



ALFIE

ASSESSMENT OF LEARNING TECHNOLOGIES AND FRAMEWORKS
FOR INTELLIGENT AND ETHICAL AI

D5.4 Initial market analysis, exploitation plan and business models

Deliverable No:	D5.4
Deliverable Name:	Initial market analysis, exploitation plan and business models
Work Package:	WP5
Task:	T5.4
Dissemination Level:	PU
Deliverable Type:	R
Lead Organization:	CTL
Submission month:	M18
Date:	31 March 2025



Funded by
the European Union

Funded by the European Union under Grant Agreement 101177912. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.



Project description

Acronym	ALFIE
Title	Assessment of Learning technologies and Frameworks for Intelligent and Ethical AI
Coordinator	CATALINK
Reference	101177912
Type	Research and Innovation Actions (RIA)
Programme	Horizon Europe (HORIZON)
Topic	HORIZON-CL2-2024-TRANSFORMATIONS-01-06 Beyond the horizon: A human-friendly deployment of artificial intelligence and related technologies
Start	01.10.2024
Duration	36 months
Website	alfie-project.eu
Consortium	Catalink Limited (CTL), Cyprus, Coordinator University of Brighton (UoB), United Kingdom Centre for Research & Technology (CERTH), Greece Distributed Analytics Solutions LTD (DAS), United Kingdom Edge Hill University (EHU), United Kingdom Kempelen Institute of Intelligent Technologies (KINIT), Slovakia Universitat Autònoma de Barcelona (UAB), Spain Robert Bosch Espana SLU (Bosch), Spain Eindhoven University of Technology (TU/e), Netherlands Diadikasia Business Consulting (DBC), Greece

Document history

Version	Date	Beneficiary	Description
0.1	15/10/2025	CTL	Created the document and the table of contents
0.2	27/11/2025	CTL	Gathered questionnaires from a few partners and start drafting Key Exploitable Results (KEA)
0.3	02/02/2026	CTL	Gathered more questionnaires from partners and continue drafting the deliverable
0.4	13/02/2026	CTL	Finalized first version of the deliverable and ask feedback from KEA owners
0.5	09/03/2026	CTL	Updated deliverable based on KEAs owners' comments
0.6	10/03/2026	CTL	Final version of the deliverable is provided to the reviewers
0.7	11/03/2026	DBC	Reviewed the document and proposed amendments
0.8	24/03/2026	Bosch	Reviewed the document and proposed amendments
1.0	27/03/2026	CTL	Updated deliverable based on reviewers' comments and final version.

Authors

Partner	Name(s)
CTL	Konstantinos Avgerinakis, Maria-Eleni Skarkala
DAS	Stelios Kapetanakis, Charlie Gadd
KInIT	Martin Tamajka
CERTH	Stamatis Samaras, Sotiris Diplaris, Alexandros Vassiliades
EHU	Yannis Korkontzelos

Contributors

Partner	Contribution type	Name
DBC	Review	Stella Theologidou
Bosch	Review	Leonardo Abad Martin

Glossary

Acronym	Definition
AI	Artificial Intelligence
AIA	EU Artificial Intelligence Act
BFSI	Financial services and insurance
BMC	Business Model Canvas
EDIHs	European Digital Innovation Hubs
EEA	External Ethics Advisor
ESG	Environmental, Social, and Governance
ETD-Hub	EthiTech Dialogue Hub
GenAI	Generative AI
HITL	Human-in-the-Loop
HLEG-AI	High-Level Expert Group on Artificial Intelligence
ICT	Information and Communication Technology
KEA	Key Exploitable Asset
LLM	Large Language Model
NGOs	Non-governmental organisations
PaaS	Platform-as-a-Service
RAG	Retrieval-Augmented Generation
SemKG	Semantic Knowledge Graph
SME	Small and medium-sized enterprise
XAI	Explainable AI

Executive Summary

The long-term success of ALFIE is fundamentally rooted in its capacity to generate impact that resonates far beyond the boundaries of the project consortium. This enduring impact is secured by maximizing diverse exploitation opportunities and ensuring that all Key Exploitable Assets (KEAs) are designed for seamless scalability across various sectors. Central to this mission is the cultivation of broad acceptability among a wide spectrum of stakeholders, including policymakers, civil society, academia, and industry leaders, ensuring that the project's methodologies remain relevant long after its conclusion.

