leading to a final list that includes the integrated EVENFLOW platform, the three core tools/ toolsets (Reasoning-Based Forecast Interpreter / DeepFA, Scalability Toolkit, Verification Toolkit) and several use-case software pipelines. Most of these components are released as open-source projects (AGPL, Apache 2.0, BSD/GPL licences) following a licence compatibility check. For each toolset and each pilot domain (healthcare, Industry 4.0, infrastructure lifecycle management), detailed exploitation models were developed, describing target users, value propositions, IPR and possible exploitation channels. At organisational level, all partners prepared individual exploitation plans. Given the consortium profile, exploitation is mainly research-driven and based on (i) reusing the toolsets as background IP in new Horizon Europe and related research projects, (ii) maintaining and extending the open-source repositories and building user communities, and (iii) offering consulting and custom development services around neurosymbolic forecasting, scalable training and verification.

Deliverable leader:	Fay Pitsilidou (NCSR)
Contributors:	Konstantinos Chisiridis (INTRA)
Reviewers:	Athanasios Poulakidas (INTRA), Hekal Abdelrahman (ICL)
Approved by:	Athanasios Poulakidas, Eleni-Vasiliki Provopoulou (INTRA)

Document History

Version	Date	Contributor(s)	Description
v0.1	1/10/2025	Fay Pitsilidou	Creation of the first version of the document, table of contents
v0.2	20/11/2025	Fay Pitsilidou	Finalised document – Communication and Dissemination
v0.3	24/11/2025	Konstantinos Chisiridis	Exploitation – Market Analysis of the Neurosymbolic AI market
v0.4	28/11/2025	Konstantinos Chisiridis	Exploitation – partners' individual exploitation plans
v0.5	05/12/2025	Konstantinos Chisiridis	Exploitation – Exploitation models/ business models per Toolset
v0.6	15/12/2025	Fay Pitsilidou, Konstantinos Chisiridis	Final version ready for quality review, incorporating comments from the internal review
v1.0	17/12/2025	Athanasios Poulakidas, Eleni-Vasiliki Provopoulou	Version to be submitted after final QA

Table of Contents

Table of Figures

List of Tables

Definitions, Acronyms and Abbreviations

Acronym/ Abbreviation	Title
ADRA-e	AI, Data and Robotics Association – ecosystem
AIoD	AI-on-Demand Platform
CSA	Coordination and Support Action
PPP	Public–Private Partnership
AI	Artificial Intelligence
GDPR	General Data Protection Regulation
DMP	Data Management Plan
SM	Social media
BG IP	Background Intellectual Property
ER	Exploitable Result
FG IP	Foreground Intellectual Property
IML	Innovation Management Log
IPR	Intellectual Property Rights

Term	Definition
Comms Team	Communications Team
Exploitation Model	Plan of how a project result will be used, monetised, transferred, or further developed after project completion. Includes target users, value proposition, ownership, and exploitation channels.
Individual Exploitation Plan	Partner-specific plan outlining how each organisation intends to use project results, considering commercial, research, consulting, or open-source pathways.
Innovation Management Log (IML)	Internal innovation management methodology used in EVENFLOW to collect, monitor and update all exploitable results, track ownership, TRL, licensing, and exploitation intent.
Open-Source Software	Software whose source code is publicly available under an approved open-source licence, allowing users to access, modify, reuse and redistribute it. It is a core exploitation and sustainability pathway in EVENFLOW, enabling partners to maintain and further develop the toolsets beyond the project. Typical licences applicable to research-driven software include MIT License, Apache License 2.0, BSD 3-Clause, and GNU GPL-v3, each offering different conditions regarding attribution, redistribution and derivative works.

1 Introduction

1.1 Project Information

EVENFLOW develops hybrid learning techniques for complex event forecasting, which combine deep learning with logic-based learning and reasoning into neuro-symbolic forecasting models. This approach combines neural representation learning techniques that construct event-driven features from streams of perception-level data with powerful symbolic learning and reasoning tools, which utilize such features to synthesize high-level, interpretable patterns for forecasting critical events.

To deal with the brittleness of neural predictors and the high volume/velocity of temporal data flows, the EVENFLOW techniques rely on novel, formal verification techniques for machine learning, in addition to a suite of scalability algorithms for training based on data synopsis, federated training and incremental model construction. The learnt forecasters will be interpretable and scalable, allowing for explainable and robust insights, delivered in a timely fashion and enabling proactive decision making.

EVENFLOW is evaluated on three use cases related to (1) oncological forecasting in healthcare, (2) safe and efficient behaviour of autonomous transportation robots in smart factories and (3) reliable life cycle assessment of critical infrastructure.

Table 1: The EVENFLOW consortium.

Number[1]	Name	Country	Short name
1 (CO)	NETCOMPANY-INTRASOFT	Belgium	**INTRA**
1.1 (AE)	NETCOMPANY-INTRASOFT SA	Luxemburg	**INTRA-LU**
2	NATIONAL CENTER FOR SCIENTIFIC RESEARCH "DEMOKRITOS"	Greece	**NCSR**
3	ATHINA-EREVNITIKO KENTRO KAINOTOMIAS STIS TECHNOLOGIES TIS PLIROFORIAS, TON EPIKOINONION KAI TIS GNOSIS	Greece	**ARC**
4	BARCELONA SUPERCOMPUTING CENTER-CENTRO NACIONAL DE SUPERCOMPUTACION	Spain	**BSC**
5	DEUTSCHES FORSCHUNGSZENTRUM FUR KUNSTLICHE INTELLIGENZ GMBH	Germany	**DFKI**
6	EKSO SRL	Italy	**EKSO**
7 (AP)	IMPERIAL COLLEGE OF SCIENCE TECHNOLOGY AND MEDICINE	United Kingdom	**ICL**

[1] CO: Coordinator. AE: Affiliated Entity. AP: Associated Partner.

1.2 Document Scope

The dissemination and communication activities that the EVENFLOW partnership has undertaken in the whole of its duration (M01-M39), aimed to provide increased visibility of the project in the audiences defined in Deliverable 2.2 by creating appropriate visual and editorial material, as well as by ensuring a regular outward flow of information about the project's progress and results achieved.

To maximise the impact and improve the exploitation potential of EVENFLOW, a communication and dissemination plan has been developed and followed up on (*Deliverable 2.2 Plan for Dissemination and Exploitation including Communication Activities*), with the objectives focusing on dissemination and communication aspects, as well as community building, engagement and adoption of EVENFLOW results.

This deliverable provides a comprehensive overview of the dissemination, communication, and awareness activities carried out under WP2 throughout the entire duration of the project. The purpose and scope of D2.3 are reflected in Tasks 2.1 Dissemination Activities, T2.2 Exploitation & Business Planning and T2.3 Liaison Activities, the objectives of which are specified in the DoA as shown below:

T2.1 Dissemination Activities (D2.1 - D2.3) [M1 - M39]
This task sets the dissemination objectives and will design and implement EVENFLOW's dissemination plan. This includes the project's website and social media channels set-up, production of dissemination material, such as newsletters, brochures and flyers and the organisation of dissemination events. Avenues for communicating the project's results and lessons learnt will primarily be top-tier conference venues and journals, but also scientific and industrial workshops, white-papers and press releases. Throughout the project, dissemination activities will be monitored, evaluated and adapted accordingly.

T2.2 Exploitation & Business Planning (D2.2, D2.3) [M1 - M39]
Task 2.2 includes all actions aimed at exploiting the project outcomes and connecting with related stakeholders across the EU. EVENFLOW's exploitation plan will be dynamic, evolving during the project, eventually leading to a sustainable business plan. A preliminary market analysis and requirements assessment will be conducted in early stages of the project, together with a technical feasibility study that will ultimately lead to the development of viable market strategies.

T2.3 Liaison Activities (D2.2, D2.3) [M1 - M39]
This task deal with communication and synergy-building activities between the related PPP on AI, Data and Robotics, the CSA funded by HORIZON-CL4-2021-HUMAN-01-02 call, as well as the AI4EU platform. Among the targets of this task are the organisation of joint workshops.

1.3 Document Structure

This deliverable is a public report documenting the various dissemination, exploitation, awareness and outreach activities and results for the whole duration of the project. Deliverable 2.3 follows a structured approach reflecting the logic of the project's dissemination and exploitation framework established in D2.2. Its methodology is based on

continuous monitoring, stakeholder mapping, and evaluation of performance against predefined Key Performance Indicators (KPIs). The structure of the report aligns with the major components of Work Package 2:

Chapter Introduction **– Introduction:** Defines the scope of dissemination, communication, and exploitation within the EVENFLOW context and summarises the work carried out across Tasks 2.1, 2.2 and 2.3.

Chapter Dissemination and Communication **– Dissemination and Communication:** Describes the overarching objectives, target audiences, and performance indicators used to measure the success of dissemination activities. It also presents detailed reporting on internal and external communication, including website analytics, social media engagement, campaign results, event participation, and publications.

Chapter Exploitation and IPR Management Strategy **– Exploitation and IPR Management Strategy:** The exploitation and sustainability dimension of the project is addressed in **Chapter 4**. It presents the final consolidated list of exploitable results (integrated EVENFLOW platform, core toolsets and use-case modules) together with ownership, TRL and licensing information, followed by a market analysis of the neurosymbolic AI domain, including an overview of similar solutions and a SWOT analysis of the EVENFLOW technologies. Subsequent sections present the exploitation models for each core toolset (Reasoning-Based Forecast Interpreter, Scalability Toolkit, Verification Toolkit) and for each use case (precision medicine, Industry 4.0 intralogistics, infrastructure lifecycle assessment), detailing target users, value propositions, IPR choices and exploitation channels. Finally, the chapter consolidates the individual exploitation plans per partner and sets out the sustainability approach and post-project roadmap for maintaining and evolving the open-source assets created in EVENFLOW and associated business and research exploitation paths.

Chapter Conclusion **– Conclusions:** Summarises overall achievements, lessons learned, and the impact of the dissemination and exploitation activities.

2 Dissemination and Communication

2.1 Overall Strategy and Objectives

The strategy outlined in *Deliverable 2.2 — Plan for Dissemination & Exploitation including Communication Activities* was implemented throughout the entire duration of the project. It placed strong emphasis on promoting both the scientific results and the overall visibility of the EVENFLOW project. Through targeted actions and communication channels, the strategy aimed to ensure that the project's findings reached the relevant audiences, enhanced awareness of its impact, and supported the long-term exploitation of its outcomes.

2.2 Defining Target Audiences and Ecosystem

In pursuit of EVENFLOW's objectives, the consortium identified the main stakeholder groups requiring regular updates on the project's developments and results. These audiences were categorised into two main groups: internal stakeholders (primarily project partners) and external ones (all other communication recipients). When engaging with scientific and policy communities, the project focused on delivering detailed and technical content, whereas outreach to the general public emphasised more accessible and widely understandable messages about the project and its achievements. Because these audiences differed in both their interests and information needs, Deliverable 2.2 sets out the tailored communication strategy designed to address each group effectively.

2.3 Key Performance Indicators (KPIs) of Dissemination Activities

To enable EVENFLOW partners and EC officials to assess the impact of the dissemination and communication strategy, a set of measurable success indicators was established, providing a clear basis for evaluating the achievement of the project's objectives. For online dissemination, data were collected through a Matomo Analytics account linked to the project website, while for social media platforms, the built-in analytics tools (e.g., LinkedIn Analytics) were used to monitor performance and engagement.

The table below presents both the initial targets and the final outcomes.

Table 2: EVENFLOW Dissemination KPIs.

Online communication activities	Success Indicators	Final outcomes (until November 2025)
Project Website: The dedicated EVENFLOW website will be the main dissemination reference with information about the project, its objectives, partners, results etc.	- AVG yearly visits: > 2,000 - Material downloads (yearly): 20	- 4,568 total website visits - 175 downloads
Social media: Accounts will be created on the major social media networks to disseminate information about the project, its updates, form communities of interest, and interact with stakeholders. Social media content strategy will be part of the EVENFLOW dissemination plan.	> 400 Twitter followers in total > 200 members on LinkedIn in total	> 70 Twitter followers > 216 LinkedIn followers

Project Video: will communicate the EVENFLOW concept and approach in a simplified and direct style appropriate to all target groups, including non-technical stakeholders.	4 videos in total	9 videos in total
Project Newsletter: will be emailed to interested recipients and published online on the project website and social media. The newsletter will include project updates, inform, and engage with end-users in the project activities.	6 Newsletters in total No. of recipients: > 200 in total	6 newsletters > 3,500 recipients
Publications		
Press releases & articles: Relevant publications will be prepared for all stakeholder communities including press releases and articles that describe project outcomes.	9 in total	2 PR and 55 news items
White papers: An authoritative guide that discusses issues on the certain subjects of EVENFLOW projects, along with a proposed solution for handling them.	4 in total	3 in total
Peer-reviewed scientific journals and conference publications: Partners from the research, academia and industry will engage in scientific publications in journals and conferences.	≥ 12 in total	21 in total
Participation in events		
Key related events will be attended by EVENFLOW partners. These events include scientific conferences, workshops, exhibitions, congresses, and commercial events.	> 20 in total	42 in total
EVENFLOW events: The consortium will organise five Industrial/Scientific/Training Workshops & Webinars dedicated to project activities, involving external stakeholders from the end user groups, industry and policy or decision makers.	> 5 in total	5 in total

2.4 Summary of Dissemination & Communication Activities

To achieve EVENFLOW goals and objectives, partners have identified key stakeholders who needed to be kept up to date with the progress and outcomes of the project. The dissemination and communication activities scheduled and realised by WP2 had dedicated target audiences which were widely separated in two groups: internal and external. In the sections below the distinction between the two shows the diversity of the ways of communication, the channels utilised as well as the conveyed messages.

2.5 Internal Communication

The EVENFLOW Partnership are classed as internal audience whilst all other stakeholders are classed as external audiences (which is the biggest part of this WP).

Communication with Partners

Partners as an internal audience, required easy to use, daily communication within the partnership, mainly via digital means. To ensure smooth interaction and safe exchange of information within the consortium, partners have agreed, and the coordinator has established, internal communication channels, as early as the kick-off meeting. A detailed overview of the tools used for internal communication within the partnership can be found in Deliverable D2.2 *Plan for Dissemination and Exploitation including Communication Activities*, on page 16.

2.6 External Communication

Communication with external audiences represented the focus of T2.1 and T2.3, as they encompassed the activities carried out by all project partners to engage the identified target groups and achieve the defined objectives. Below follows an overview of all communication activities undertaken, highlighting the range of tools and channels used in line with the strategy outlined in the communication plan (D2.2).

2.7 Dissemination & Communication Channels and Activities

This section presents an overview of the dissemination and communication activities carried out throughout the project's duration to ensure broad visibility, stakeholder engagement, and effective knowledge transfer. It outlines the activities, tools, and channels employed to reach diverse audiences, including scientific communities, industry stakeholders, policymakers, and the general public. The activities described below were designed to maximise the project's impact, enhance awareness of its objectives and results, and foster collaboration and uptake of outcomes beyond the project's lifetime. The current deliverable reflects the information available until month 39; any additional Dissemination and Communication actions and/or updated figures will be presented during the final review meeting.

2.7.1 Project Website

Among the project's various communication tools, the official website played a pivotal role as the main digital hub for visibility and outreach. It provided a flexible environment for presenting project objectives, activities, and results, while facilitating interaction with stakeholders across the project's ecosystem. The website acted as both an information repository and a dynamic interface for engagement. A comprehensive overview of its design, structure, and functionalities is available in *Deliverable 2.1 Project Presentation and Website*, page 14–20.

The EVENFLOW project website has been made publicly available in November 2022, under the URL https://evenflow-project.eu/. The website was frequently updated depending every time on the stage of the project and will remain to be so until after the end of the project by WP2, Comms Lead, NCSR, who is responsible for the maintenance and for sourcing content from all EVENFLOW partners. The most populated areas of the project were the sections *News* and *Publications*.

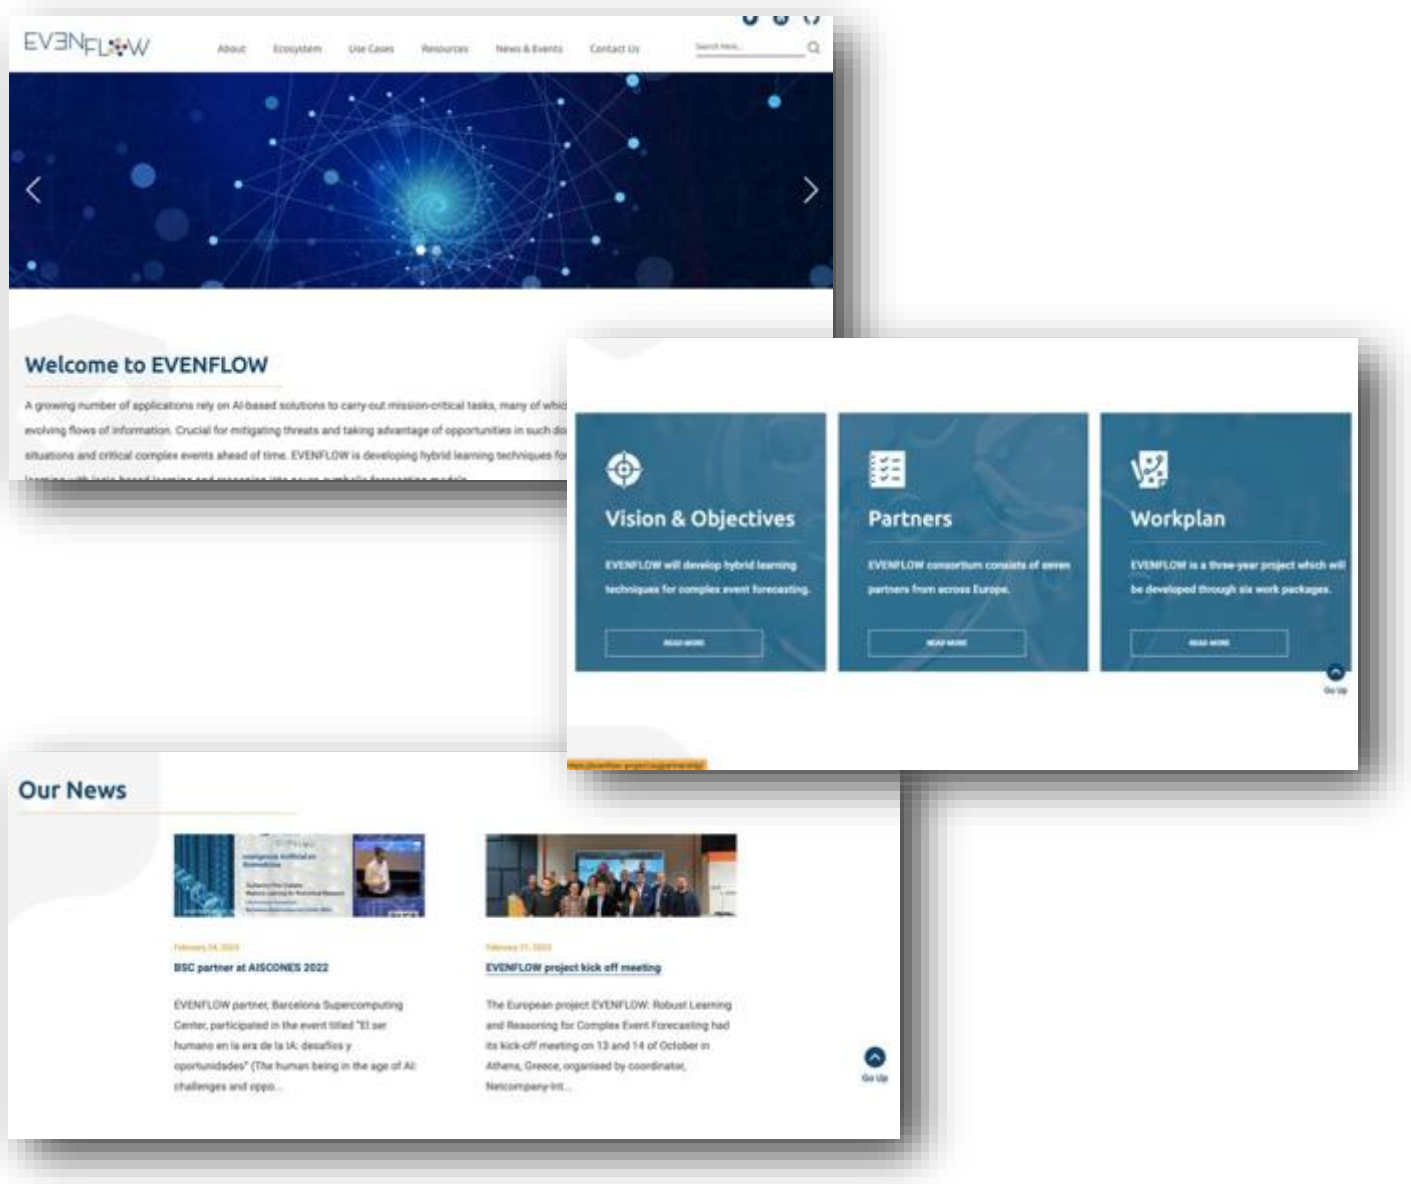

Figure 1: EVENFLOW Website.

To monitor website usage and visitor behaviour, WP2 has employed website analytics tools (Matomo Analytics) linked to the website since its launch. Such tools helped us monitor user behaviour and extract statistics so much so for reporting purposes as well as for extracting conclusions and then amending the communication and dissemination strategy as required. For the reference period between November 2022 (first operation of the website) until 19 November 2025 (date of this deliverable) the website has attracted more than 4,568 visits.

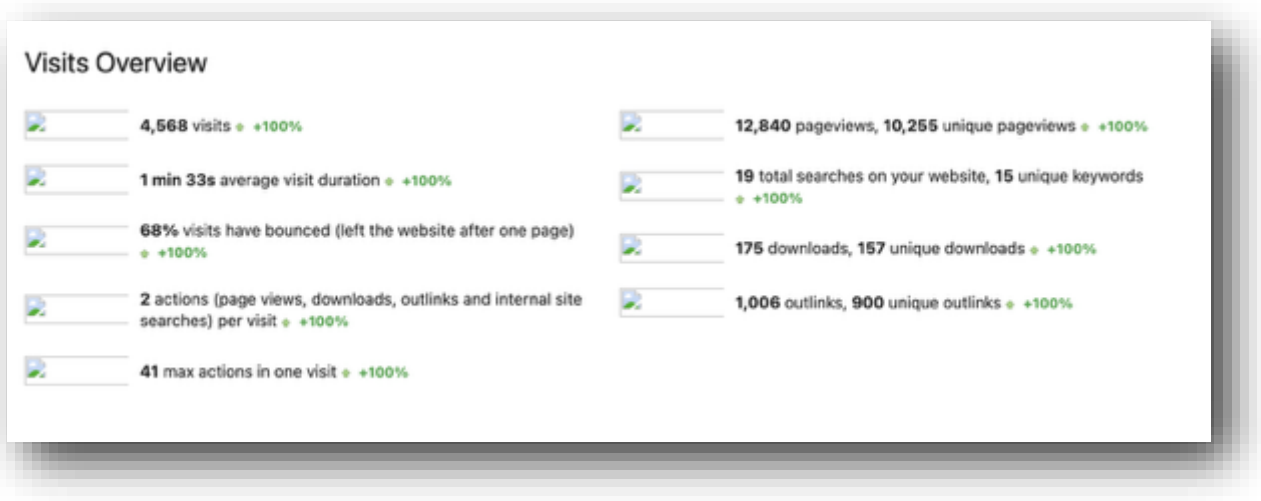

Figure 2: Website analytics February 2022 - November 2025.

The average engagement time within the website is approximately four minutes which is considered very good, when compared to the field average. A considerable number of users landed on the website via a direct link, provided through our campaigns. Organic traffic via search engines (Google etc.) and social media campaigns seem to have worked well in driving traffic to the EVENFLOW website.

As indicated in the analytics, the most visited page on the EVENFLOW website (except the Home page), is the Publications page, followed by them News, Partnership and TrustWorthy AI Cluster pages.

Channel Types

CHANNEL TYPE	VISITS	ACTIONS	ACTIONS PER VISIT	AVG. TIME ON WEBSITE	BOUNCE RATE
Direct Entry	4,107	9,374	2.3	2 min 12s	66%
Search Engines	1,383	3,268	2.4	1 min 45s	54%
Websites	318	813	2.6	2 min 13s	52%
Social Networks	234	554	2.4	2 min 3s	44%

Figure 3: Website analytics showcasing direct entries.

Page titles

PAGE TITLE	PAGEVIEWS	UNIQUE PAGEVIEWS	BOUNCE RATE	AVG. TIME ON PAGE	EXIT RATE	AVG. PAGE LOAD TIME
Evenflow	4,418	3,448	56%	00:00:45	64%	3.8s
Publications | Evenflow	815	516	53%	00:01:19	69%	8.74s
News | Evenflow	874	481	61%	00:00:41	38%	2.14s
partnership | Evenflow	486	413	59%	00:01:13	54%	1.91s
trustworthy ai cluster | Evenflow	463	380	68%	00:01:20	73%	3.35s
Vision & Objectives | Evenflow	366	319	76%	00:00:55	48%	1.8s
Meet the Partner | Evenflow	416	314	73%	00:01:07	81%	7.14s

Figure 4: Most visited pages on EVENFLOW website.

2.7.2 Social Media Channels

When the time came to choose from the variety of social media channels available in the beginning of the project, the Communications Team considered two main factors:

- What does the domain and its stakeholders use? We researched what social media the ecosystem, sibling projects, key stakeholders, policy makers, governmental bodies and the EC utilise.
- What do our partners use? Following thorough investigation of partners' social media, it was decided to create accounts on social media channels that our partners would be able to follow and share content from. Thus, the choice was made to create accounts on LinkedIn and X (former Twitter).

LinkedIn

Following the above-mentioned decision-making process, the creation of a LinkedIn page has been established at https://www.linkedin.com/company/evenflow-project/, which currently counts 216 followers—surpassing the KPI target.

A LinkedIn account was created to enhance visibility among professionals, researchers, and stakeholders in the field, fostering networking opportunities for collaboration and partnership. Through LinkedIn, EVENFLOW disseminated research outputs, publications, and project updates to a broad audience, thereby increasing the impact and visibility of the project's activities. The platform also facilitated engagement with key stakeholders such as other European projects, universities, and researchers, enabling discussions, feedback, and interaction. Maintaining a professional presence on LinkedIn helped strengthen the project's brand and reputation within the research community and beyond, showcasing its expertise, achievements, and contributions to the field. Based on the level of engagement and interaction observed, it became evident that LinkedIn was the project's main and most effective communication channel, where its target audience was most active.

Figure 5: EVENFLOW LinkedIn profile

X (former Twitter) and BlueSky

The consortium also agreed to create an X (former Twitter) account in November 2022 (@EvenflowProject – 70 followers). The platform was used to share consortium and project updates, promote events and outputs, and build a wider community around the project's core area of interest. Twitter (now X) also served as a channel to direct audiences to the project website through targeted links and timely posts.

Even though maintaining an active presence for two years, recent developments on the platform—such as changes in ownership, reduced organic reach, and the decision of many media outlets, universities, and stakeholders to limit or cease their activity—have significantly

impacted engagement. Growth in followers has been slow, and the discontinuation of analytics in the free version of X further limited our ability to monitor performance.

In response to these changes and following internal discussions, the consortium decided in March 2025 to establish a presence on Bluesky (not foreseen in the project's Description of Action) while continuing to maintain the X account. The same handle (@evenflowproject.bsky.social – 13 followers) was used to ensure continuity and facilitate discovery by existing followers. Initial efforts focused on connecting with users who had migrated to the new platform.

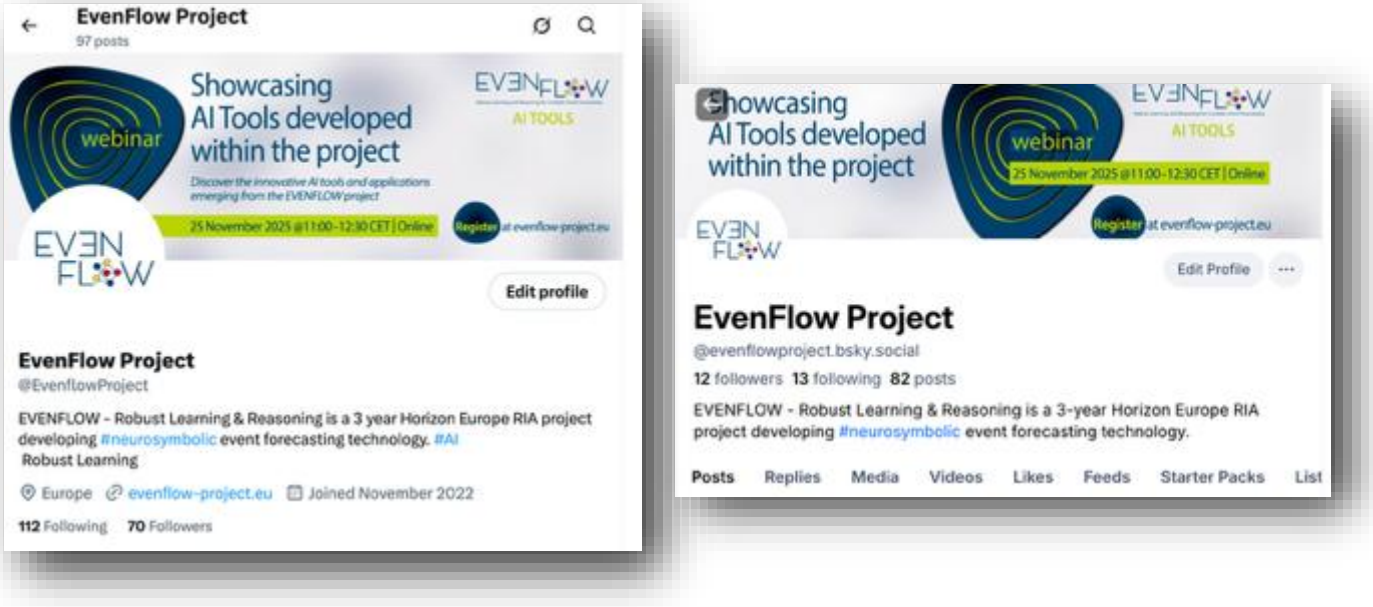

Figure 6: X and BlueSky profiles.

Despite consistent and high-quality communication efforts, the project did not fully meet its KPIs on X and Bluesky. The main reasons include the structural and algorithmic changes to X that reduced visibility for institutional accounts, and the early-stage nature of Bluesky has a limited active user base. Furthermore, the project's specialised focus and limited resources for targeted promotion naturally constrained audience growth. Nevertheless, the project continues to maintain strong visibility through regular updates on its LinkedIn account, its website, partner channels, newsletters, and other dissemination tools.

YouTube

After careful consideration, we have decided against creating a new YouTube channel for our three-year European project. Establishing a substantial subscriber base from scratch is both a challenging and time-consuming process that may not be the most efficient use of our resources. Instead, we leveraged the existing YouTube channel of NCSR (WP2 Leader), which already boasts a strong follower base (483 subscribers). To ensure our content is easily accessible and well-organised, we created a dedicated project playlist on this established channel. This approach allowed us to maximise our outreach and engagement efforts more

effectively. All videos produced as part of the project's activities are also available on the EVENFLOW website, under the dedicated *Videos* subcategory. A total of nine videos has been uploaded to date, accumulating more than 289 views overall.

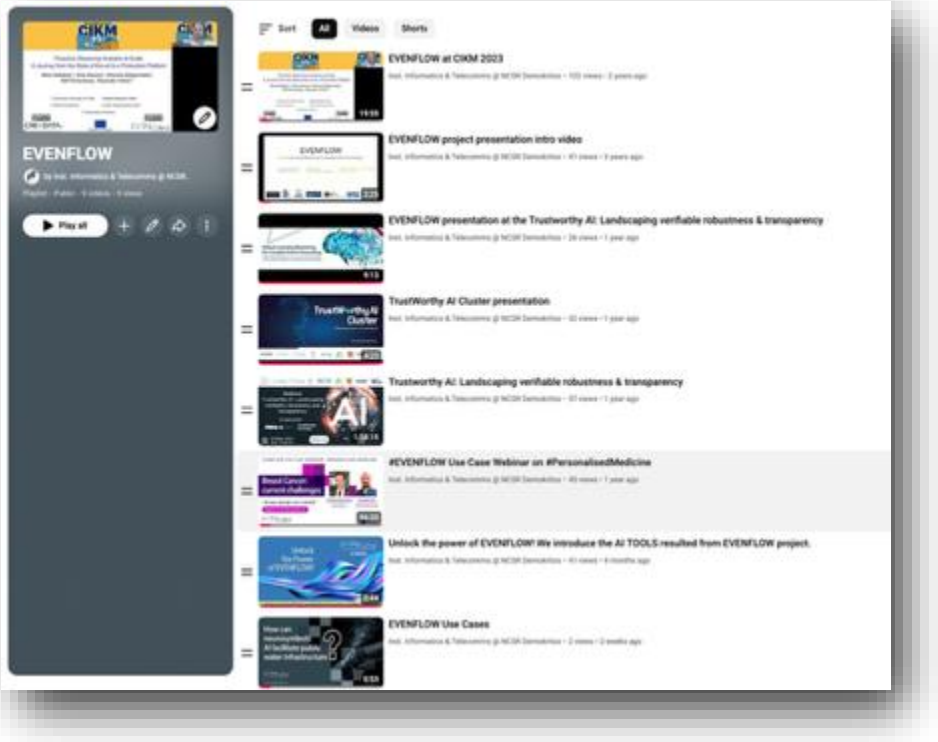

Figure7:EVENFLOW YouTube playlist.

2.7.3 Communication & Dissemination Campaigns

To maximise the visibility and impact of the project's results, a series of targeted communication and dissemination campaigns were implemented throughout the reporting period. These campaigns aimed to raise awareness among key stakeholders, engage relevant communities, and promote the project's achievements to a wider audience. Activities included the production and publication of news items on EVENFLOW website, the design of appealing visual materials, and the regular dissemination of content through social media channels and newsletters. Together, these efforts contributed to strengthening the project's identity, enhancing public understanding of its objectives, and fostering knowledge exchange across the European ecosystem.

Meet the Partner

In the frame of creating visual content in the first six months and to raise awareness about the project, a communication initiative named *Meet the Partner* was up and running for seven weeks (one partner per week). This online activity was an initiative with the purpose to promote partners, their role and work within the EVENFLOW project. Respective banners were designed, introducing to the wider public each organisation and its lead person working on the project. This activity was communicated once a week (every Tuesday starting from

March 2023) via the EVENFLOW project website and social media. A dedicated webpage was created on the website under the *Partnership* section, where all the partners' information were stated and where the audience was able to refer to. The campaign appeared to have been highly successful, as evidenced by our website analytics, which show increased traffic and indicate that the *Meet the Partner* webpage was the second most visited page on the EVENFLOW website for that period of time.

Figure 8: Meet the Partner visuals and social media post.

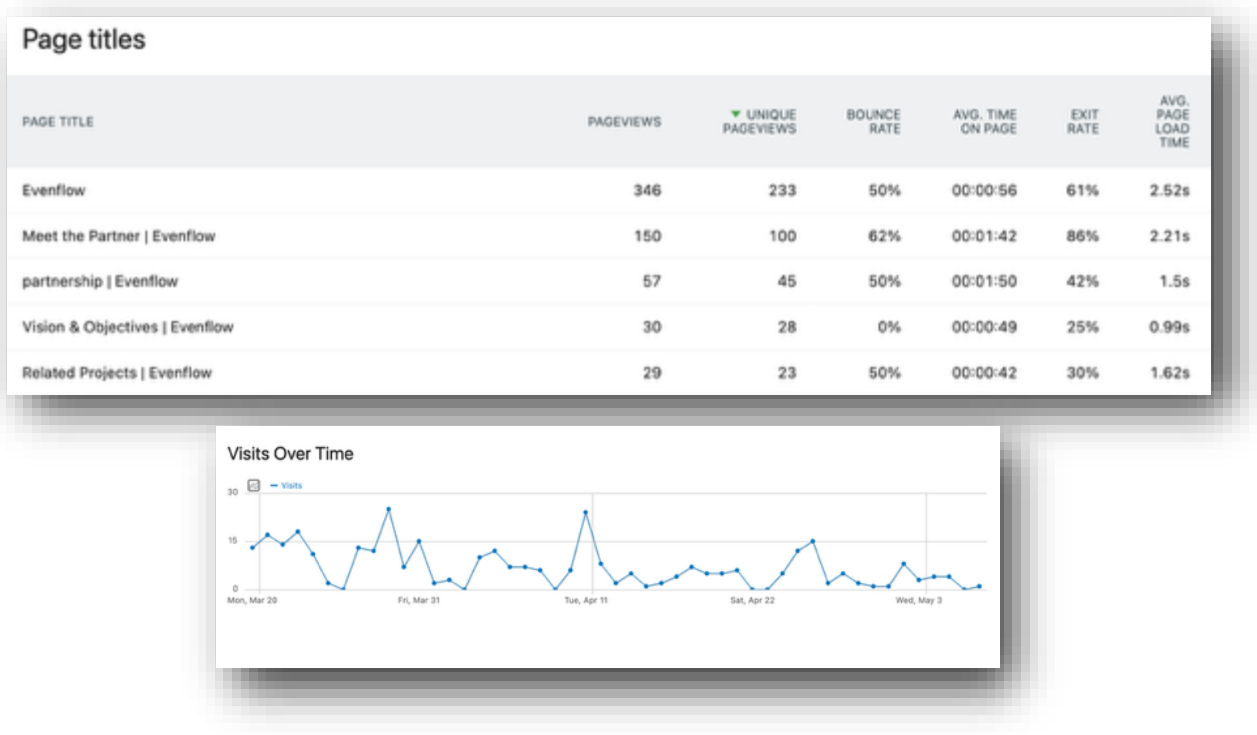

Figure 9: Meet the Partner analytics.

Use Cases Campaign

As part of the project's communication and dissemination strategy, a dedicated *Use Cases Campaign* was launched to present the project's applications in an accessible and engaging way. The campaign aimed to familiarise a wider audience with the project's real-world impact through clear storytelling and appealing visuals. A dedicated webpage – *Use Cases in a Nutshell* – was developed to serve as the central hub, featuring an introductory overview page and three subpages highlighting each individual use case. To maximise outreach, tailored visual banners were designed for each use case, and a coordinated social media plan was implemented, featuring one use case per week across the project's channels. This approach allowed for consistent engagement, increased visibility, and a deeper understanding of how the project's outcomes translate into tangible benefits for different sectors and end-users.

Figure 10: Use Cases visuals.

Figure 11: Dedicated Use Cases page on EVENFLOW website.

Through this campaign, we not only raised awareness of the project's results but also increased traffic to our website and expanded our follower base.