The exploitation strategy is built upon a foundation of rigorous market analysis and innovative business modeling. This strategic framework is designed to fortify the European Union's position within the global ethical AI market, driving European industries and research institutions to the global forefront by transforming ethical considerations into a distinct competitive advantage. Central to this effort is the identification of the project outcomes, ensuring they address complex technological and ethical requirements while proactively identifying critical market gaps and stakeholder needs.

To gain insights into the individual KEAs, the project gathers primary data from consortium partners through specialized questionnaires, to capture essential details regarding the competitive landscape, and facilitate transparent information sharing across the consortium. A comprehensive market analysis of the ethical AI domain is conducted, and SWOT and PESTEL analytical frameworks evaluate the strengths, weaknesses, opportunities, and threats of the ALFIE KEAs against political, economic, social, technological, environmental, and legal factors. The exploitation strategy proposes actionable plans, with an emphasis on stakeholder engagement, establishing efficient knowledge transfer mechanisms, and formulating strategies to protect and optimize the commercial potential of the project.

By identifying target stakeholder needs early on, ALFIE ensures widespread awareness. Furthermore, the strategy includes financial roadmaps. These financial structures are vital for bridging the gap between research and commercialisation, ultimately positioning ALFIE as a frontrunner in the ethical AI domain and laying the groundwork for sustainable, market-leading innovation.

Contents

Executive Summary	5
1 Introduction	10
1.1 Overview.....	10
1.2 Relation to other tasks and deliverables.....	11
1.3 Scope of deliverable D5.4	11
1.4 Structure of the deliverable.....	12
2 Key Exploitable Assets (KEA) of ALFIE	13
2.1 Questionnaire	13
2.2 Description of KEAs and its value proposition.....	15
2.2.1 KEA1. ETD-Hub (Owner: DAS).....	16
2.2.1.1 Description.....	16
2.2.1.2 Value Proposition.....	16
2.2.1.3 Potential research and market exploitation.....	17
2.2.1.4 Potential customers and engagement activities	18
2.2.2 KEA2. AutoML platform (Owner: CTL, Co-owner: All).....	21
2.2.2.1 Description.....	21
2.2.2.2 Value Proposition.....	21
2.2.2.3 Potential research and market exploitation.....	21
2.2.2.4 Potential customers and engagement activities	22
2.2.3 KEA3. Data warehouse (AutoDW) of ethical data (Owner: CERTH).....	23
2.2.3.1 Description.....	23
2.2.3.2 Value Proposition.....	23
2.2.3.3 Potential research and market exploitation.....	23
2.2.3.4 Potential customers and engagement activities	23
2.2.4 KEA4. Semantic knowledge graph (SemKG) of ethical data (Owner: EHU, Co-owner : CTL)	25
2.2.4.1 Description.....	25
2.2.4.2 Value Proposition.....	25
2.2.4.3 Potential research and market exploitation.....	25
2.2.4.4 Potential customers and engagement activities	26
2.2.5 KEA5. Agentic Core (Owner: KINIT)	27
2.2.5.1 Description.....	27
2.2.5.2 Value Proposition.....	27
2.2.5.3 Potential research and market exploitation.....	28