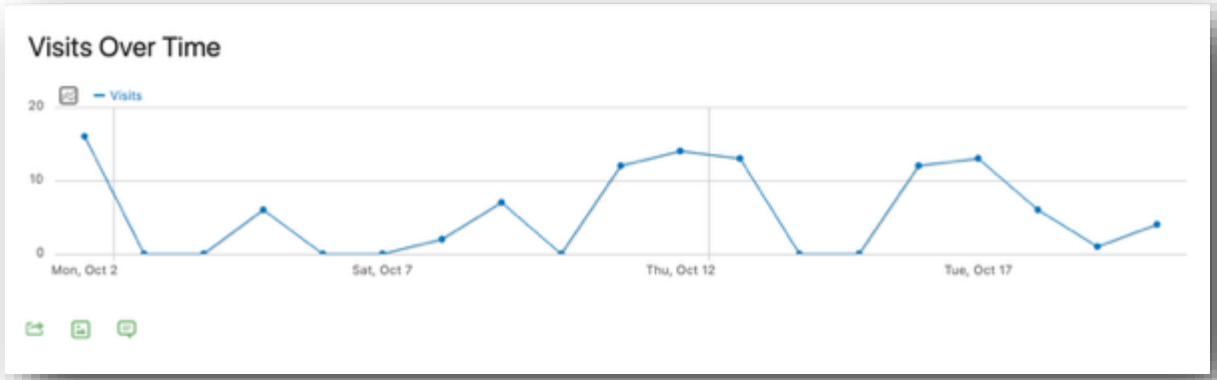

Figure 12: Website traffic during the Use Cases campaign.

Subscribe to our Newsletter Campaign

To further strengthen community engagement and ensure continuous information flow, a *Subscribe to our Newsletter* campaign was developed and implemented. The main objective was to expand the project's audience and encourage stakeholders, partners, and interested individuals to stay informed about ongoing activities and results. Dedicated visual banners were designed and disseminated across the project's website and social media platforms to promote subscriptions and increase visibility. Through this campaign, the project successfully fostered a growing network of followers and enhanced the regular dissemination of updates, events, and key achievements.

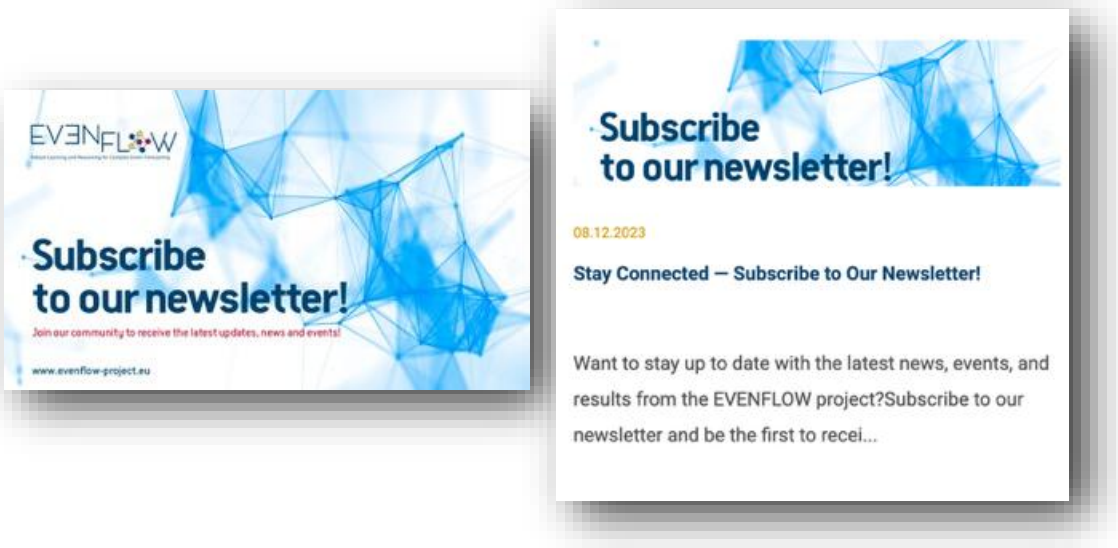

Figure 13: Subscribe to our Newsletter visuals and news item on website.

EVENFLOW at Bluesky Campaign

Dedicated visuals and introductory posts were created to announce the project's presence in Bluesky, attract followers, and stimulate interaction around project updates and outcomes. This effort contributed to expanding EVENFLOW's visibility within new online environments and fostering dialogue with a wider spectrum of stakeholders.

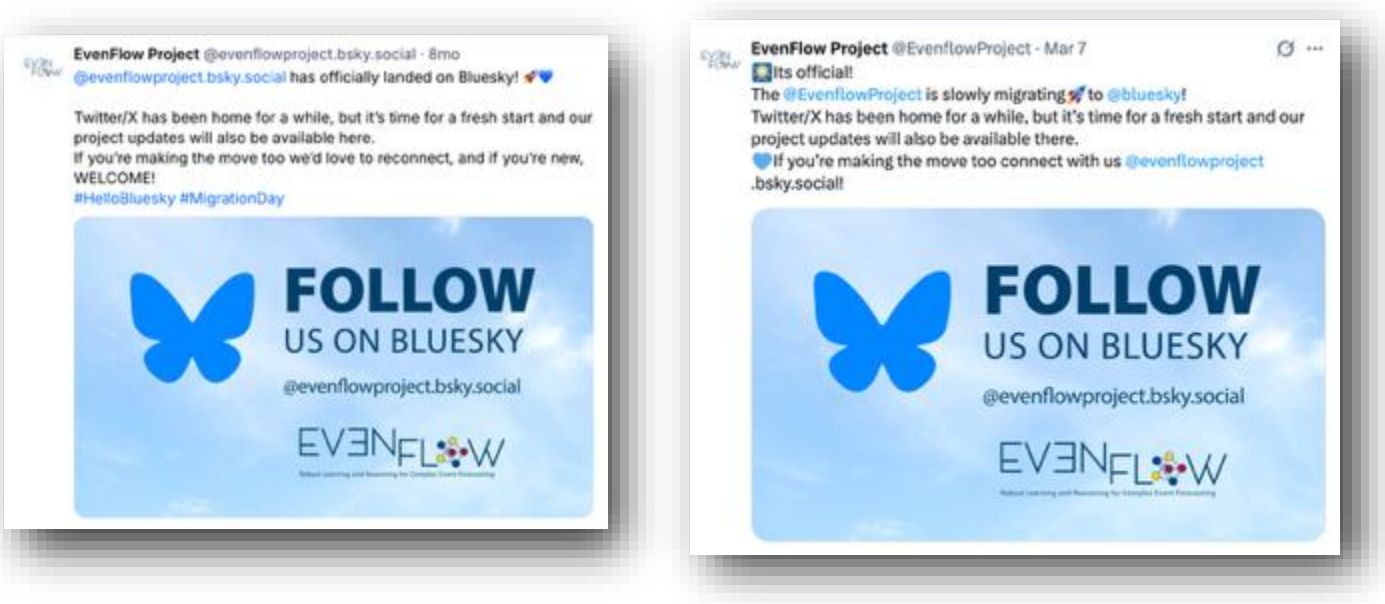

Figure 14: Posts on Bluesky and X about EVENFLOW migration.

EVENFLOW AI Tools Campaign

A key pillar of the project's dissemination strategy was the *EVENFLOW AI Tools Campaign*, designed to showcase the suite of AI solutions developed within the project and highlight their potential for real-world application and exploitation. To maximise visibility and accessibility, a dedicated webpage titled *EVENFLOW AI Tools* was created under the *Resources* section of the project's website featuring three subcategories – one for each tool –, serving as a central hub for information and materials related to the AI tools. The campaign included a coordinated social media rollout introducing each AI tool in an accessible and engaging way, supported by eye-catching visuals, a dedicated newsletter, and a short video showcasing their main features and benefits. This dynamic communication approach helped bring the tools closer to the audience, clearly showing how they can be applied in real contexts. By presenting the tools' value and usability, the campaign played a key role in paving the way for their future exploitation and in sparking interest among potential users and partners.

Figure 15: EVENFLOW AI Tools visuals and dedicated page on website.

Meet the TrustWorthy AI Cluster campaign

As part of its joint dissemination efforts, the TrustWorthy AI Cluster launched the *Meet the Cluster* campaign, developed collaboratively by EVENFLOW and its sibling projects. The campaign aimed to introduce the Cluster to the wider public and highlight the shared mission of the nine Horizon Europe projects working on trustworthy AI. To achieve this, the partners co-created a unified set of communication materials, including a dedicated news item, coordinated social media posts, a joint press release, and a suite of visual assets designed to strengthen the Cluster's collective identity. These coordinated actions not only enhanced visibility but also showcased the strong collaborative spirit driving the Cluster's activities. More information on the collaboration with the sibling projects is provided in Section 2.11.1.

Figure 16: Meet the TrustWorthy AI Cluster campaign - comms material.

White Papers Campaign

As part of EVENFLOW wider communication efforts regarding its results, the Comms Team planned and delivered a dedicated campaign to highlight the three white papers developed within the project by the partners. The scope of these documents was to be both appealing and easy to navigate, so the Comms Team enhanced them with clear visuals and designs that reflect the EVENFLOW brand. To support their dissemination, it was created a set of engaging banners and graphics to share across EVENFLOW social media channels, newsletter, and website. Through this more approachable and visually consistent communication, the project's findings were intended to be made more appealing to the targeted audiences.

Figure 17: Visuals for White Paper, banner, cover and back cover.

2.8 Communication Material

2.8.1 Visuals: Project Logo | Templates | Banners | Brochures

The EVENFLOW project logo has been created early in the project, and three versions of the logo have been presented to the partners during the kick-off meeting of the project in October 2022. The partners chose the logo shown below through a voting procedure. An extensive analysis of the project's visual identity, including the logo and the produced material for the project (electronic banners, roll up banners, virtual backgrounds etc.) is provided in *D2.1 Project Presentation and Website*, page 9-11.

Figure 18: EVENFLOW logo.

To embed the project's brand identity across communications, several project templates have been produced to ensure consistency across partner usage including:

- A PowerPoint presentation
- Meeting agenda
- Meeting minutes
- Deliverables

The templates have been made available in the joint folder for ease of access by all partners since the very start of the project.

A series of electronic banners have been created to help promote the project online, bearing the project's brand identity, logo and colours. These banners were to use on social media, organisational websites, announcements about the project, etc. The project main logo, together with different versions are also available on the EVENFLOW website under the section *Media kit* for use by journalists or other stakeholders.

A virtual background has been created for online meetings (through platforms such as Zoom, Webex, Google Meet and Microsoft Teams) to enhance the look and feel of the project when participating in meetings with external stakeholders or virtually presenting at events.

Promotional material plays a pivotal role in enhancing recognition and engagement. It boosts brand visibility, and it fosters a sense of community among stakeholders, encouraging engagement and support. Additionally, these items act as effective promotional tools, spreading awareness and sparking interest in the project's mission. Two roll up banners have been created and used through the several physical events that EVENFLOW partners have

attended, as well as an informative brochure about the project and the Use Cases in its first year of operations which has been distributed online. This flyer included information on:

- About the project info
- Info regarding the TrustWorthy AI Cluster
- Introduction to the Use Cases and a QR code for further info
- The partnership and contact information
- QR code to the website and handles to social media

Figure 19: EVENFLOW e-flyer.

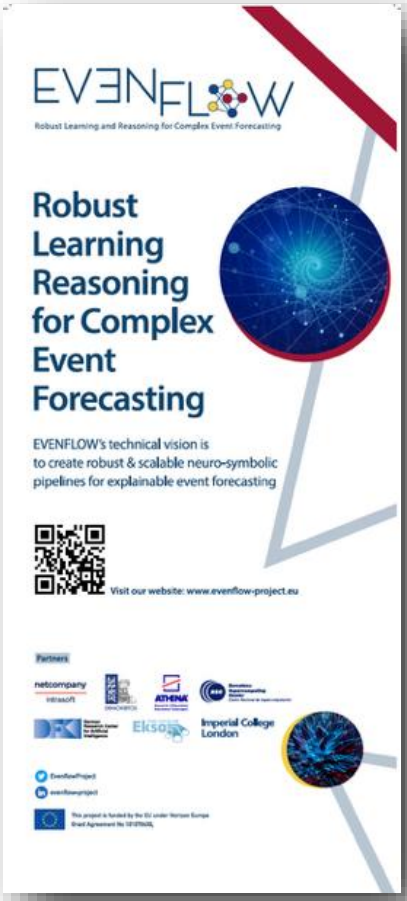

Figure 20: Roll up banner design and on-site placement at an event.

2.8.2 Newsletters

Newsletters have been sent for various reasons. The newsletters were informing about publications, events and highlighted project progress. These electronic newsletters have been disseminated to users who have willingly subscribed to the project newsletter through the relevant section on the website which is linked to a secure database (Mailchimp) and to NCSR's organisational mailing lists (3500 recipients in total). Newsletters are a tool that can assist with creating a community around the project and can help establish its sustainability and impact in the long term. A total of six newsletters have been sent to date, and published versions are available on the EVENFLOW website under the *Newsletter* subcategory.

Figure 21: Newsletters subcategory on EVENFLOW website.

2.8.3 Media Relations

A *Media kit* has been created and is accessible for public use via the website. The kit includes the EVENFLOW logo in various formats, project press releases, as well as banners and visuals for social media in various sizes.

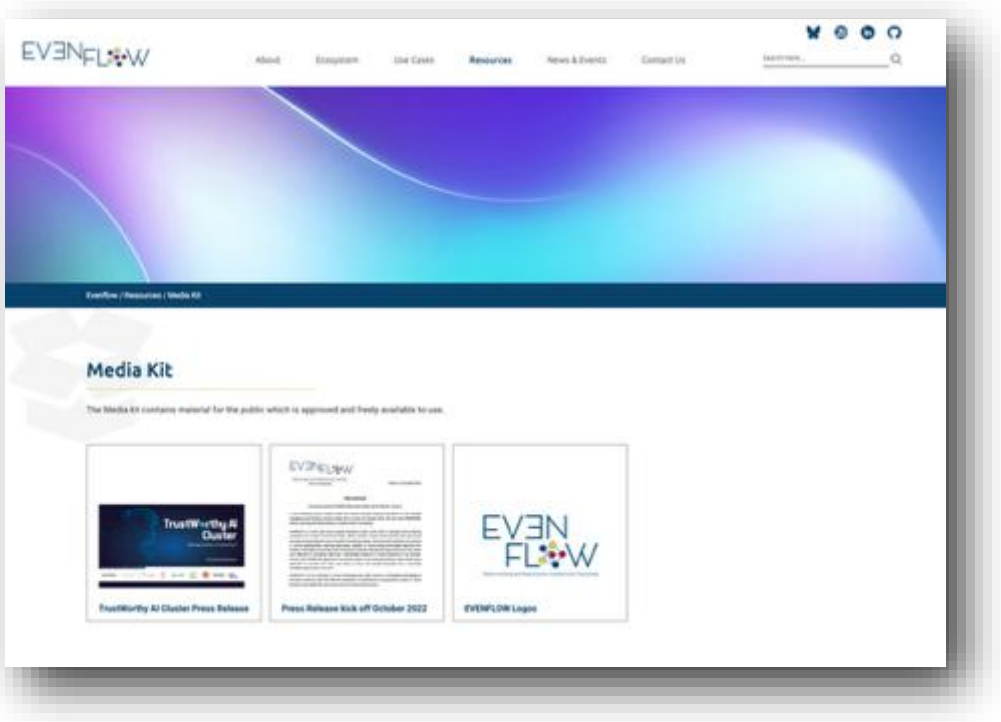

Figure 22: Media Kit subcategory on EVENFLOW website.

2.9 Event Participation / Organisation: Conferences | Workshops | Webinars | Meetings

Scientific work presentation in workshops, conferences, webinars and other occasions is yet another dissemination and communication activity and a key mechanism of engagement with the communities. In its whole duration, EVENFLOW project has already participated and presented in several virtual events, physically attended events and organised workshops. A total number of 42 participations of EVENFLOW partners in conferences, workshops, webinars and other events was achieved, which is a significant result relating to the dissemination KPI set (>20 in total).

EVENFLOW has been actively represented across a wide range of international conferences, workshops, and symposia, showcasing its advancements in AI, data science, and bioinformatics. To name a few, in 2023, partners such as BSC, ARC, and NCSR promoted the project's objectives and research results at multiple venues. Highlights include presentations at the AI Conference in Coruña (Spain), Generative Modeling Summer School - GeMSS 2023 in Copenhagen, and SMiLe– Spring workshop on Mining and Learning 2023 (Netherlands). EVENFLOW was also featured at 31st Annual Intelligent Systems For Molecular Biology and the 22nd Annual European Conference on Computational Biology (ISMB/ECCB 2023) (France) and during a visit from a Chinese Delegation on AI at NCSR. The project's publications gained visibility through presentations at the 26th European Conference on Artificial Intelligence – ECAI 2023 (Poland) and the 32nd ACM International Conference on Information and Knowledge Management (UK), with additional contributions from ARC partners at 10th IEEE International Conference on Data Science and Advanced Analytics (DSAA'23, Greece) and

EMBO 2023 (Spain) and the 30th International Symposium on Temporal Representation and Reasoning (TIME 2023, Greece) along with NCSR partner.

The project continued its strong dissemination in 2024, contributing to high-profile venues such as the 27th International Conference on Extending Database Technology (EDBT 2024, Italy), the 12th International Conference on Learning Representations (ICLR 2024, Austria), and the IEEE Big Data 2024 (USA). EVENFLOW partners also engaged in community outreach at the 5th ACM Europe Summer School on Data Science and the HIAS AI Summer School (Greece). Scientific results were presented at the 23rd European Conference on Computational Biology (ECCB 2024, Finland), Data Week 2024 (Luxembourg), and the 32nd Mediterranean Conference on Control and Automation (MED'24, Greece), emphasising the project's interdisciplinary impact.

In 2025, EVENFLOW has maintained a strong international presence with contributions to the 39th Annual AAAI Conference on Artificial Intelligence (AAAI 2025, USA), the 51st International Conference on Very Large Data Bases 2025 (VLDB 2025, UK), and recognition at the Neurosymbolic Learning and Reasoning Conference (NESY 2025, California). Additional engagements include the 1st German Robotics Conference (Germany), the SMiLe 2025 Workshop (Netherlands), the ISCB Student Council Symposium 2025 (UK), and the 60th Summer School of NCSR (Greece), further extending its influence across AI, robotics, and computational biology communities.

Table 3: EVENFLOW participation in events.

	Event	Date	Partner participated
1	Paving the way towards the next generation of R&I excellence in AI, Data and Robotics Launch Event, online	17 October 2022	NCSR
2	AI conference in Coruña, ES	29 December 2022	BSC
3	Clustering – Connection Bridge, online	2 March 2023	NCSR
4	BSC Severo Ochoa International Doctoral Symposium 2023, ES	9-10 May 2023	BSC
5	SMiLe 2023 – Spring workshop on Mining and Learning 2023, NL	30 May – 1 June 2023	NCSR
6	19th IEEE International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT 2023), CY	20 June 2023	Netcompany
7	GeMSS 2023, DK	26-30 June 2023	BSC
8	AI Ecosystem Forum 2023, GR	26-30 June 2023	Netcompany, NCSR
9	31st Annual Intelligent Systems for Molecular Biology and the 22nd Annual European Conference on Computational Biology (ISMB/ECCB 2023), FR	23 – 27 July 2023	BSC

	Event	Date	Partner participated
10	Chinese Delegation of AI during visit to partner NCSR Demokritos, GR	26 September 2023	NCSR
11	30th International Symposium on Temporal Representation and Reasoning (TIME 2023), GR	25 –26 September 2023	NCSR, ARC
12	26th European Conference on Artificial Intelligence – ECAI 2023, PO	30 September – 4 October 2023	NCSR
13	10th IEEE International Conference on Data Science and Advanced Analytics – DSAA 2023, GR	9-13 October 2023	ARC
14	32nd ACM International Conference on Information and Knowledge Management, UK	21-25 October 2023	ARC
15	AIoD Community Forum, IT	13-14 November 2023	NCSR
16	DIH4AI: Enhancing Trustworthy AI and the AIoD Platform, online	29 November 2023	NCSR
17	EMBO 2023 workshop, hybrid	26 November – 1 December 2023	ARC
18	ULTIMATE project: Trustworthy AI workshop, online	2 February 2024	NCSR
19	AIoD Technical Contributors Board meeting, online	23 February 2024	Netcompany
20	27th International Conference on Extending Database Technology – EDBT 2024, IT	24-28 March 2024	ARC
21	12th International Conference on Learning Representations – ICLR 2024, AUT	7-11 May 2024	ICL
22	Trustworthy AI: Landscaping verifiable robustness & transparency, online	29 May 2024	NCSR
23	32nd Mediterranean Conference on Control and Automation (MED'24), GR	11-14 June 2024	DFKI
24	European Convergence Summit 2024, Online	19 June 2024	NCSR
25	5th ACM Europe Summer School on Data Science, GR	11 July 2024	ARC
26	HIAS AI Summer School 2024, GR	1-3 July 2024	NCSR
27	23rd European Conference on Computational Biology (ECCB 2024), FI	16-20 September 2024	BSC
28	EVENFLOW Use Case webinar on Personalised Medicine	1 October 2024	BSC, NCSR

	Event	Date	Partner participated
29	European Big Data Value Forum (EBDVF 2024), HU	2-4 October 2024	NCSR
30	Data Week 2024, LU	10 Dember 2024	Netcompany
31	2024 IEEE International Conference on Big Data (IEEE BigData '24), USA	15-18 December 2024	ARC
32	Future-Ready: On Demand Solutions with AI, Data, and Robotics, hybrid, BE	18-19 February 2025	NCSR
33	39th Annual AAAI Conference on Artificial Intelligence (AAAI-25), USA	25 February – 4 March 2025	ARC
34	1st German Robotics Conference, DE	13-15 March 2025	DFKI
35	SMile 2025 workshop, NL	4-6 June 2025	NCSR
36	60th Summer School of NCSRD, GR	8 July 2025	NCSR
37	ISCB Student Council Symposium 2025, UK	20 July 2025	BSC
38	51st International Conference on Very Large Data Bases 2025 (VLDB 2025), UK	1-5 September 2025	ARC
39	Neurosymbolic Learning and Reasoning Conference (NESY 2025), USA	8-10 September 2025	NCSR, ICL
40	1st International Workshop on Advanced Neuro-Symbolic Applications (ANSyA 2025), IT	25-30 October 2025	NCSR
41	ARC hosts TU Berlin/BIFOLD for a deep dive into post-cloud streaming technologies, GR	10-11 November 2025	ARC
42	EVENFLOW Webinar: Showcasing AI Tools developed within the project, online	25 November 2025	Netcompany, ICL, ARC, NCSR

Overall, the project's participation in these events has helped disseminate its research achievements, present innovative use cases and publications, and strengthen collaborations within the European and global AI and data science ecosystems. For every participation in workshops, conferences, webinars and other occasions, a news item on the EVENFLOW website has been created and posts on all EVENFLOW social media channels, to inform wider audiences and promote EVENFLOW activities further.

Figure 23: EVENFLOW participation in various events.

The EVENFLOW project has also organised and co-hosted a suite of events aimed at disseminating key research outcomes and tools. In March 2023, EVENFLOW together with NCSR hosted a hybrid talk titled *Selected Techniques for Relational & Text Data Analysis*, featuring invited experts from the Jožef Stefan Institute in Ljubljana. In May 2024, the project co-organised a webinar titled *Trustworthy AI: Landscaping Verifiable Robustness & Transparency*, in collaboration with the ADRA-e ecosystem and the wider TrustWorthy AI Cluster, offering guidance on open-innovation and trustworthy AI practices. Later, in October 2024, EVENFLOW held a Use Case Webinar on *Personalised Medicine – Breast Cancer: Current Challenges*, in which clinical and AI experts discussed the project's work on one of its major use-cases. Most recently, in November 2025, ARC partner organised a visit from the Technische Universität Berlin and BIFOLD – Berlin Institute for the Foundations of Learning and Data, where EVENFLOW results were presented to researchers and students. Lastly, also in November, the project organised, a webinar *Showcasing AI Tools Developed Within the Project*, featuring live demonstrations of the Scalability Toolkit, the Reasoning-Based Forecast Interpreter and the Verification Tool for hybrid neuro-symbolic AI approaches. Across all these events, EVENFLOW has demonstrated a commitment to fostering dialogue among stakeholders and showcasing research outputs.

Figure 24: Visuals created for organised events.

2.10 Publications & Deliverables

Publications represent an important outcome of the EVENFLOW project**,** showcasing the scientific and technical advances achieved throughout its duration. A total of 21 publications has been produced by the EVENFLOW partners. To increase the visibility and impact of EVENFLOW work, the consortium had established a dedicated Zenodo community, enabling the scientific community and the wider public to easily access and engage with EVENFLOW results. All project publications were also available directly on the EVENFLOW website**,** ensuring transparency and open access to EVENFLOW research outputs. To further highlight these achievements, we launched a *Publications Campaign*, through which each new publication was promoted across EVENFLOW's communication channels, helping to broaden our reach and disseminate EVENFLOW findings to a wider audience.

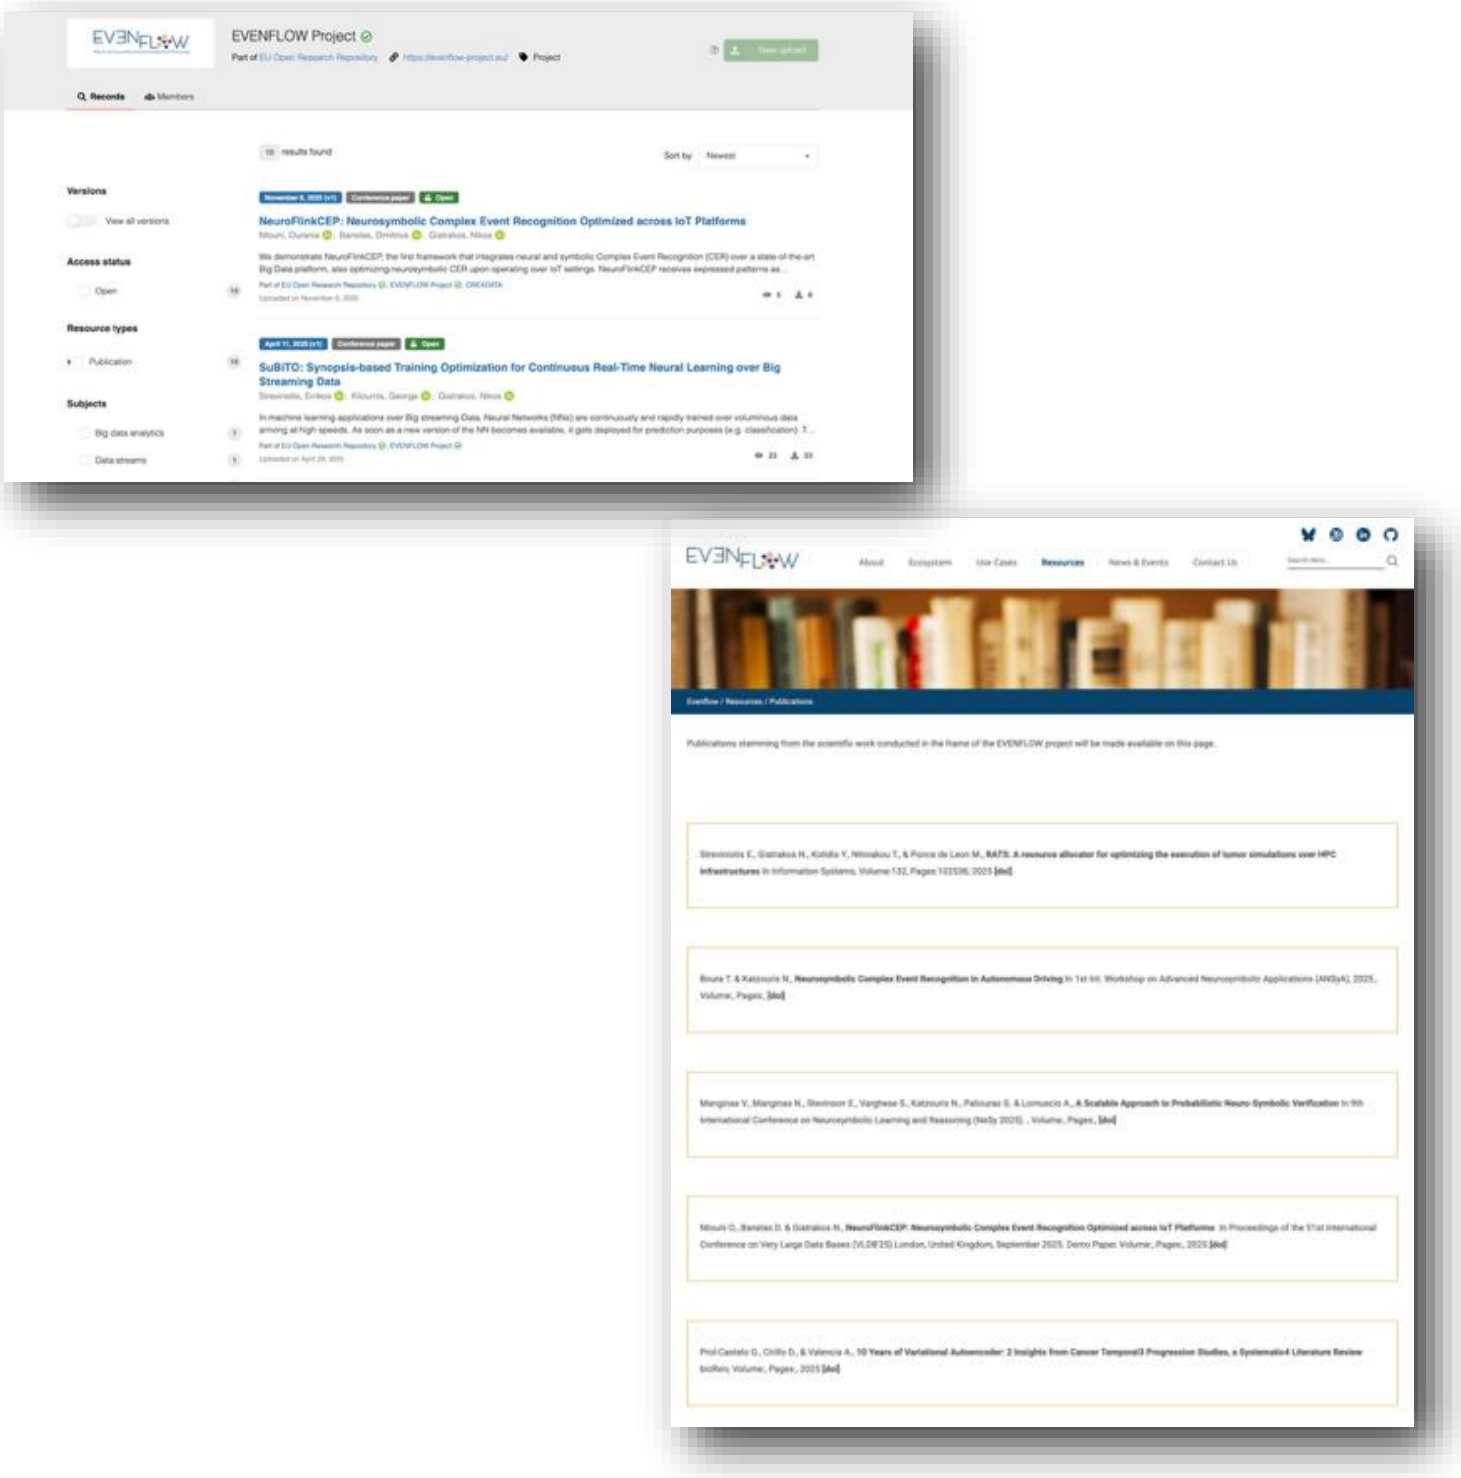

Figure 25: EVENFLOW Zenodo community and Publications on website under Resources category.

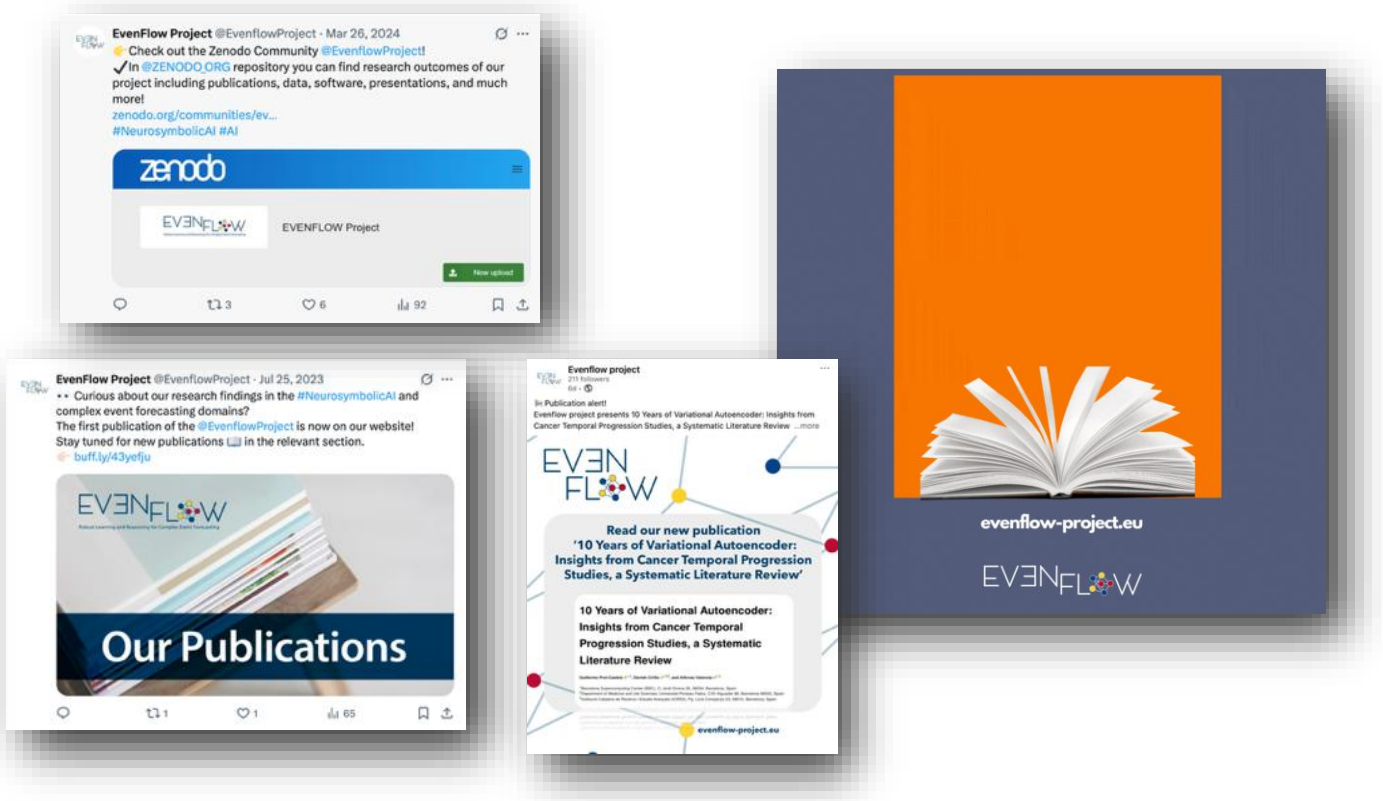

Figure 26: Social media post about publications and created visuals.

The public deliverables produced by the EVENFLOW project were uploaded on the project website under the *dedicated sub-section* category upon review approval by the EC.

2.11 Collaboration with the Ecosystem

2.11.1 The TrustWorthy AI Cluster

Nine projects funded under the Horizon Europe call HORIZON-CL4-2021-HUMAN-01-01, helped pave the way for the greater acceptance of Artificial Intelligence (AI) across Europe. These projects shared a common goal of shaping AI solutions that were innovative, human-centric, and dependable. Among others, their main focus was on verifiable robustness, energy efficiency, and transparency. Each of these projects brought a unique perspective to the broad realm of trustworthy AI. Driven by the collective aim of fostering trust and confidence in AI technology, these initiatives were geared towards enhancing transparency, explainability, accountability, safety, and performance in systems using AI.

The following eight projects were part of the call along with EVENFLOW, and each project provided solid scientific solutions complemented by tools and processes for design, testing, validation, certification, software engineering methodologies, and real-world applications.

- AUTOFAIR | Exploring the ethical use of AI
- ENEXA | Efficient Explainable Learning on Knowledge Graphs
- REXASI-PRO | REliable & eXplainable Swarm Intelligence for People with Reduced mObility

- SAFEXPLAIN | Safe and Explainable Critical Embedded Systems based on AI
- SUSTAINML | Sustainable Machine Learning
- TALON | Autonomous and Self-organized Artificial Intelligent Orchestrator for a Greener Industry 4.0
- TUPLES | TrUstworthy Planning and scheduling with Learning and ExplanationS
- ULTIMATE | mUlti-Level Trustworthiness to IMprove the Adoption of hybrid arTificial intelligencE

All nine projects joined forces to form the **TrustWorthy AI Cluster**, where, by pooling their collective expertise, visions, and communication channels, they aimed to amplify their overall impact. The collaborative spirit of the initiatives was evident as they engaged in joint efforts to understand each other's objectives and identify opportunities for synergy. Through these coordinated "cluster" activities, the projects worked to establish a cohesive strategy aligned with the objectives of the Horizon Europe call and the broader Digital Europe Strategy**.**

EVENFLOW took the initiative to bring together and coordinate these nine projects, fostering a unified identity and joint communication activities. EVENFLOW developed the cluster's logo and visual identity, including banners, roll up banner for physical events, presentation templates, and a shared repository to facilitate collaboration. It also organised and chaired monthly coordination meetings to ensure smooth communication among the sibling projects. Within this collaboration, a comprehensive outreach campaign was launched to introduce the TrustWorthy AI Cluster to the wider public (as mentioned above in Section 2.7.3Communication & Dissemination Campaigns). This included the creation of a joint press release, cross-references across all sibling projects' websites, and a long-running social media campaign where the projects promoted each other's work and shared regular updates.

Figure 27: Visuals and comms material created by EVENFLOW for the TrustWorthy AI Cluster (roll up banner, presentations templates, logo, shared folder, banners).

Figure 28: Dedicated page for presenting the cluster on EVENFLOW website.

Moreover, the projects participated collectively in several events under the TrustWorthy AI Cluster umbrella, further strengthening their visibility and impact within the European AI research community. These participations served to enhance the cluster's visibility, promote synergies, and showcase its contribution to Europe's trustworthy AI ecosystem.

The collaboration was first discussed during the first Siblings Cluster online event held in March 2023, which highlighted the shared goals of the participating projects. Following this, EVENFLOW represented the cluster in several high-level European initiatives, including the DIH4AI event on Enhancing Trustworthy AI and the AI-on-Demand (AIoD) Platform in November 2023 (online), and the AIoD Community Forum also in November 2023 (Italy), where project representatives met with AI practitioners and stakeholders to discuss transparency, verifiable robustness, and explainability in AI systems.

Figure 29: EVENFLOW representing the cluster in events.