2.2.5.4	Potential customers and engagement activities	28
2.2.6	KEA6. Web interface with XAI features (Owner: CTL, Co-owner: CERTH, KInIT) 29	
2.2.6.1	Description	29
2.2.6.2	Value Proposition.....	29
2.2.6.3	Potential research and market exploitation.....	30
2.2.6.4	Potential customers and engagement activities	30
3	Market analysis	33
3.1	Overview.....	33
3.2	Market landscape	34
3.3	KEAs market analysis.....	38
3.4	SWOT analysis.....	45
3.5	PESTEL analysis.....	53
4	Exploitation Plan	56
4.1	Business Model Canvas	57
4.2	Target stakeholders and engagement.....	58
4.3	Exploitation impact indicators	58
4.4	Exploitation strategy for ALFIE KEAs	59
4.4.1	KEA1. ETD-Hub (Owner: DAS).....	59
4.4.1.1	Market Entry and Expansion.....	59
4.4.1.2	Product Differentiation	60
4.4.1.3	Revenue Generation.....	61
4.4.1.4	Partnership Development	62
4.4.2	KEA2. AutoML platform (Owner: CTL, Co-owners: All).....	62
4.4.2.1	Market Entry and Expansion.....	62
4.4.2.2	Product Differentiation	63
4.4.2.3	Revenue Generation.....	63
4.4.2.4	Partnership Development	64
4.4.3	KEA3. Data warehouse of ethical data (Owner: CERTH)	64
4.4.3.1	Market Entry and Expansion.....	64
4.4.3.2	Product Differentiation	64
4.4.3.3	Revenue Generation.....	65
4.4.3.4	Partnership Development	65
4.4.4	KEA4. Semantic knowledge graph (SemKG) of ethical data (Owner: EHU, Co-owner: CTL)	66



4.4.4.1	Market Entry and Expansion.....	66
4.4.4.2	Product Differentiation	66
4.4.4.3	Revenue Generation.....	66
4.4.4.4	Partnership Development	67
4.4.5	KEA5. Agentic Core (Owner: KINIT)	67
4.4.5.1	Market Entry and Expansion.....	67
4.4.5.2	Product Differentiation	67
4.4.5.3	Revenue Generation.....	67
4.4.5.4	Partnership Development	68
4.4.6	KEA6. Web interface with XAI features (Owner: CTL, Co-owners: CERTH, KInIT) 68	
4.4.6.1	Market Entry and Expansion.....	68
4.4.6.2	Product Differentiation	68
4.4.6.3	Revenue Generation.....	69
4.4.6.4	Partnership Development	70
5	Impact of ALFIE	71
5.1	Social, economic, and technological impacts of ALFIE across various domains ...	71
5.2	Standards that influence ALFIE's solutions	72
6	Conclusions and Future plans.....	74
	Appendices.....	75
	Appendix A: PESTEL analysis for ETD-Hub	75
	Appendix B: PESTEL analysis for AutoML platform.....	79
	Appendix C: PESTEL analysis for Data warehouse of ethical data	80
	Appendix D: PESTEL analysis for Semantic knowledge graph of ethical data	81
	Appendix E: PESTEL analysis for Agentic Core.....	82
	Appendix F: PESTEL analysis for web interface with XAI features	83



Figures

Figure 1. AI Market Size estimates in the period between 2024 and 2034	34
Figure 2. Global Responsible AI market size estimate between 2024-2034.....	36
Figure 3. AI governance market size (2025 - 2034)	37

Tables

Table 1. D5.4 Input from other tasks and deliverables.....	11
Table 2. D5.4 Output for other tasks and deliverables	11
Table 3. Exploitation of Technical Components questionnaire	14
Table 4. Exploitation of Scientific Components questionnaire	15
Table 5. ETD-Hub potential customers and activities	19
Table 6. AutoML platform potential customers and activities.....	22
Table 7. AutoDW potential customers and activities	24
Table 8. SemKG potential customers and activities.....	26
Table 9. Agentic Core potential customers and activities.....	29
Table 10. Web interface with XAI features potential customers and activities.....	31
Table 11. SWOT analysis for ETD-Hub	46
Table 12. SWOT analysis for AutoML platform.....	48
Table 13. SWOT analysis for Data warehouse (AutoDW) of ethical data	49
Table 14. SWOT analysis for Semantic knowledge graph of ethical data	49
Table 15. SWOT analysis for Agentic Core.....	51
Table 16. SWOT analysis for web interface with XAI features.....	51

1 Introduction

1.1 Overview

The ALFIE project is anchored by two primary objectives designed to harmonize human-centric values with advanced technological development. The first objective focuses on the creation of a collaborative forum that invites active participation from a diverse spectrum of specialists. Through the EthiTech-Hub (ETD-Hub), the project establishes an environment where citizens, ethicists, legals, policymakers, and AI professionals, can directly contribute to the evolution of the European Union's legal and ethical frameworks regarding Artificial Intelligence (AI). This space ensures that regulatory standards are informed by a broad consensus and reflect the experiences and concerns of the public. The second core objective is the delivery of the AutoML platform, a platform engineered to generate AI models and tools that are inherently compliant with the EU legal framework. This platform is designed mitigating algorithmic bias and ensuring that technical performance does not come at the expense of fairness or legal integrity.

To realize this ambitious vision, ALFIE extends beyond these two pillars to develop an integrated suite of high-value technological products. These Key Exploitable Assets (KEAs) include a Data Warehouse (AutoDW), which serves as a curated repository for high-quality information, and Semantic Knowledge Graph that maps the complex relationships between ethical principles and technical requirements. Additionally, ALFIE is developing advanced multilingual AI tools (Agentic Core) to ensure accessibility across different linguistic regions and a sophisticated user interface equipped with Explainable AI (XAI) features, which provide transparency by making the logic behind AI decisions understandable to human users. Because these products possess significant utility both as an integrated ecosystem and as stand-alone solutions, they form the cornerstone of the project's broader exploitation and marketisation strategy.

This deliverable provides a market analysis and a series of business plans tailored to each of these KEAs. To ensure these plans are grounded in commercial reality, the project engaged in a rigorous collaborative process with consortium partners, utilizing detailed questionnaires to harvest their marketisation and exploitation prospects. This data collection effort captured the unique value propositions of each asset, identified potential customers, and provided an assessment of competitive advantages and potential market barriers. By evaluating how these tools compare to existing solutions, the project has formulated specific strategies to maximize both research impact and commercial revenue. Ultimately, by detailing these exploitation paths, this deliverable expands upon the initial commercialisation roadmap, providing a path toward sustainable market entry and long-term viability for ALFIE's innovations.

1.2 Relation to other tasks and deliverables

This deliverable is related to the following other ALFIE tasks and deliverables:

Receives inputs from:

Table 1. D5.4 Input from other tasks and deliverables

Deliverable	Due Date	Input for D5.4
D5.1	31 Jan 2025	Project website, market analysis, dissemination and exploitation plan

Provides outputs to:

Table 2. D5.4 Output for other tasks and deliverables

Deliverable	Due Date	Output from D5.4
D5.5	30 Sep 2027	Final market analysis, exploitation plan and business models

1.3 Scope of deliverable D5.4

The scope of this deliverable encompasses an evaluation of the ethical AI market, a rigorous assessment of individual project assets, and a long-term strategy for ensuring the ALFIE's sustainable impact across diverse sectors.

Detailed profiles for each of the project's Key Exploitable Assets (KEAs) are provided, featuring comprehensive SWOT analyses that identify internal technical strengths and weaknesses alongside external market opportunities and competitive threats. To bridge the gap between innovation and adoption, the deliverable identifies specific customer segments and outlines tailored engagement modalities.

The deliverable provides an in-depth analysis of the ethical AI landscape, a market projected to grow significantly as organisations transition from theoretical principles to the practical operationalisation of AI governance. This section explores the rising demand for tools that facilitate algorithmic bias detection, transparency, and regulatory automation, positioning ALFIE's innovations within a sector driven by the global expansion of generative AI and increasingly stringent compliance mandates like the EU AI Act¹. This is complemented by an exploitation plan that details the progression of ALFIE's technologies from research prototypes to market-ready solutions.

Finally, the deliverable quantifies the projected multi-dimensional impacts of the project. This includes social impacts such as fostering public trust through democratic participation in AI law-making, economic impacts derived from the creation of new high-value jobs and reduced compliance costs for EU industries, and technological impacts through the advancement of state-of-the-art AutoML and explainable AI features. By highlighting the international standards that influence ALFIE's design, the deliverable ensures that the project's outputs are not only

¹ <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689>



innovative but also globally interoperable and future-proofed against an evolving regulatory environment.