The cluster also co-organised an online webinar titled *Trustworthy AI: Landscaping Verifiable Robustness & Transparency* in May 2024, which aimed to show the progress in the development of Trustworthy AI technology and introduce verifiable robustness and transparency in the AI domain, while it was addressed professionals, researchers, and stakeholders interested in AI, Data, and Robotics, but also welcomed participants from academia, industry, and government sectors seeking insights into Research and Innovation Actions (RIAs) projects. Evangelia Markidou, from the European Commission, welcomed the participants, while the three flagship projects **ADRA-e**, **AI4Europe** and **DeployAI** elaborated on *How is Trustworthy AI* manifested in these flagship projects.

Figure 30: Webinar Trustworthy AI: Landscaping verifiable robustness & transparency.

Continuing its outreach, EVENFLOW and its sibling projects jointly participated in the European Convergence Summit 2024, an online event held in June 2024, and later in the European Big Data Value Forum (EBDVF) 2024 in October (Budapest), where they presented their latest achievements at the AI-on-Demand booth.

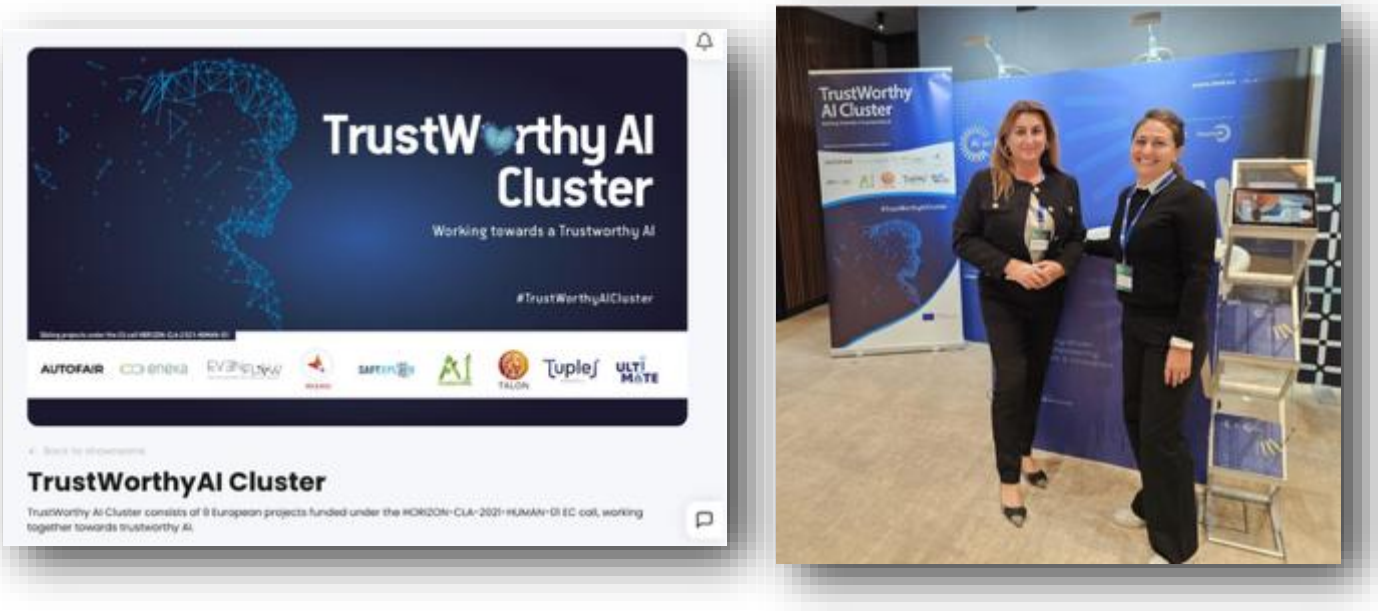

Figure 31: TrustWorthy AI Cluster at European Convergence Summit 2024 and EBDVF 2024.

The collaboration extended to additional joint activities, including participation in a Trustworthy AI workshop with sibling project ULTIMATE in February 2024 (online), where ethical and technical dimensions of hybrid AI systems were discussed, and the ADRA-e event *Future-Ready: On Demand Solutions with AI, Data and Robotics* in February 2024 (Brussels and online), in which the cluster presented its progress to the wider European AI community.

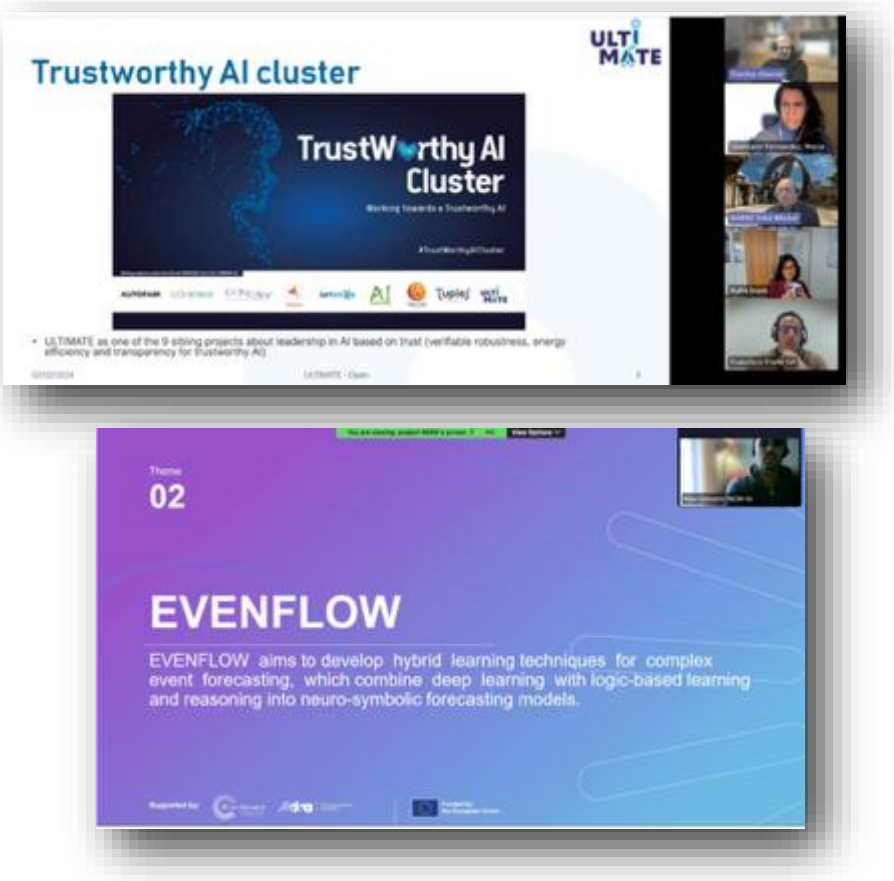

Figure 32: TrustWorthy AI Cluster at Trustworthy AI workshop with sibling project ULTIMATE and Future Ready event.

Through these joint participations, the TrustWorthy AI Cluster strengthened its presence across Europe, fostered exchange between projects, and contributed to shaping a coherent narrative around human-centric, transparent, and dependable AI.

2.11.2The AI on Demand Platform

At a very early stage of the project, EVENFLOW was included and presented—with an introductory video—in the Launch Event**:** *Paving the Way Towards the Next Generation for R&I Excellence in AI, Data and Robotics*, co-organised by ADRA-e and AI4Europe Coordination and Support Actions (CSAs). The event aimed to introduce and take stock of the newly funded R&I landscape and to identify common themes and challenges for future activities and collaboration with the AI-on-Demand (AIoD) Platform.

From the very beginning of the project, EVENFLOW also established and actively maintained its profile on the AIoD Platform, regularly updating its page with news, events, tools, and project outcomes to ensure continuous visibility and engagement with the wider AI community. EVENFLOW has also created a dedicated page on its website and a dedicated section in its Newsletter, to highlight its connection with key European lighthouse initiatives —including the **AI-on-Demand (AIoD) Platform** and **ADRA-e**— and to provide updates from

the community. By maintaining this visibility, EVENFLOW helps strengthen collaboration, knowledge exchange, and interoperability with major strategic initiatives in trustworthy AI.

Figure 33: EVENFLOW profile on AIoD and SM post for content upload on AIoD.

Figure 34: Dedicated page under Ecosystem category on EVENFLOW website.

Beyond its online presence, EVENFLOW participated in key AIoD-related activities, including a presentation at the AIoD Community Forum held in November 2023 in Italy, where the project's objectives and contributions were introduced to the platform's stakeholders. Furthermore, EVENFLOW took part in the AIoD Technical Contributors Board meeting in February 2024 (online), reinforcing its technical collaboration and commitment to the AIoD ecosystem. As part of the joint webinar organised with the TrustWorthy AI Cluster**,** EVENFLOW also invited representatives from the AIoD Platform (AI4Europe and DeployAI), as well as from ADRA-e, to present their perspectives—fostering closer ties and knowledge exchange between the cluster and the broader European AI community.

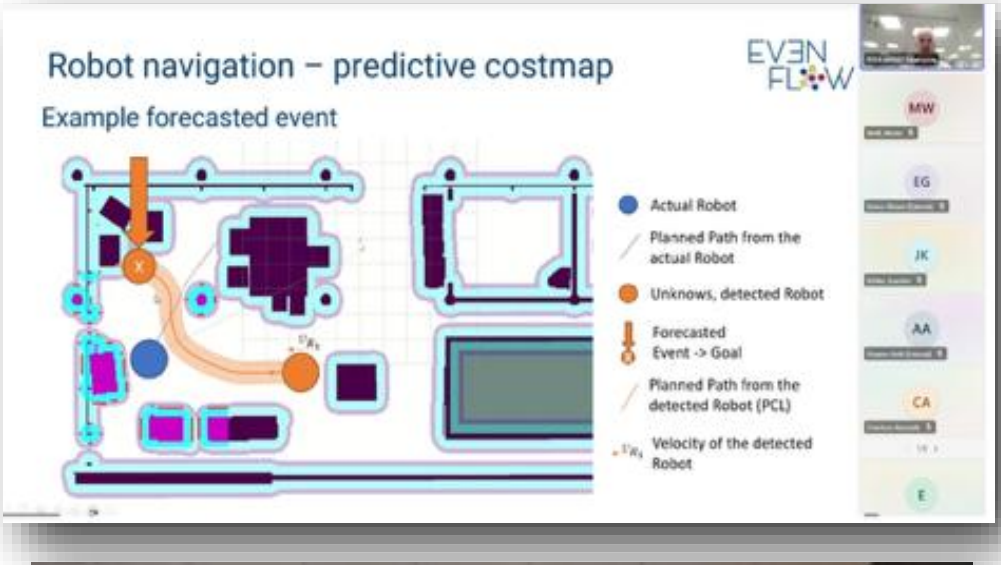

Figure 35: EVENFLOW at the AIoD Community Forum and the at the AIoD Technical Contributors Board meeting.

3 Exploitation and IPR Management Strategy

3.1 Progress towards the Exploitation and IPR management strategy of EVENFLOW

In Deliverable D2.2 "*Plan for Dissemination & Exploitation including Communication Activities*", the project defined a dynamic exploitation and IPR strategy which was refined as the it progressed. It basically set the principles for: (i) early identification and monitoring of exploitable results, (ii) a combination of open-source and proprietary IPR paths depending on partner needs and market readiness, (iii) link between exploitation, communication, and liaison activities, and (iv) market and stakeholder analysis to prepare viable post-project business and research pathways.

Based on this plan, throughout the period from M6 to M39, the consortium implemented it and progressively translated it into exploitation activities and models, as reported in the following sections of the current deliverable. The key activities/ exploitation outcomes are:

- A central activity was the systematic identification and tracking of exploitable results via the EVENFLOW innovation management log and regular partners' exchange. This work led to a consolidated Final List of Exploitable Results (Results Ownership List) that covers the integrated EVENFLOW platform, the core toolsets (Reasoning-Based Forecast Interpreter – DeepFA, Scalability Toolkit, Verification Toolkit), and use-case specific modules. Each result is described with its TRL, ownership, and licensing scheme, with the majority of the components being released as open-source under AGPL, Apache 2.0 or GPL-type licences.

- A license compatibility analysis was done in the frame of EVENFLOW (via the innovation management log), in order to analyse the licenses of the inbound technologies and carefully select the outbound licenses for the open-source projects created.

- On top of this inventory work, INTRA, in the frame of T2.2 organised a set of exploitation workshops and structured questionnaires with technical WPs and use-case owners (in total, 6 online internal exploitation workshops were organised). These workshops provided the insights for the detailed exploitation models per toolset and per use case, covering target users, value propositions, potential channels, IP strategies, among others. In parallel, use-case oriented exploitation models were developed to capture domain-specific opportunities.

- Based on the outcomes of the exploitation workshops, 6 different exploitation models were developed in the frame of EVENFLOW, namely 3 for the core toolsets (Reasoning-based forecast interpreter, scalability toolkit, and verification toolkit) and 3 for the exploitation of the uses cases (in healthcare, Industry 4.0 and infrastructure management).

- At organisational level, individual exploitation plans were collected from all partners using a template and consolidated into a table of exploitation intentions. These plans confirm that, given the consortium composition (mostly research centres and universities), the dominant pathways are:

- reuse of the EVENFLOW toolsets as **background IP in new Horizon Europe and national R&I projects**,
- **maintenance and evolution of the open-source repositories** (DeepFA, Scalability Toolkit, Verification Toolkit and selected use-case code) to build communities of users and contributors, and
- **consulting and custom development services** around neurosymbolic forecasting, scalable training and verification, occasionally linked to existing or future spin-offs.

- To support the adoption of the open-source projects in EVENFLOW, in collaboration with the dissemination task, the marketing material of the core project tools and toolsets was updated, presenting who can use them, what is the value and benefit and how the toolsets can be accessed, as presented below:
  - Reasoning-Based Forecast Interpreter
  - Scalability Toolkit
  - Verification Tool

These exploitation activities were supported by the broader dissemination and liaison work of WP2. The dedicated "EVENFLOW AI Tools" campaign, the presence of the tools and datasets in the project website and Zenodo community, and the collaboration with the TrustWorthy AI Cluster and AI-on-Demand platform, all contribute to the long-term visibility and adoption of the exploitable results and are consistent with the initial D2.2 commitment to leverage European ecosystems and open-access channels. The work performed in T2.2 has remained aligned with the initial exploitation and IPR management strategy defined in D2.2.

3.2 Final list of the EVENFLOW Exploitable Results

The final list of the EVENFLOW Exploitable Results (ER) follows the latest technical developments in the project; the results' identification process was facilitated by the exploitation manager in the project (namely, INTRA) involving all partners. To decide and conclude on the results, an internal innovation management tool has been used, namely the EVENFLOW innovation management log (.xlsx template). The innovation management log was circulated to all partners, and it was used as a standardised way to record the various project technologies that could be classified as exploitable results.

Innovation Management Log - v3

EVENFLOW

Note: Inputs depicted in the following table are preliminary and not Final. Partners retain the right to revise accordingly until the end of the project and the final submission of the underlying deliverable.

			IP Management						
#	Exploitable Result Name	Current TRL	Name of Owner(s)	Open source project admin	Contributors	Opensource/ proprietary	Related Background IP as per the EVENFLOW GA and DoA	Type of Open-Source License used for production	License to be used at the Open-Source project (released version)
1	EVENFLOW Platform	TRL 3 - experimental proof of concept	INTRA	NA	NCSR, ARC, BSC, DFKI, EKSO, ICL	Proprietary	Pre-existing knowledge for the ML models for the EKSO use case	Apache 2.0 (Kafka, kubernetes, strimzi) MIT License (Influx DB)	Open documentation to be offered
2	Verification Toolkit for learning-based CEF	TRL 3 - experimental proof of concept	ICL	ICL	ICL	Open source	Verification toolkits - Venus, VeriNet	BSD 2-Clause License, Restricted License	BSD 3-Clause License (dependency with NeSy - might use GPL3)
3	Reasoning-based forecast interpreter	TRL 3 - experimental proof of concept	NSCR	NSCR	-	Open source		GPLv3 (GNU General Public License, version 3)	GPLv3 (GNU General Public License, version 3)
4	Scalability Toolkit	TRL 3 - experimental proof of concept	ARC	ARC	-	Open source	Apache License 2.0 for the synopses data engine as a service	Apache License 2.0 (for the synopses data engine as a service); Creative Commons Licenses; GNU GPL Version 3 for the rest	

Figure 36: EVENFLOW Innovation Management Log.

The innovation management log was also used as a point of reference to achieve consensus among partners on the leaders and owners per result, as well as on other Intellectual Property (IP) management areas like the potential protection measures and the software licenses. It was also used to track and guide the implementation roadmap of the EVENFLOW open-source projects, e.g., the selection of the open-source license, the README.MD files, the CONTRIBUTING.MD files, and the code of conduct files.

The innovation management log was communicated regularly to all partners (via emails and online project calls) and its content was validated by all partners. Based on this work, the final list of the EVENFLOW exploitable results is presented below. This list can also be used as the project's Results Ownership List (ROL).

Table 4: Final List of the EVENFLOW Exploitable Results.

#	Name of ER	Description
1	**EVENFLOW Platform – Integration Scenarios**	The (loosely) integrated platform of the project which combines integrated suites on all core project components (neuro-symbolic learning, verification, scalability, interpretable forecasting) connected via data sharing and management modules. It also includes documentation on all core project components. The suites cover major integration scenarios for developing and deploying CEF technologies in several application domains.

#	Name of ER	Description
	Leader(s)	INTRA
	TRL	4
	License	Copyright (software and documentation), trade secrets (integration logic)
2	**Reasoning-Based Forecast Interpreter (DeepFA)**	Neuro-symbolic forecast interpreter using symbolic finite automata to explain the sequence of conditions leading to a forecast. Produces Most Probable Explanations (MPE) and enables transparent temporal reasoning in applications such as robotics, infrastructure and medical simulations.
	Leader(s)	NCSR
	TRL	5
	License	Copyright; open source (GLPv3)
3	**SDEaaS – Synopses Data Engine as a Service**	SDEaaS is a stream-processing component implemented on Apache Flink and Dask that maintains multiple synopses in parallel to produce compact, real-time summaries of high-volume data streams. It optimizes resource usage by reusing synopses across workflows and allocating lightweight tasks that maximize CPU efficiency. The engine supports runtime creation of new synopses, dynamic plug-ins, and cross-platform JSON/Kafka-based input–output operations.
	Leader(s)	ARC
	TRL	5
	License	Copyright; open source (Apache 2.0)
4	**Synopses-Based Training Optimization**	This component accelerates the training of ML models by using data synopses instead of full data streams, significantly reducing computational load while preserving accuracy. It applies Bayesian Optimization to automatically select optimal hyperparameters such as sampling rates, synopsis type and number of epochs.
	Leader(s)	ARC
	TRL	5
	License	Copyright; open source (AGPL3)
5	**Advanced Distributed-Parallel Training**	This component accelerates large-scale ML training by distributing computation across multiple machines while maintaining (or slightly improving) model accuracy. Its main innovation is a data-driven synchronisation mechanism that significantly reduces communication delays compared to standard parallel or federated training approaches. As a result, overall training time is greatly reduced without compromising model performance.
	Leader(s)	ARC

#	Name of ER	Description
	TRL	5
	License	Copyright; open source (AGPL3)
6	**Subito Pipeline Orchestrator**	Production-level workflow orchestrator integrating SDEaaS, training optimisation and distributed-parallel training. Enables fast model retraining and deployment in streaming environments.
	Leader(s)	ARC
	TRL	5
	License	Copyright; open source (AGPL3)
7	**Verification Toolkit**	The verification toolkit can be used to assess whether forecasting models (especially neural and neuro-symbolic ones) remain reliable under small input changes, sensor noise or other perturbations. The tool identifies failure cases and stability issues before deployment, ensuring safe and trustworthy predictive behaviour.
	Leader(s)	ICL
	TRL	4
	License	Copyrights; open source
8	**Simulation & Trainable Forecasting Toolkit for Oncological Precision Medicine**	A toolkit which is comprised of two tools deployed at the health use case of EVENFLOW, namely: (i) **Variational auto-encoder (BSC):** tool that produces synthetic molecular profiles and reconstructs pseudotemporal patient trajectories from static cancer data. It helps model tumour progression when longitudinal samples are missing, enabling richer datasets for analysis and model development. (ii) **Neuro-Symbolic Forecasting on Molecular Trajectories (NCSR):** A forecasting tool that applies neuro-symbolic Complex Event methods to the synthetic patient trajectories generated by BSC.
	Leader(s)	BSC, NCSR
	TRL	4 (for both tools)
	License	Copyright; open source
9	**Prototype Neuro-Symbolic AGV Controller**	A toolkit which is comprised of three tools deployed at the industrial use case of EVENFLOW, namely: (i) **NeSy tool (NCSR)**: module that predicts AMR deadlocks by combining neural distance estimation from onboard camera images with a symbolic automaton learned from past deadlock patterns. It uses these neural inputs to forecast approaching

#	Name of ER	Description
		deadlocks more accurately than purely neural models, especially in unseen scenarios. (ii) **Multiple-robot simulation tool (DFKI)**: A 3D simulation of the SmartFactory Lab built in NVIDIA Isaac Sim, allowing realistic testing of multiple robots, machine modules, and navigation behaviours. It uses a ROS-based stack, enabling robots in simulation to react to forecasted events as they would on real hardware. (iii) **Data generation pipeline (DFKI)**: An automated pipeline that generates robot trajectories, images, and labelled events such as deadlocks and collisions. It includes task allocation, scenario triggering, and logging tools, producing the data needed to develop and evaluate forecasting models at scale.
	Leader(s)	DFKI. NCSR
	TRL	4
	License	Copyrights; open source
10	**Infrastructure Lifecycle Assessment Module (Water Pipes)**	An ML-based leakage detection and forecasting component, used to identify anomalies and analyse sensor correlations trained on datasets provided by EKSO. The component processes real-world datasets, using sensors deployed in the pipes to identify incident leakages and extract sensor correlations, and also to improve the leakage localisation precision. This component can be used as a tool for early detection and predictive maintenance to support water infrastructure operators.
	Leader(s)	EKSO, INTRA
	TRL	4
	License	Copyright; trade secrets (sensor calibration, model parametrisation)

3.3 Market Analysis of the application of Neurosymbolic AI

The EVENFLOW technologies could be classified, from a market research point of view, at the intersection of several AI markets, targeting both enabling technologies (such as neuro-symbolic AI, explainable and scalable machine learning, and AI verification) and application verticals including healthcare (precision medicine), smart manufacturing, and infrastructure monitoring (e.g., water networks). Hence, based on this segmentation, this section summarizes market sizes of the EVENFLOW technologies, with a focus on the European landscape.

Neuro-symbolic and Explainable AI are becoming more important as AI deployments move into more regulated environments. While the neuro-symbolic AI niche is still emerging, it can be classified under the broader Explainable AI (XAI) market, which was valued at $7.8 billion

in 2024 and is projected to reach $21 billion by 2030, growing at 18% CAGR.[2] Europe is driving this trend, mainly due to regulatory frameworks like the EU AI Act, which emphasizes transparency and human oversight in AI systems. In the domain of scalable and federated ML, the market for federated learning (FL) was valued at $138 million in 2024, expected to more than double to $297 million by 2030 at a 14.4% CAGR.[3] Adoption is especially strong in Europe, where GDPR and data sovereignty concerns drive uptake in healthcare, banking, and public sector AI deployments. In parallel, AI-based verification tools are becoming more important since they support the reliability and safety in AI applications. Although for this market to the best of our knowledge there were not niche-related data, it overlaps with AI governance and compliance software, a growing segment as enterprises adopt high-risk AI applications. The XAI market growth also reflects the need for verifiable systems. Tools for formal verification are increasingly used in finance (e.g., model risk auditing), aerospace, and autonomous systems, and Europe's AI Act will further mandate conformity assessments for such tools.

Among application sectors, **AI in precision medicine** is expected to grow from $2.3 billion in 2024 to $14.5 billion by 2030 (approx. 36% CAGR), driven by advances in genomics, oncology, and personalized treatments.[4] Europe is the fastest-growing region in this market due to strong public research support and initiatives such as more than one million genomes. In smart manufacturing, Europe is a global leader due to its industrial base and Industry 4.0 investments. The **European AI in manufacturing market** was valued at $1.3 billion in 2024, with a projected CAGR of 45% through 2034.[5] Finally, in **infrastructure monitoring**, especially smart water management, the global market is projected to grow from $23.7 billion in 2025 to $43.7 billion by 2030 (13% CAGR), driven by urban resilience initiatives and climate adaptation.[6] In Europe, municipalities are investing in AI-driven leakage detection, failure prediction, and asset optimization. Explainable and verifiable AI systems are particularly valued in public utilities, making EVENFLOW's interpretable forecasting technology well-suited for these contexts.

3.3.1 Competition analysis and similar solutions in the market

In this section, we identify and compare key companies and startups (especially those coming from Europe) that are developing technologies or products similar to EVENFLOW's focus or targeting the same application areas. We basically group competitors into sub-categories for clarity, though there might be some overlaps. These sub-categories include start-ups in AI (who have commercialised neurosymbolic/scalability technologies), as well as established players who use neurosymbolic AI technologies for new services.

In neurosymbolic and XAI, many European startups are early-stage (seed or Series A) focusing on specific use-cases (research automation, sales AI, etc.), often using a SaaS model. They

[2] https://www.grandviewresearch.com/industry-analysis/explainable-ai-market-report
[3] https://www.grandviewresearch.com/industry-analysis/federated-learning-market-report
[4] https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-precision-medicine-market-report
[5] https://www.polarismarketresearch.com/industry-analysis/europe-artificial-intelligence-in-manufacturing-market
[6] https://finance.yahoo.com/news/smart-water-management-global-markets-080400308.html

differentiate by offering explainability and logical consistency out-of-the-box, unlike conventional "black-box" AI solutions. In parallel, larger companies integrate similar principles into enterprise AI platforms (usually as part of a bigger sale). For example, EY applies neuro-symbolic AI via its Growth Platforms to deliver explainable decision intelligence tools that support strategic forecasting, regulatory analysis, and operational optimization across consumer, industrial, and financial sectors.[7] Open-source also plays a role, since for example, some startups build on open technologies. From EVENFLOW's perspective, these players validate a market need for hybrid AI.

At a global level, there are several startups (Extensity AI, SynaLinks, and Lakmoos AI) who offer platforms and tools in the market that combine symbolic reasoning with neural networks in various use cases, in order to improve transparency, compliance, and reliability in decision-making systems.[8] For instance, SynaLinks has built a neuro-symbolic NLP framework to support compliance-focused applications and complex automation workflows, while Lakmoos AI has developed a neuro-symbolic simulation platform for consumer behaviour modelling.

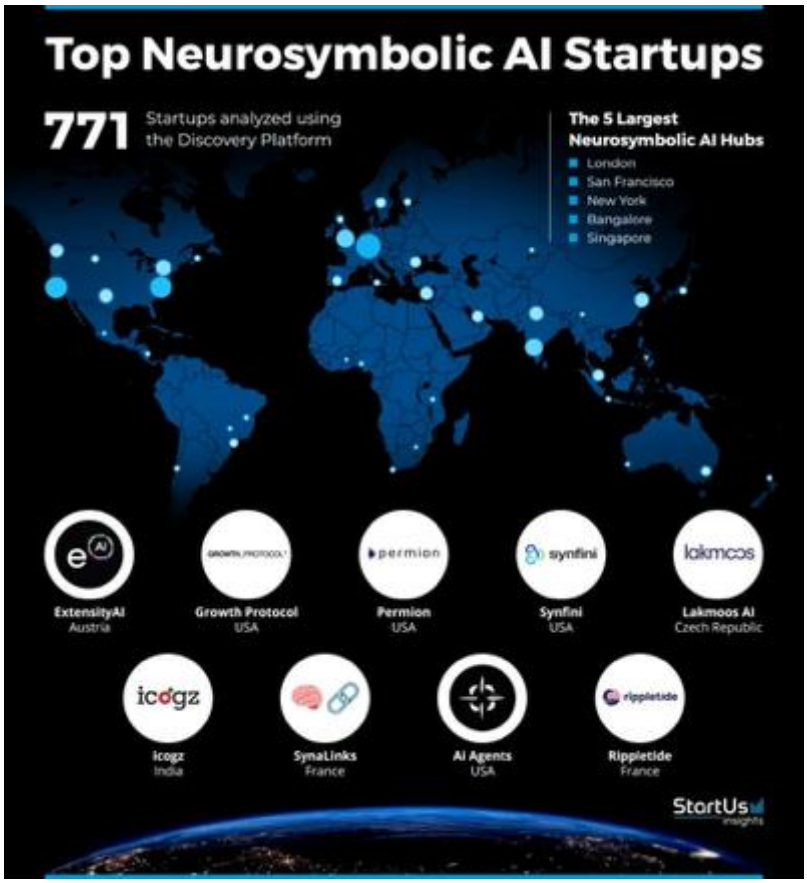

Figure 37: Top Neurosymbolic AI Startups and Hubs (source: https://www.startus-insights.com/innovators-guide/neurosymbolic-ai-companies/).

Among the most relevant (indirect) competitors to the EVENFLOW scalability toolkit are FedML, Owkin, and Seldon, each offering overlapping capabilities in distributed training, model orchestration, and scalable machine learning deployment. FedML is an open-source framework designed for federated and distributed ML, enabling collaborative model training across devices and resource-efficient scheduling, primarily targeting cross-institutional research, edge AI, and collaborative ML use cases. Owkin, a French scale-up operating in the biomedical sector, has developed Owkin Connect, a federated learning platform that allows hospitals and pharmaceutical companies to collaboratively train AI models on decentralized clinical data.[9]

[7] https://www.ey.com/en_us/services/strategy/ey-growth-platforms-neurosymbolic-ai
[8] https://www.startus-insights.com/innovators-guide/neurosymbolic-ai-companies/
[9] https://www.owkin.com/federated-learning

On the deployment side, Seldon provides an open-source MLOps infrastructure (Seldon Core) and enterprise-grade tools for scaling, deploying, and monitoring machine learning models in production.

Table 5: EVENFLOW Analysis of Similar Solutions in the Market along with their Offerings and Business Models.

Company	Offering	Business Model	Target Industries
ExtensityAI (AT)	Neuro-symbolic AI research automation platform	SaaS (cloud platform); subscription.	Knowledge work automation/ cross industry
SynaLinks (FR)	Neuro-symbolic NLP framework	Likely enterprise software licensing or SaaS.	Enterprise NLP applications (digital assistants, compliance chatbots) in regulated sectors (finance, insurance, etc. needing controlled language).
Imandra (UK/US)	Automated reasoning engine for code and AI verification; neurosymbolic logic	Proprietary SaaS platform; enterprise licensing.	Finance; software companies needing formal verification
Owkin (FR)	Federated learning platform for collaborative AI on biomedical data	Collaboration-based SaaS	Healthcare & Pharma – hospitals, pharmaceutical companies, research networks
Seldon (UK)	Open-source MLOps platform (Seldon Core) for scalable model deployment.	Open-source and Enterprise Support; on-prem or cloud deployment.	Cross-industry MLOps – used in finance, IT, e-commerce
HULO (NL)	AI leak detection for water networks; analyses sensor data to identify leaks early	B2B SaaS	Water utilities (municipal water companies, private utilities managing distribution networks worldwide).
Ekydra (FR)	AI solution suite for water management (joint venture of Veolia).	Likely solution sales to utilities – could be software license or service contract.	Water utilities (municipal water companies, private utilities managing distribution networks worldwide).

3.3.2 SWOT Analysis for the project Toolsets

Based on the market analysis above, a SWOT analysis has been designed for the EVENFLOW technologies, treating them as a whole, as a package of different technological components.

Table 6: EVENFLOW - SWOT Analysis.

INTERNAL FACTORS

STRENGTHS +	WEAKNESSES –
• EU AI Act compliant technologies, compliant-by-design, interpretable forecasting pipelines. • Exploitable results are open source (AGPL/Apache licenses), enabling community expansion, faster adoption. • The scalability toolkit achieves up to 10× reductions in training time while maintaining accuracy and lowering energy/computational costs.	• Early TRL levels (Mostly TRL 4) limiting immediate commercial adoption • Implementing end-to-end neuro-symbolic pipelines requires expertise not widely available in industry teams. • Although validated in the three use cases, longitudinal evidence (e.g., in real factories, water utilities, hospitals) is still limited, which may delay trust and large-scale adoption.

EXTERNAL FACTORS

OPPORTUNITIES +	THREATS –
• EVENFLOW aligns with upcoming regulatory requirements for high-risk AI systems, especially interpretability, robustness, and traceability areas. • EVENFLOW technologies can address broad markets where temporal, multivariate, or image-based forecasting is relevant (cross-domain application creates opportunities for new use cases/ further adoption).	• Water utilities, manufacturing, and healthcare are risk-averse and require long validation cycles and certified technologies. • The rise of multimodal foundation models (vision-language-time series) may reduce demand for specialized neuro-symbolic solutions.

3.4 Exploitation Models per EVENFLOW Toolset

3.4.1 Reasoning-based forecast interpreter

The exploitation model of ER #2 - Reasoning-Based Forecast Interpreter (DeepFA) is presented below.

Table 7: Reasoning-Based Forecast Interpreter exploitation model.

Tool/ product description	The reasoning based forecast interpreter is mainly comprised of the deepfa library, which is an open-source tool developed in the frame of EVENFLOW. The toolset symbolic finite automata (which use logical rules) and neural networks (which learn patterns from data), offering enhanced predictive capabilities via combining logical reasoning with pattern recognition. This could be beneficial when dealing with image/video timeseries data inputs, where neural networks can extract features and patterns from complex image/video representations. Deepfa is at a stage where it has been validated in a relevant environment (TRL 5). The library offers a robust framework for defining automata and transitions based on logical conditions, performing operations, and computing acceptance probabilities, including the Most Probable Explanation (MPE). The library also offers an API for defining automata, computing probabilities, and integrating neural networks. This makes the toolset useful for applications requiring a combination of logical reasoning and pattern recognition. This allows users to understand the most likely sequence of events leading to a particular outcome.
Exploitation/ use models	Sustainability of the tool will be ensured as it will be a living open-source project, maintained by NCSR after the end of EVENFLOW. Hence, there will be three main exploitation actions. First, the tool will be proposed as background IP in future Horizon Europe and EIC Transition proposals, for further validation in new industrial and scientific pilots and making its TRL to progress toward TRL 6 - 7. Second, NCSR will support a user community around the library, ensuring continuous refinement and an expanding set of use cases and technical extensions. Third, NCSR will publish peer-reviewed papers and demonstrations to position the library as a reference implementation for neuro-symbolic temporal prediction. Complementarily, NCSR may offer consulting services for organisations requiring tailored deployments, integration support or advanced configuration of the interpreter in specialised domains such as autonomous systems or infrastructure monitoring.
Background IP and Ownership Status	• NCSR acts as the sole owner of the result in the frame of the EVENFLOW project. • The tool and its components are protected by Copyrights (i.e., source code, documentation, configuration files) • The toolkit is offered open source via GitHub - GLP3 license

Target users/ customers	The user profiles identified are: • Data scientists who want to build and evaluate forecasting models for various applications (relevant researchers in general) • Start-ups/ businesses for decision making and planning Some of the use case areas identified are: • It is mostly domain-agnostic/ flexible, can be used in various use cases. • Specifically, it can address event recognition and forecasting in video/image analysis (cases which combine these 3 elements: video/images analysis & neural networks & time). • Some indicative examples could be: ✓ Autonomous Guided Vehicle (AGV) navigation in smart factories. ✓ Enhancing autonomous driving systems to recognize and predict road events. ✓ Analysing traffic camera feeds to recognize and predict traffic congestion and accidents.
Value proposition analysis	• Leverages both symbolic reasoning and neural network predictions, thus, it offers explainability (it can say basically which is the sequence of events). • The neurosymbolic approach correctly predicts the outcomes 85% of the time on the test data. ✓ For comparison, Convolutional Neural Networks (CNNs) and Long Short-Term Memory networks (LSTMs) correct predictions only 35% of the time. ✓ Transformer models for prediction have a test set accuracy of 0.68, which is better than CNN LSTM but still lower than the neurosymbolic approach. ✓ Neurosymbolic approach results in substantial cost savings due to accurate predictions and timely maintenance scheduling (2.43 times more savings compared to CNN LSTM, and 1.25 more savings compared to transformer models). • Can be applied to various types of time series data (flexible) – domain agnostic.

	• Allows customization and improvement by the community (it is an opensource project stemming from EVENFLOW).
Events and channels for promotion	The plan is to show the key features, documentation, use cases, and how contributors can support in the following channels: • Social media channels of EVENFLOW • Promotion via scientific conferences • Scientific promotion via papers • NeSy community portal • Horizon Results Platform • Y combinator portal via posts • We will also target partners' networks and find relevant collaborators to engage to test and contribute to the open-source project – they could be professors from universities and relevant research teams. • PhDs and MSc students will be also engaged to contribute.
Open-source project maintenance and governance	NCSR will be the open-source project admin, and it will do the maintenance and future support.
Exploitation next steps to increase the TRL	• Will be proposed as background IP in upcoming EU-funded R&I proposals for further validation and extension in new domains. • Targeted effort will be made to grow a developer and user community around the forecast interpreter tool. • Publishing peer-reviewed papers based on its components and performance gains; present work in conferences.
Impact assessment	Extra people involved in the implementation of the tool: • 1 PhD student + 1 MSc students • 1 research member

3.4.2 Scalability Toolkit

The exploitation model of ERs/ separate tools that comprise the Scalability Toolkit (namely ER #3 - SDEaaS – Synopses Data Engine as a Service, ER #4 - Synopses-Based Training Optimization, ER #5 - Advanced Distributed-Parallel Training, and ER #6 - Subito Pipeline Orchestrator) is presented below.

Table 8: Scalability Toolkit Exploitation Model.

Tool/ product description	The Scalability Toolkit is a set of 4 different tools (Synopses Data Engine as a Service, Synopses-based training optimization, Advanced Distributed-Parallel Training, and Subito) which are designed to enhance the scalability and efficiency of ML and data processing workflows. The expected TRL of the solution under EVENFLOW is 5. Each of the four components in the toolkit serves a distinct purpose to address a variety of scalability challenges. More specifically: • **Synopses Data Engine as a Service (SDEaaS)**: it is a component implemented both on Apache Flink and Dask, and its purpose is to maintain stream synopses, in parallel, efficiently. In that, it can rapidly provide concise summaries of incoming streams for synopses-based training optimization purposes. SDEaaS on Flink further optimizes resource usage by reusing common data synopses across multiple workflows, which as a consequence reduces redundancy and computational overhead. Each new synopsis reserves new tasks without monopolizing entire task slots, optimizing CPU core usage. The SDEaaS API allows for building and stopping synopses at runtime, dynamic code plug-ins for new synopses, and it also supports continuous and ad-hoc queries. It facilitates cross-platform execution via JSON and Kafka input/output formats. • **Synopses-based training optimization**: This component optimises the training process of ML models via leveraging data synopses. It employs advanced techniques like Bayesian Optimization to select optimal hyper-parameters, including sampling rates and the number of epochs. It can integrate with Kafka for real-time data handling; this helps fast training of models with reduced computational load, but at the same time maintaining high accuracy. • **Advanced distributed-parallel training**: This component improves the efficiency of training large-scale ML models via the distribution of the workload across multiple machines or processors. It supports parallel processing, therefore, training times can be reduced; in parallel, the model performance is maintained or even slightly improved. The key advancement compared to the prior art in the field is that the advanced, distributed training techniques built in EVENFLOW use data-driven synchronization mechanisms which reduce the communication lags and, thus, training times.