1.4 Structure of the deliverable

Section 2 describes ALFIE's Key Exploitable Assets (KEAs) in detail, highlighting the functional features and technical innovations, defining the value propositions of these assets. Section 3 provides a market analysis and employs two strategic frameworks to evaluate the environment in which ALFIE operates (PESTEL and SWOT analysis). Section 4 presents the exploitation plan, serving as the strategic blueprint for moving ALFIE from a research context into commercial and practical application. Section 5 focuses on the broader influence of the project, highlighting its multi-dimensional impact across various. Finally, Section 6 concludes the deliverable.

2 Key Exploitable Assets (KEA) of ALFIE

The ALFIE project is anchored by two core components: the EthiTech Dialogue Hub (ETD-Hub) and the AutoML platform. These serve as the foundational pillars of the project, embodying an innovative and integrated solution designed to significantly advance the field of ethical AI.

The ETD-Hub and the AutoML platform encompass a suite of tools and services, referred to as Key Exploitable Assets (KEAs), which are developed to address the ethical and legal implications of AI. These tools include a Data Warehouse (AutoDW) for secure data storage and management, a Semantic Knowledge Graph (SemKG) for enhanced knowledge representation and reasoning, the Agentic Core providing natural language processing (NLP) with multilingual support and a Web interface with explainable AI (XAI) features, empowering users to understand and trust AI decision-making processes, a critical aspect for ethical compliance and user acceptance.

To maximize the exploitation potential of ALFIE's KEAs, a questionnaire was developed and circulated among KEA owners across the consortium (Section 2.1). This questionnaire was designed to gather vital information regarding each KEA's technical specifications, market readiness and potential applications. The collected data is instrumental in identifying target market segments, refining the exploitation strategy, and ensuring that each KEA is utilized to its fullest potential. This systematic approach aims to enhance the commercial viability of ALFIE's assets, foster broader adoption, and ultimately maximize the project's impact within the ethical AI landscape. Following this overview, the questionnaire results and an initial analysis of the KEAs are presented, providing valuable insights into their strengths, opportunities, and pathways for successful market deployment.

2.1 Questionnaire

Several Key Exploitable Assets (KEAs) were identified and documented. The process involved the collaboration with consortium partners to gather detailed insights and understand each KEA's unique contributions, such as information on the market, SWOT analyses, and examined external factors using PESTEL. This includes recognizing market segments, current competitors, and strategies to overcome them. For each KEA, detailed data was collected including the description of the technology, existing and future intellectual property rights (IPRs), value propositions, target audiences, and competitive positioning. This process ensures each KEA's innovation and potential impact is highlighted while creating a plan for how they can be used and introduced to the market. The following sections provide detailed profiles of these KEAs, showing how they can support the project's goals and benefit the wider industry.

The questionnaire of the technical components is presented in Table 3, and for the scientific components is shown in Table 4.

Table 3. Exploitation of Technical Components questionnaire

Technology Name (KEA)	<i>[Name of the key exploitable asset.]</i>
Description of Technology	<i>[Description of the core technology, including its functionality, features, and technical specifications]</i>
IPR existing (background)	<i>[Explanation of existing intellectual property rights related to the technology or background IP.]</i>
IPR to be developed (foreground)	<i>Tentative (may change in the future):</i> <i>[Description of planned intellectual property rights to be developed during the project.]</i>
IPR to be commonly exploited by the consortium	<i>Tentative (may change in the future):</i> <i>[Explanation of intellectual property rights intended for common exploitation by the consortium partners.]</i>
Value Proposition	<i>[Detailed statement highlighting the value proposition of the technology and its benefits to stakeholders.]</i>
Potential Audiences you refer	<i>[Identification of potential target audiences or end-users for the technology.]</i>
Potential application in other market segments further than the ones specified in the project	<i>[Exploration of potential applications of the technology beyond the project's specified market segments.]</i>
Existing Competition (for market analysis section)	<i>[Overview of existing competitors in the market relevant to the technology.]</i>
Advantages against Competition	<i>[Explanation of the advantages the technology possesses over existing competition.]</i>
Disadvantages against Competition	<i>[Identification of potential disadvantages or limitations compared to existing competition.]</i>
How you plan to outmaneuver the competition	<i>[Strategies planned to overcome competition and gain market advantage.]</i>
How you plan to support entrance to market	<i>Tentative (may change in the future):</i>

Table 4. Exploitation of Scientific Components questionnaire

Project Offering (if technological, please fill the technology partners table)	<i>We are a university medical center, so comply with Academic but our offering is Technological (see above).</i> <i>[Description of the academic project offering, including its objectives and scope.]</i>
Description of Offering	<i>[Elaboration on the details and components of the academic project offering.]</i>
IPR existing	<i>[Explanation of any existing intellectual property rights associated with the academic project.]</i>
IPR to be developed	<i>[Description of planned intellectual property rights development within the academic project]</i>
IPR to be commonly exploited by the consortium	<i>[Explanation of intellectual property rights intended for common exploitation by project consortium members.]</i>
Value Proposition	<i>[Statement highlighting the value proposition of the academic project and its contributions]</i>
Potential Audiences you refer	<i>[Identification of potential audiences or beneficiaries of the academic project.]</i>
Plans to reach your audiences	<i>[Strategies and plans to effectively reach and engage with the identified audiences.]</i>

2.2 Description of KEAs and its value proposition

The KEAs are core parts of the ALFIE project that have the potential to be exploited beyond the project's initial goals. These KEAs demonstrate innovative methods and technologies developed through extensive research and development efforts. In this initial stage of the ALFIE project, we focus on highlighting the key KEAs, providing details about their functionalities, plans for research and market development, and their main value propositions. These KEAs are vital because they increase the project's impact by creating new market opportunities and helping set industry standards, ensuring that the research results can transition from the lab to real-world applications. To do this, a detailed market analysis was conducted to identify where the KEAs can be applied, who the target audiences are, and how to exploit them commercially. This approach aims to foster innovation, promote economic growth, and lay the groundwork for future advancements. The following sections present the main KEAs of ALFIE, emphasizing their key innovations and potential for commercialisation and market use. These KEAs not only showcase the project's innovative strengths but also highlight the value they can bring to stakeholders and the broader market.

2.2.1 KEA1. ETD-Hub (Owner: DAS)

2.2.1.1 Description

The ETD-Hub is a multilingual digital discussion and deliberation platform designed to capture, structure, and provide data enabling the analysis of public and expert opinion on artificial intelligence (AI) and AI ethics. It uses a moderated discussion forum format to collect analytics that can be explored to surface trends, areas of consensus and disagreement, and emerging concerns.

The platform offers:

- **Structured discussion spaces** (e.g. by topic, policy file, use case, stakeholder group) with configurable moderation rules.
- **Multi-lingual participation** via an extensible language layer and support for EU official languages.
- **Role-based access** for administrators, moderators, and general users; with granular permissions.
- **A novel dataset** that aggregates contributions and can offer unprecedented insight (e.g. can be analysed to derive sentiment, argument clusters, stakeholder perspectives, evolution over time).
- **Private data export interfaces** (APIs) for integration with research.

Technically speaking, ETD-Hub is a modular, web-based platform with a microservices-based architecture, featuring a scalable security-first backend, secure authentication/authorisation (Keycloak), and a responsive front-end optimised for both desktop and mobile use.