	• **Subito**: Subito integrates the above to boost the performance of production training and prediction pipelines. Each component of the toolkit has specific purpose and functionalities, and they can be used independently or in combination to address a set of scalability cases. The toolkit is provided as an open source project to the community via the EVENFLOW GitHub page: https://github.com/EVENFLOW-project-EU/Scalability-Toolkit/tree/main.
Exploitation/ use models	The main exploitation path for the Scalability Toolkit will be based on open-source business models. The toolkit is provided as an open-source project in GitHub to increase its impact and exploitation potential after the end of EVENFLOW. It is released via 4 components, under two types of licenses (AGPL and Apache 2.0). The goal is to promote the open-source project to key stakeholders (including researchers and developers) and create a community around it. We will actively promote the toolkit to academic institutions and research organizations, via conferences, workshops and new publications, to encourage its adoption. Our goal is to support and grow this community and investigate the use of the tool in other research and commercial use cases. On top of that, ARC could offer consulting services for the toolkit's configuration and custom developments, via service contracts with potential customers. The open-source nature of the toolkit eases its way to marketing channels for acquiring new customers who would like configurations and/or new functionality for scalability. Moreover, the research exploitation path will be followed by ARC. The Scalability Toolkit will be used as background IP in other Horizon Europe research projects to develop and test new functionalities and test it under different business cases. At the moment, we do not envision to establish a spin-off company that will be based on this technology. Instead, our focus will remain on supporting and expanding the open source community and further enhance the TRL of the toolkit via new research projects. Another exploitation path is that some features of the scalability toolkit are available as a free add-in in the Altair AI studio suite, a pipeline for future commercial engagements.
Background IP and Ownership Status	• ARC acts as the sole owner of the Result in the frame of the EVENFLOW project. • The toolkit and its components are protected by Copyrights (i.e., source code, documentation, configuration files) • The toolkit is offered open source via GitHub. The following open-source licenses are applied:

	○ SDEaaS: Apache 2.0 ○ Synopses-based training optimization: AGPLv3 ○ Advanced distributed-parallel training: AGPLv3 ○ Subito: AGPLv3
Target users/ customers	• The primary targets are research teams and start-ups. But we have identified potential uses by large companies as well. More specifically, some potential use cases include: ✓ **Content moderation** (e.g., on internal platforms of large companies) for large and smaller companies who develop internal platforms where they share information with their employees. The scalability toolkit is an open-source solution that could be used to address such use case, via classifying the videos/ images shared among employees, which content might need to be reviewed or flagged. ✓ **Portfolio management for start-ups and large companies in finance** could use the toolkit to make a trade-off for this case via enabling data summaries and smarter training methods. ✓ **Banking case for fraud detection** could use the toolkit to process transaction data in real time and update fraud detection models efficiently. ✓ **Telecommunications, mobile fraud detection;** it is an open-source solution so telecom operators could customize it for their systems without relying on costly vendor solutions.
Value proposition analysis	• In principle, it enables faster model training, since: ○ EVENFLOW synchronization protocols could be up to 10x faster than Vanilla protocols (benchmark) ○ From 961sec reduced to 91sec ○ With the same or better accuracy • Allows trade-offs among accuracy/ time/ monetary costs (e.g., computation costs) • Less energy consumption over other federated learning solutions • Less maintenance costs for the sensors of the end user
Events and channels for promotion	We will show the key features, documentation, use cases, and how contributors can support in the following channels: • Social media channels of EVENFLOW

	• Promotion via scientific conferences that ARC and the Technical University of Crete participate. • Scientific promotion via papers that are based on the Scalability Toolkit. • We will also target partners' networks and find relevant collaborators to engage to test and contribute to the open-source project – they could be professors from universities and relevant research teams. • PhDs and MSc students will be also engaged to contribute.
Open-source project maintenance and governance	After the completion of the EVENFLOW project, the Scalability Toolkit will continue to be actively maintained, ensuring its long-term sustainability as a valuable open-source result. The responsible parties for post-project maintenance are ARC and the Technical University of Crete (TUC). These two entities will jointly oversee the evolution of the toolkit by integrating it into ongoing academic and research activities. In particular, ARC will use the toolkit as background IP in future research, and both entities will engage students and PhDs to contribute new features, fix bugs, and publish academic papers based on the toolkit. Both ARC and TUC will monitor the GitHub repository, respond to community feedback, and address technical issues. Most probably on an annual basis. The financing for maintenance and evolution of the open-source project will primarily come from in-kind contributions (e.g., using internal students and researchers to do research), as well as EU research projects.
Open-source contributions needed	Contributors are encouraged to propose bug fixes, new features, or enhancements to the existing modules (e.g., SDEaaS, Subito, training optimization). For more details check the CONTRIBUTING.md file - Scalability Toolkit: https://github.com/EVENFLOW-project-EU/Scalability-Toolkit/blob/main/CONTRIBUTING.md.
Exploitation next steps to increase the TRL	• The scalability toolkit will be proposed as background IP in upcoming EU-funded R&I proposals for further validation and extension in new domains (e.g., federated learning, edge AI, mobile fraud detection, and industrial process optimization) – the goal here is to reach TRL 6 (or even 7) via use in large-scale demonstrators and applied research environments. • Targeted effort will be made to grow a developer and user community around the Scalability Toolkit.

	• Publishing peer-reviewed papers based on its components and performance gains; present work in conferences.
Impact assessment	• 1 PhD student and 2 MSc students were involved during the project for the toolset implementation. • Potential impact quantification: the 10.56x reduction in training time achieved by EVENFLOW synchronization protocols translates roughly into 10x lower computational costs, energy consumption (kWh), and CO_2 emissions per task, based on a linear model, which could imply faster and more cost-effective federated learning (for the synchronous case). For example, a project budgeted for 100 Federated Learning cycles with the baseline protocol could potentially afford 1,000 cycles using the scalability toolkit for the same computational cost. But this analysis is just an internal exercise, as an effort to quantify the potential impact and value proposition of the Scalability Toolkit.

3.4.3 Verification Toolkit

The exploitation model of ER #7 – Verification Toolkit is presented below.

Table 9: Verification Toolkit Exploitation Model.

Tool/ product description	The Verification Toolkit is a software designed to check if complex AI models used for forecasting events behave reliably and correctly, even when they face tricky or unexpected situations. It tests and verifies whether the neural part of the system (the AI model): ✓ Is robust; it gives stable and reliable predictions, even if inputs change slightly (e.g., small sensor noise). ✓ Does not break under pressure. For instance, if the model wrongly forecasts an event just because a minor change happened in one of the inputs, that is a problem the tool helps detect. In such scenarios, the tool checks if the model is robust or not. Models that are not robust can cause the system forecast to change even with slight changes. ✓ We aim to ensure and provide robustness guarantees of the system, by simulating possible perturbations / noisy inputs within a specified limit and computing the worst-case accuracy. ✓ The verification toolkit provides an end-end verification of Neuro-Symbolic systems, similar to the EVENFLOW use

	cases - might also have use case specific versions (similar to EVENFLOW cases but in general it is domain agnostic). It deals with Complex Event Forecasting (CEF). This means predicting events that are not simple, like detecting when a machine is about to fail based on many sensor readings or forecasting financial crashes from streams of market data. These are events that depend on patterns over time and multiple variables.
Exploitation/ use models	• The main exploitation path for the Verification Toolkit will be based on open-source business models. • The toolkit is provided as an open-source project in GitHub to increase its impact and exploitation potential after the end of EVENFLOW. • Offer consulting services for the toolkit's configuration and custom developments, via service contracts with potential customers. This can be done also via a start-up company linked with ICL which provides consulting and already has existing customers. • The open-source nature of the toolkit eases its way to marketing channels for acquiring new customers who would like configurations and/or new functionalities. • The research exploitation path will be followed; will be used as background IP in other Horizon Europe research projects to develop and test new functionalities and test it under different business cases. • Focus will remain on supporting and expanding the open source community and further enhance the TRL of the toolkit via new research projects. • The toolkit will be available for academic researchers and packaged within the EVENFLOW Framework.
Background IP and Ownership Status	• ICL acts as the sole owner of the result in the frame of the EVENFLOW project. • NCSR is also involved in developing the Neuro-symbolic component.) • The toolkit and its components are protected by Copyrights (i.e., source code, documentation, configuration files). • The toolkit is offered open source via GitHub. (The current repository is - https://github.com/EVENFLOW-project-EU/nesy-veri/.)

Target users/ customers	• The primarily target are research teams and start-ups. But we have identified potential uses by large companies as well. • The toolkit will also be used by the Neuro-symbolic research community to advance the verification of similar systems. • AI engineers who develop and deploy neural network models in systems that require time-based event predictions. • Verification experts in critical domains (e.g., manufacturing (digital twins), energy, healthcare, finance) who need assurance that the AI won't fail unexpectedly. Primarily working on Neurosymbolic systems.
Value proposition analysis	• Via this toolkit, AI models deployed will not change predictions even in the presence of an acceptable noise. AI model will not make bad predictions due to minor noise or changes. • The toolkit can be used to identify adversarial examples from the input dataset. Thus, we can further improve the AI models' resistance to changes in input (due to noise, attacks, sensor faults or environmental changes.) • The system is verifiably reliable, not just assumed to be. Formal guarantees of the models' robustness and a worst-case accuracy of the model on the test dataset can be obtained. • Can be used for images/ videos, could be also extended to other data formats, for example, textual data as a future improvement.
Events and channels for promotion	• Social media channels of EVENFLOW • Promotion via scientific conferences ○ Like NeSy 2025 in Santa Cruz, California • Scientific promotion via papers • We will also target partners' networks and find relevant collaborators to engage to test and contribute to the open-source project – they could be professors from universities and relevant research teams. • Collaboration with other EU projects.

	• PhDs and MSc students will be also engaged to contribute. • University Newsletter • Emails to other departments/ publish info on ICL's website
Open-source project maintenance and governance	• ICL and NCSR will be the admins of the open source project on GitHub (mainly ICL)
Exploitation next steps to increase the TRL	• Background IP in upcoming EU-funded R&I proposals for further validation and extension in new domains – the goal here is to reach TRL 6 (or even 7) via use in large-scale demonstrators and applied research environments. • Targeted effort will be made to grow a developer and user community • Publishing peer-reviewed papers based on its components and performance gains; present work in conferences. • Promoting the tool within the Neuro-Symbolic community and research forums.
Impact assessment	• 1 PostDoc, 1 PhD and 1 MSc student

3.5 Exploitation Models per EVENFLOW Use Case

3.5.1 Forecasting toolkit for oncological precision medicine use case

Based on the exploitation workshop organised, the exploitation model of the precision medicine use case in EVENFLOW is presented below.

Table 10: Exploitation Model of the Precision Medicine Use Case.

Description of the use case and the tools used	The implementation of this pilot is based on two tools: • A pipeline for the creation of pseudo-time cancer trajectories of transcriptomics profiles using a Variational AutoEncoder (VAE). • A NeSy forecasting tool that applies neuro-symbolic Complex Event methods to the synthetic patient trajectories generated by BSC. The combined solution (in case of a joint offering) could integrate BSC's generative AI-based synthetic data generator

	with NCSR's Neuro-symbolic forecasting engine, to create a service capable of producing realistic pseudotemporal patient datasets and generating interpretable event forecasts on top of them. The service then will enable privacy-preserving predictive analytics for healthcare and biomedical research, supporting early diagnosis, risk stratification, and personalized treatment planning with explainable AI. The application of the tools in this use case have reached TRL 4.
Which actor(s) would be the paying customer for a use case-based solution? (customers, stakeholders)	• Companies and start-ups in the healthcare domain who want AI-driven forecasting and data generation tools to enhance clinical decision support systems, patient monitoring, and disease progression modelling. • Pharma companies could benefit from high-quality synthetic patient data that mirror real clinical patterns, enabling them to conduct early drug-response simulations, identify biomarkers, and optimize clinical trial design. • Technology companies engaged in biomedical research could use the synthetic datasets to train and validate AI models for precision medicine. • Hospitals can enhance early diagnosis and personalized care pathways with this data. • Insurance companies can refine risk assessment models, improving healthcare cost management.
End-users needs / problems	• Scarcity of longitudinal datasets (which capture patient evolution over time) prevents the development of predictive and personalized models. The pseudotemporal data generation component reconstructs temporal patient trajectories filling data gaps over time. • Critical lack of realistic and domain-specific synthetic data capable of accurately representing actual patient cases, disease evolution, and population-specific variations. Existing datasets are fragmented and restricted due to privacy and ethical constraints. • Healthcare providers and researchers have a pressing need for early and reliable diagnosis tools to identify high-risk patients and personalize treatment strategies
What value does the solution bring to the user	• Enables early identification of high-risk patients via predictive modelling and personalised screening simulations.

based on the pilot results so far?	• Generates realistic, time-series medical datasets that replicate patient evolution • Validated synthetic datasets that accelerate R&D cycles and improve AI model performance without access to real patient data • Supports proactive healthcare decisions by forecasting disease progression and treatment response ahead of time, recognizing that some cancers progress slowly over years or decades, while aggressive types can advance from stage 1 to stage 4 in months to a few years. • Identification of key molecular and clinical indicators driving disease outcomes • Able to make simulations – informed decision making/ better patient journey, monitoring, follow-up
Early Adopters	• Hospitals already collaborating with BSC (such as IGTP)
Competitors	• Direct competitors in neuro-symbolic, latent temporal trajectory cancer forecasting are extremely limited. • Adjacent competition comes from synthetic temporal data tools, deep learning survival models, and digital twin or simulation platforms. Examples include the Synthetic Data Vault (SDV) library (https://sdv.dev/) and deep survival models that forecast progression and survival via multi-omics (reviewed in https://doi.org/10.1093/bib/bbaf440). • There is a unique niche for combining latent temporal trajectories with symbolic reasoning for personalized clinical predictions.
Timing	• More than 5 years – health tech is a difficult domain • Use the EIC transition and new Horizon Europe projects as a channel
IP Strategy	• BSC holds sole owner of the variational auto-encoder tool. As for the IP strategy, to ensure the most effective technology protection strategy, the BSC protocol has been followed, in which all resources subject to protection are identified and documented using an **IP Disclosure Form** that provides detailed information for an in-depth analysis via a due diligence procedure of the new invention and its legal implications. After this internal registration, evidence

	of the authorship and ownership is done through a block-chain tool. • NCSR is sole owner of the forecasting CE tool.
What commercial channels are realistic for scaling (OEM embedding, DIHs/innovation hubs, systems integrators, cloud marketplace)? What is a potential revenue/ commercialization model?	• There is already a spinoff from BSC (which could use the knowledge gained in EVENFLOW in commercial activities). • AI Factories integration, since BSC could include these data in the factories, supply with synthetic data the AI factories – companies could use infrastructure and data (synthetic data from the BSC tool). • Generate more data – in collaboration with hospitals. (Data generation service as an exploitation model) • Open-source business model – consulting for customization of the open source tools.
Which parts of the pilot output must remain proprietary (e.g. data, models, controller code) and which can be opened or shared (e.g. FAIR datasets, APIs)?	• Open-source component: Application of the VAE tool to Kidney cancer – dynamic scenarios (temporal ones) • Open-source component: Application of the VAE tool to Medulloblastoma – static case (not temporal aspect) • All synthetic dataset generated using the BSC's VAE tool will be released in Zenodo (the Medulloblastoma case is already deposited: https://zenodo.org/records/14516824).
What partnerships are needed to scale?	• Hospitals (physicians), pharmas, tech research companies • AI factories • Conferences and associations for cancer research • EBAN and relevant investor networks
What are the minimal product features for a paid PoC (e.g. for latency, forecast accuracy, explainability outputs)?	• Software as a medical device – class 2 (2a or 2b) • MDR regulation • Define a regulatory roadmap – to get a certification • European regulation • IPR strategy – potential patent application
Events and promo channels (e.g., conferences, internal promo, publications, via	• Conferences and associations for cancer research • GitHub page for promo to the open source community

mailing list, via the website)	
What follow-on investment (internal budget, co-funding, external investors) would you need if pilot KPIs are achieved?	• Use funding from EIC transition • New research projects – Horizon Europe and national funding
Tell me what the exploitation steps foreseen after the end of the project are to sustain the tool and promote its TRL	• EIC transition • Look for private funding opportunities • HE projects – use the result as BG IP • National funding opportunities • New pilots to test and demonstrate the tools – usability, application domains

3.5.2 Prototype neuro-symbolic AGV controller use case

Based on the exploitation workshop organised, the exploitation model of the industry 4.0 use case in EVENFLOW is presented below.

Table 11: Exploitation Model of the Industry 4.0 Use Case.

Description of the use case and the tools used	The use case focuses on improving the performance, safety, and reliability of Autonomous Mobile Robots (AMRs) in smart factory environments. It utilizes the EVENFLOW neuro-symbolic learning and reasoning engine to forecast and prevent incidents in AMR navigation, such as collisions, deadlocks, and blockades. The solution integrates machine learning (ML) and symbolic reasoning to interpret complex sensor data, predict deadlocks, and proactively adapt routes. The result is an AMR deadlock forecasting and path optimizing framework that can be integrated in a "plug-and-produce" manner, without hardware modification on industrially available AMRs. • **Neural-based learning** is used to make robots capable of estimating each other (x,y) coordinates and subsequently, their Euclidean distance, from images of the robots that they obtain via their on-board cameras. A symbolic pattern in the form of a finite state machine (automaton), which captures incidents of the robots entering a deadlock situation in terms of their distance

	evolution through time, has been learnt in an offline fashion from historical data. The robots' neural predictions are fed to the pattern, and the approach uses an off-the-shelf even forecasting tool (Wayeb) to derive forecasts of imminent deadlocks. Our results show that this simple NeSy approach outperforms a purely neural baseline (forecasting deadlocks with a NN) in out-of-distribution settings. • **Multi-mobile Robot Simulation**: The tool is developed in NVIDIA Isaac Sim to create a 3D version of the Smart-Factory Lab in DFKI. It allows to spawn machine modules, assign tasksets and spawn two or more robots to fulfil them. The ROS (Robot Operating System) based navigation stack control is used so that it can use the forecasted information when the model is ready. Simulation provides ability to reproduce and emulate complex situations on demand. • **Data Generation pipeline**: The Simulation is extended with an automated task allocator, initiator as well as data logger to generate trajectory and perception (image) data. Along with this data, additional data on robot task status, deadlock flag and collision flag are also recorded. This forms as the comprehensive synthetically generated dataset for development of forecasting tools.
Which actor(s) would be the paying customer for a use case-based solution? (customers, stakeholders)	• Logistics and warehousing operators relying on AGVs for material handling. • Factories operating with fleets of AGVs and robotic transportation systems. Plants/ Factories who use heterogeneous fleets – fleet management operators. • Technology suppliers who can bundle the EVENFLOW neuro-symbolic optimization software with their sensor packages to offer enhanced perception and fault detection capabilities as part of an integrated product offering. **Potential procurement constraints:** • Solutions must comply with factory safety standards (ISO 3691-4) and be compatible with existing industrial automation frameworks (ROS, OPC-UA).

	• Procurement is typically B2B, requiring proof of ROI (through reduced downtime or maintenance costs) and integration with existing control systems.
End-users needs / problems	• Inefficient route planning leading to bottlenecks and delays. • Safety risks and damage caused by unexpected collisions or object misclassification.
What value does the solution bring to the user based on the pilot results so far?	• Reduction in AGV task completion time, due to predictive incident forecasting • Minimized collision risk via anticipatory route adjustment • Enables more efficient management of heterogenous fleets
Early Adopters	• Smart manufacturing plants using autonomous robots for material transport.
Competitors	**Potential competitors:** • AGV software vendors offering collision avoidance or route optimization modules (e.g., BlueBotics, MiR, OTTO Motors). • Intrinsic from Alphabet.
Timing	More than 3 years
IP Strategy	**Background IP – by NCSR** • Wayeb: European patent EP 3 955 176 B1 - a system and method for forecasting complex events in real-time by defining event patterns (via symbolic regular expressions and symbolic finite/state automata) and applying probabilistic modelling to streams of data to predict if and when such patterns will occur. **Background IP – by DFKI** Robotino: an autonomous transportation vehicle, equipped with software packages for AI-based perception tasks, written in the Robot Operating System (ROS). **Foreground IP:** • NeSy – owned by NCSR, copyright protection, open source.

	• Multi-mobile Robot Simulation – owned by DFKI, copy-right protection • Data generation pipeline: – owned by DFKI, copyright protection
What commercial channels are realistic for scaling (OEM embedding, DIHs/innovation hubs, systems integrators, cloud marketplace)? What is a potential revenue/ commercialization model?	**Sales Channels:** • Licensing to AGV manufacturers and robotics system integrators. • Open-source exploitation of the software **Revenue model:** • License fee per AGV unit or per deployment site. • Optional consulting for AI adaptation and integration.
Which parts of the pilot output must remain proprietary (e.g. data, models, controller code) and which can be opened or shared (e.g. FAIR datasets, APIs)?	• NCSR will have its software components open source
Any commercial activities planned based on this pilot? (e.g., consulting, software product)/ By when, how many years after the project?	N/A in this pilot.
What partnerships are needed to scale?	Create partnerships with industrial stakeholders for pilot deployments that showcase measurable gains in safety, throughput, and coordination efficiency, supporting the progression toward higher TRL levels.
What are the minimal product features for a paid PoC (e.g. for latency, forecast accuracy, explainability outputs)?	Forecast accuracy accepted at industrial levels/ KPIs Compatibility with standard AGV control frameworks (ROS/ROS2).

Events and promo channels (e.g., conferences, internal promo, publications, via mailing list, via the website)	• Hannover Messe, Advanced Factories, IEEE ICAR for showcasing industrial applications. • Publications in IEEE Transactions on Industrial Informatics and AI in Manufacturing journals. • Targeted outreach via DFKI website, LinkedIn, and specialized robotics newsletters.
Tell me what the exploitation steps foreseen after the end of the project are to sustain the tool and promote its TRL	• Mature the prototype into a stable, industry-ready software component by extending its testing across varied intralogistics scenarios and validating performance with real operational data. • Establish integration pathways with commercial robotics controllers, fleet-management systems, and digital-twin environments to demonstrate interoperability and shorten adoption cycles.

3.5.3 Infrastructure lifecycle assessment use case

Based on the exploitation workshop organised, the exploitation model of the infrastructure lifecycle use case in EVENFLOW is presented below.

Table 12: Exploitation Model of the Infrastructure Lifecycle Use Case.

Description of the use case and the tools used	The **Infrastructure Life Cycle Assessment** use case is working on a sensorised potable water pipe section, whereas through the datasets gathered via vibration sensors installed on the pipe is able to identify anomalies and incidents related to the pipe status with high accuracy in terms of location of the anomaly, typically leakages. The prototype has already demonstrated the ability to identify leakage events with high localisation accuracy, based on the datasets generated from EKSO's real-world deployments and refurbished test pipes (e.g., Burgas dataset). To support these capabilities, the use case integrates multiple EVENFLOW software components: • **ML-based leakage detection and forecasting models developed by INTRA**, with regressors and classifiers, process the vibration signals to detect abnormal signatures and estimate the distance from the leakage point. • **ML-based analysis modules developed by NCSR**: NCSR contributes data analysis pipelines and neural models tailored to time-series classification and regression tasks. The work includes exploration of correlations

	across multi-sensor configurations (notably the Burgas dataset) and the application of deep learning approaches for leakage detection and distance estimation, achieving solid predictive accuracy even under noisy conditions. These components are applied on datasets collected from EKSO's deployments, including the 4-tap sensor dataset, forming the basis for the interpretation of sensor correlations, anomaly signatures, and forecasting patterns. The current TRL of the integrated solution is **TRL 4**, validated in a laboratory-controlled but relevant environment.
Which actor(s) would be the paying customer for a use case-based solution and what are their procurement constraints? (customers, stakeholders)	The target market is the pipe infrastructure management and production, targeting: • Smart cities (municipalities) and critical infrastructure operators • Water utility companies • Infrastructure management companies **Short analysis of the underlying national market**: Italy operates one of Europe's most extensive water distribution systems, with over 300,000 km of potable water pipelines and 160,000 km of sewer networks. However, aging infrastructure remains a critical issue: 60% of pipes exceed 30 years, and 25% are over 50 years old, contributing to 42% water losses—equivalent to the annual needs of 43 million people. Historically, investment levels have been low, averaging €600 million/year in the early 2000s, but recent years show a sharp increase. Under the National Recovery and Resilience Plan (PNRR), Italy allocated €4 billion for leakage reduction and modernization, including 25,000 km of new or rehabilitated pipelines by 2026. Household expenditure on water services reached €24.4 billion in 2023, projected to grow to €26 billion by 2028, reflecting rising tariffs and infrastructure costs. The sector is transitioning from capital-intensive upgrades to digital and IoT-based solutions. Smart water metering penetration is currently 17%, with 13.5 million units slated for replacement. Recent tenders worth €250 million and partnerships (e.g., Publiacqua deploying ultrasonic meters) highlight the push toward real-time leak detection, pressure control, and predictive maintenance. These technologies aim

	to reduce non-revenue water, optimize energy use, and ensure compliance with EU directives on wastewater treatment. Italy's water infrastructure faces structural fragility and high leakage rates, but significant investments and smart technologies promise improved efficiency, sustainability, and resilience against climate-induced water stress. EKSO is expecting a relevant benefit from this growing market with special consideration of its intrinsic low impact maintenance technologies and services.
End-users needs / problems	• Infrastructure managers cannot accurately locate and forecast when pipes will require maintenance, replacement, or refurbishment, resulting in either premature replacements or failures. • Water utility operators face unpredictable pipe failures causing emergency repairs, service disruptions, and higher costs compared to timely effective or planned maintenance.
What value does the solution bring to the user based on the pilot results so far? (value proposition)/ Do we have any measurable benefits stemming from this pilot?	• Higher accuracy in the leakage detection with less sensors • Reduced resources loss due to a more precise and timely effective maintenance intervention through anomaly detection. • Based on the EVENFLOW results, we identified that we could achieve the same results by using 5-times less sensors in the pipeline deployment (every 50 meters, compared to the baseline deployment every 10 meters). This could have implications for the maintenance costs of the deployed sensors.
Early Adopters	• Critical pipe infrastructure owner/managers • EKSO use case in EVENFLOW
Competitors/ similar alternative solutions	Several companies including, for example: • HULO (NL) • Ekydra (FR)
Timing	*5 years*
IP Strategy	**Ownership**: INTRA (ML software component, technology integration), NCSR (NeSy expert, software component) and EKSO (domain expertise, infrastructure, data)

	IP Rights: Proprietary license could be offered by INTRA to use its software under commercial agreements.
What commercial channels are realistic for scaling (OEM embedding, DIHs/innovation hubs, systems integrators, cloud marketplace)? What is a potential revenue/ commercialization model?	Channels: Direct marketing towards customer base already serviced both in civil and petrochemical sectors Two ways marketing: • Hardware sale • Software licencing
Which parts of the pilot output must remain proprietary (e.g. data, models, controller code) and which can be opened or shared (e.g. FAIR datasets, APIs)?	• Proprietary: Core forecasting algorithms, trained models, customer-specific configurations • Use case data could be available on Zenodo for re-use.
Any commercial activities planned based on this pilot? (e.g., consulting, software product)/ By when, how many years after the project?	There are already in talks regarding testing pipe sections to be sensorised in the civil sector.
What partnerships are needed to scale?	• Sensor/IoT hardware vendors (manufacturing/ distribution/partnerships) • Water utility companies (Hardware supply, software licence/consultancy)
What are the minimal product features for a paid PoC (e.g. for latency, forecast accuracy, explainability outputs)?	Still to be defined based on customer needs and specs, potentially they would like to have around: • 95%+ forecast accuracy for critical events • <5 second response time for real-time alerts
Events and promo channels (e.g., conferences, internal promo, publications, via	Direct contact of existing customers and vendors, relevant industrial fairs, EVENFLOW website

mailing list, via the website)	
What follow-on investment (internal budget, co-funding, external investors) would you need if pilot KPIs are achieved?	Both internal budget for preliminary marketing, EKSO expects investing a budget of 200.000 Euro over the next 3 years, and co funding technical investors for wider developments to be evaluated in 2026. Discussions with big European manufacturers are already taking place.
Tell me what the exploitation steps foreseen after the end of the project are to sustain the tool and promote its TRL (at least 3 actions)	• Advance to TRL 6-7 with industry pilots and regulatory compliance • Gather dataset with real leakages to be analysed, and build a NeSy solution based on them • Extend the solution for predictive maintenance use case • Use as Background IP in other projects (INTRA's ML component)

3.6 Individual exploitation paths per partner

The individual exploitation plans of all EVENFLOW partners were collected and analysed during the current reporting period. Each organisation has outlined how it intends to re-use the project's technological results, methods, and datasets, reflecting its specific role and operational context. A template Word-based template was used to gather all inputs, and a consolidated table with the detailed exploitation intentions of each partner is provided in this section.

Table 13: Individual Exploitation Plans per Partner in EVENFLOW.

Partner name	Partner exploitation path (exploitation intention of the organization based on the knowledge and results gained/ achieved in EVENFLOW)
INTRA	INTRA will exploit the EVENFLOW results by offering IT consulting services (via service contracts) focused on the integration and deployment of neuro-symbolic solutions, capitalising on the expertise gained via the platform integration activities. Also, INTRA will continue investigating the utilisation of the cutting-edge EVENFLOW results to improve our commercial offerings to our customers. In parallel, INTRA will pursue the further exploitation of the ML component for the water pipe solution and analyse its impact on water pipe cost maintenance and deployment.
NCSR	NCSR will exploit the reasoning-based forecast interpreter by positioning it as an open-source project for neuro-symbolic technology. The tool will be maintained and further matured by NCSR as background IP, leveraging its combination of symbolic automata and deep learning to deliver explainable forecasting across time-series, image/video and complex

Partner name	Partner exploitation path (exploitation intention of the organization based on the knowledge and results gained/ achieved in EVENFLOW)
	event data. NCSR will sustain the tool's evolution via (i) active participation in Horizon Europe proposals, where the interpreter will be reused and extended as background IP for validation in new domains, (ii) community-building around the open-source repository to attract external contributors and academic users, and (iii) targeted consulting and custom development services for organisations who require similar configuration or new features on neurosymbolic applications.
ARC	ARC will use the Scalability Toolkit as an open-source asset to strengthen its positioning in research and attract collaboration opportunities. ARC intends to capitalize on the toolkit as background IP in upcoming EU-funded R&I proposals. This approach will enable ARC to advance the tool's TRL beyond its current level of 5, aiming to reach TRL 6 or 7 by integrating the toolkit in large-scale demonstrators across domains such as federated learning, mobile fraud detection, edge AI, and industrial AI use cases. Beyond research reuse, ARC will support the open-source community around the toolkit via engaging academic collaborators, MSc and PhD students, and technical contributors to maintain and extend the toolkit's components (e.g., SDEaaS, Subito, training optimization). ARC also plans to have a roadmap of peer-reviewed publications and conference presentations to increase the toolkit's visibility, demonstrate its technical advantages (e.g., 10x training speedup, 5–10x time-to-deploy gains), and reinforce its market relevance. Finally, ARC envisions to provide consulting and configuration services as a secondary exploitation path, via service contracts, using the toolkit's open-source adoption as a lead generator for tailored deployments, especially for teams in banking, telecom, and AI-intensive sectors seeking lightweight and scalable ML solutions.
BSC	BSC aims to maximize the potential commercial and strategic value of EVENFLOW outcomes by exploring a structured exploitation strategy. Possible pathways include licensing, spin-offs, trade secret protection, and consulting services; however, each option will require careful evaluation to assess feasibility, market potential, and IP considerations before implementation.
DFKI	The complex-event forecaster will be carried forward as a predictive intelligence layer for industrial intralogistics systems, providing early warnings of collisions, deadlocks, and other disruptive events that can be directly consumed by mapping, coordination, and motion-planning components. By offering a standardized interface and modular deployment model, the tool can be integrated into existing robotics frameworks with minimal disruption, allowing operators to enhance fleet responsiveness and overall system performance. Scaling through pilot deployments, partner integrations, and targeted commercial demonstrations will support continued refinement and adoption, enabling the technology to mature into a high-TRL component of advanced automated intralogistics solutions.

Partner name	Partner exploitation path (exploitation intention of the organization based on the knowledge and results gained/ achieved in EVENFLOW)
EKSO	EKSO will use the knowledge gained in EVENFLOW into its engineering and maintenance services. EVENFLOW helped EKSO with practical expertise in configuring sensorised pipe sections, collecting and labelling vibration data, and interpreting ML and neuro-symbolic diagnostic outputs. This helps EKSO to provide more accurate condition-based maintenance advice to water utilities and infrastructure operators. Although EKSO will not commercialise the EVENFLOW platform, it will leverage the acquired know-how on neuro-symbolic forecasting to strengthen the reliability of its predictive maintenance workflows.
ICL	ICL aims to advance and promote the EVENFLOW verification toolkit, a suite of tools designed to support the formal verification of complex neurosymbolic systems. The tools developed within the EVENFLOW project are actively incorporated into academic publications, conference contributions, and collaborative research outputs. By showcasing the toolkit via peer-reviewed venues, ICL helps to establish EVENFLOW as a reference point for researchers investigating verification methods across machine learning, logic-based reasoning, and integrated neurosymbolic approaches. Going forward, ICL intends to place greater emphasis on building an active and participatory community around the EVENFLOW framework. A key objective is to attract contributions from external researchers, practitioners, and industry partners who are working on verification challenges within neurosymbolic AI. In particular, ICL aims to expand the toolkit to address specialised verification requirements within neurosymbolic architectures, like reasoning-aware robustness analysis, hybrid constraint propagation techniques, and interpretable verification pipelines that unify symbolic structure with neural uncertainty. These efforts will support the long-term objective of establishing EVENFLOW not only as a research prototype but also as a sustainable, community-driven ecosystem for neurosymbolic verification.

3.7 Sustainability of the EVENFLOW Toolsets and way forward

EVENFLOW's sustainability is mostly achieved via the development and maintenance of open-source projects, and it is primarily research-driven. The core toolsets (Reasoning-Based Forecast Interpreter, Scalability Toolkit, Verification Toolkit and several use-case pipelines) are released as open-source projects that partners will maintain and extend as background IP in future research collaborations. Building on these assets, the consortium will pursue further maturation and validation via Horizon Europe and other funding schemes, while via potential spin-offs the consortium could explore service-based and productization opportunities in domains such as precision oncology, industrial intralogistics (industry 4.0) and water infrastructure.

Table 14: EVENFLOW open-source projects.

#	EVENFLOW Open-Source Project Name	Admin	License	Link	Potential target groups
1	Scalability Toolkit	ARC		Here	AI/ML engineers seeking scalable training; research groups on distributed ML; MLOps teams; cloud/edge AI developers; organisations optimising training cost and energy efficiency; companies exploring federated learning.
1.1	Synopses Data Engine as a Service	ARC	Apache 2.0	Here	
1.2	Synopses-based training	ARC	AGPL v3	Here	
1.3	Advanced Distributed-Parallel Training	ARC	AGPL v3	Here	
1.4	SuBiTO	ARC	AGPL v3	Here	
1.5	NeuroFlinkCEP	ARC	AGPL v3	Here	Event processing engineers; time-series analytics developers; researchers building predictive maintenance or anomaly detection pipelines.
2	NeSyA / DeepFA (Reasoning-Based Forecast Interpreter)	NCSR	GPL v3	Here	Researchers in neuro-symbolic AI; organisations working on explainable forecasting and trajectory prediction.
3	Verification Toolkit (NeSy-Veri)	ICL	BSD 3-Clause License (dependency with NeSy - might use GPL3)	Here	AI verification and safety teams; organisations in regulated sectors (healthcare, finance, mobility); SMEs providing compliance, certification or AI governance services.
4	Synthetic Medulloblastoma Data Generator & Pipeline	BSC	Apache 2.0	Here	Hospitals and cancer research centres; biomedical AI teams; pharma companies.

To ensure sustained adoption, the consortium's research partners will continue to use these repositories as reference implementations and as components in future pilots. Their open-source nature significantly increases the likelihood of community contributions, cross-project

reuse, and uptake by related EU initiatives such as AI-on-Demand, Trustworthy AI clusters, and even national AI research programmes. While the project is primarily research-driven, it also opens several promising commercialisation channels. For instance, spin-offs related to consortium partners could leverage the insights gained in EVENFLOW to enhance their existing service portfolios in neurosymbolic forecasting, scalable ML deployment, verification and MLOps. At the same time, selected technical components are transitioning into industrial ecosystems. Notably, specific features of the Scalability Toolkit have been integrated into the Altair AI Studio suite and made available as a free value-adding extension, increasing exposure to an industrial user base.

Another exploitation strategy for most partners centres on using EVENFLOW tools as background IP in future Horizon Europe, Digital Europe, and national research initiatives. Several exploitable results currently at TRL 4 - 5 have clear trajectories towards TRL 6 - 7 under new collaborative pilots. Moreover, selected technical results (particularly synthetic data generation, scalable training methods, verification engines, and neuro-symbolic forecasting) could also be candidates for EIC Transition calls.

4 Conclusions

The dissemination, communication, and exploitation activities carried out under EVENFLOW have supported the project's objectives and helped establish a clear identity. Through coordinated actions, the consortium achieved visibility across scientific and stakeholder communities, while laying the groundwork for the project's results to be taken forward.

The project effectively used multiple communication channels, producing high-quality content and achieving engagement on the website and social media. The campaigns and participation in leading conferences significantly enhanced project recognition and strengthened collaborations across Europe's AI ecosystem. Through coordination of the TrustWorthy AI Cluster and its involvement in the AI-on-Demand platform, EVENFLOW contributed to the creation of a coherent European narrative on trustworthy AI.

Concerning the exploitation activities and the future sustainability of the project results, a structured process for identifying and monitoring exploitable results enabled the consortium to establish a final, validated list of assets, covering the integrated platform, three core toolsets (DeepFA, Scalability Toolkit, Verification Toolkit) and several use-case modules. Most components are released as open source to support their re-use. Via six targeted exploitation workshops and continuous engagement with technical and pilot partners, detailed exploitation models were developed for each toolset and for the three use cases, defining user groups, value propositions, licensing approaches and potential routes to market. In parallel, all partners prepared individual exploitation plans confirming that future exploitation will rely mainly on further research projects, maintenance and evolution of the open-source repositories, and consulting or custom development activities.

Overall, the dissemination, communication, and exploitation efforts ensured that EVENFLOW's work reached the appropriate audiences and remained visible throughout the project's duration. These activities now provide a solid basis for sustaining interest in the project's outcomes and for supporting their further development and uptake beyond the project's formal end.