2.2.1.2 Value Proposition

The ETD-Hub provides a trusted, evidence-oriented environment to understand how different stakeholders think and feel about AI and AI ethics. It delivers value on several levels:

- **For policymakers and regulators:** the ability to request and obtain (via open discussion) timely, structured insights into public and expert views on AI policies and regulatory initiatives, supporting impact assessments, public consultations, and legitimacy of decisions.
- **For researchers and ethicists:** rich, analysable datasets to study norms, concerns, and acceptance of AI systems across countries and stakeholder groups.
- **For industry and innovators:** an early-warning and co-creation space to obtain feedback on test AI use cases, understand societal and regulatory expectations, and co-design ethical safeguards with affected communities.
- **For civil society and citizens:** an inclusive forum to have a voice in AI governance with transparent rules, feedback mechanisms, and connections to decision-making processes.
- **For AI engineers:** ETD-Hub data is connected to AutoML via AutoDW and SemKG. This means insights from public and expert deliberation are surfaced by SemKG and translated into concrete design constraints and evaluation criteria for developers in AutoML. This helps support model selection, hyperparameter search and deployment

decision approaches that are grounded in the expectations and concerns expressed by affected stakeholders.

- **For non-specialists building AI:** SemKG insights based on ETD-hub data can advise users on creating more ethical AI and apply constraints to the AutoML model creation process. Additionally, users can raise new questions on ETD-Hub via AutoML to ensure their model aligns with the most up to date policy and public sentiment.

The core value proposition for ETD-Hub is capturing debate about AI and through SemKG transforming this into structured, actionable knowledge; reducing the gap between technical development, ethical reflection, and policy decisions.

2.2.1.3 Potential research and market exploitation

As an EU-Horizon funded technology, developed in a consortium of experts, there are multiple clear pathways through which ETD-Hub can gain market advantage:

- **Leverage the consortium's credibility:** use the academic, policy, and civil society reputation of agreeable partners to build trust, secure early adopters, and obtain pilot deployments with high-visibility institutions.
- **Demonstrate impact through pilots:** we are in the process of running high-quality pilots to obtain data that can support clear stories of policy influence, workflow impact, or organisational change as it relates to specific partner test use cases for Bosch, CTL, and UAB. We will document these as case studies to showcase added value.
- **User-centred design and support:** through collaboration with UAB we are investing in accessibility, usability, training materials, and facilitation support to reduce complexity barriers and make the platform accessible to disabled and non-technical users.
- **Multilingual support built-in:** competitors operate their solutions in English. ETD-Hub will have partner language support validated by translation experts at UAB. This expands the serviceable market for ETD-Hub and offers greater accessibility as a value-add to prospective users.
- **Offer unique functionality:** ETD-Hub data from AutoDW connects to SemKG and AutoML, meaning debate directly informs model creation and AI engineering for ethically minded developers. Consequently, ETD-Hub offers a direct impact on the course of AI projects within the EU that competitors cannot yet match.
- **Focus on a clear niche:** we position ETD-Hub as *the* reference platform for AI and AI ethics deliberation in Europe, rather than a generic participation tool, thereby avoiding direct competition with broad platforms.
- **Flexible business and licensing models:** we may consider a mixed model (free core and value-added services; or tiered SaaS offering) for any monetisable elements in order to meet the needs of public sector, research, and private clients with different budgets and constraints.

We aim to perform some pilot deployments and reference cases for ETD-Hub (mostly in UK and Spanish Universities) during and immediately after ALFIE. This will enable us to support our planned route-to-market by:



- Developing **high-quality flagship case studies** and testimonials evidencing value and ROI, citable in stakeholder engagement.
- **Expanding datasets** to improve ROI for the entire integrated AutoML system from the outset.
- **De-risking adoption** for future customers by proving technical feasibility and impact in realistic contexts.
- Devising a **feature roadmap** to further optimise ETD-Hub for market need.

This clear roadmap for technical maintenance, feature development, and governance of the platform beyond the project reassures potential customers that ETD-Hub is a sustainable, actively maintained solution. This supports market access by reducing perceived vendor and technology risk, enabling evolution with regulatory and market needs, and providing the basis for long term strategic planning.

Longer term, we aim to deliver training programmes for administrators, moderators, policymakers, and facilitators. These programmes can support capacity-building and potentially operate as paid services (workshops, online courses, certification). This supports market entrance through:

- Addressing **organisational capability** barriers to ensure effective use.
- Improving user satisfaction to support **longer term adoption and reduce churn**.
- Opening an **additional revenue line** through professional services during the likely least commercialisable phase of the ETD-Hub route-to-market.

Project dissemination channels, conferences, policy briefings, and scientific publications can further increase visibility and show ETD-Hub's unique role in responsible AI governance. We have already led an ALFIE contingent at AIEOU2025 in service of this. Scientific dissemination activities can boost market access for ETD-Hub via:

- Positioning ETD-Hub as a **thought leader solution** in the emerging market for responsible AI tools, generating inbound demand.
- **Strengthening the brand** among technical and policy audiences.

After the project, we may also look to offer a stable private SaaS instance of ETD-Hub, able to support service-level agreements for clients. This offers various route-to-market benefits:

- Offers a use-case specific version of ETD-Hub as a **productionised service** with clear pricing, support, and risk allocation.
- Lowers barriers to entry for customers looking to utilise ETD-Hub to canvas advice on particularly **sensitive use cases**.

2.2.1.4 Potential customers and engagement activities

Due to its open forum nature, ETD-Hub offers value to an extremely diverse audience. This can be divided into two groups based on level of involvement in the initial development process; primary audiences (engaged in ETD-Hub development under ALFIE) and secondary audiences (not involved under ALFIE but likely to get value from ETD-Hub).

We expect primary audiences to be seeking insight from, or to lead, or to contribute to open AI policy, ethics, or technological discussions. Their involvement during development means this group will have a larger impact on ETD-Hub functionality and thus it may be more optimised for their use cases. This group may include:

- **EU and national policy makers and regulators** working on AI, data, digital rights, and related domains.
- **Public administrations** running consultations on the deployment of AI in public services (e.g. health, education, welfare, policing).
- **Research institutions including universities** focussed on AI, ethics, law, policy, and social sciences exploring debate or seeking AI engineering support via AutoML.
- **Educational institutions including universities** using ETD-Hub as a teaching tool for courses on AI ethics, digital democracy, or policy making.
- **Civil society organisations and Non-governmental organisations (NGOs)** focused on digital rights, consumer protection, labour rights, equality, and fundamental rights.
- **Industry actors** (AI developers, tech companies, sectoral players deploying AI) seeking structured engagement with stakeholders or ethical AutoML.

We anticipate secondary audiences that are seeking the same access to discourse and insight as our primary audience. The key difference is due to lack of involvement during development, the initial ETD-Hub offering is likely to be less optimised for their use cases. This group may include:

- **Professional associations and standardisation bodies** dealing with AI guidelines and sectoral codes of conduct.
- **Media and fact-checking organisations** interested in accurate, contextualised public opinion insights on AI.

Table 5. ETD-Hub potential customers and activities

Sector/ Community	Potential Customers	Activities
AI professionals and researchers / scientific and research community	- AI laboratories and confederations (e.g. CAIRNE, ELLIS) - AI Initiatives (e.g. appliedAI, AI4Europe/AIOnDemand) - Technology research institutes (e.g. Ada Lovelace Institute, Alan Turing Institute) - Universities with a strong focus on AI and Ethical AI (e.g. Universidad Complutense de Madrid, Middlesex University, Royal Holloway University of London, University of Greenwich)	- Continue to support development of domain taxonomies, sentiment labels and recommendation schemas for AI ethics datasets by providing data to ALFIE partners. - Run joint workshops / summer schools on participatory AI governance and integrating deliberation outputs into AutoML/model evaluation. - Plan to co-author scientific publications using ETD-Hub or ETD-Hub data (e.g. comparative ethics perceptions across use cases).

Legal and regulatory community	- AI governance and regulatory boards (e.g. EDPB, UK ICO, UK OFCOM) - Digital and AI commission offices (e.g. EU AI Office) - Rights agencies interested in AI (e.g. EU Agency for Fundamental Rights)	- Plan to use ETD-Hub to run demo structured stakeholder consultations for project test use case AI systems for partners Bosch, CTL, and UAB. - Encourage expert partners DBC to continue to provide legal/ethical feedback on ETD-Hub features (privacy-by-design, consent, retention policies).
Polymakers, advisors, and government bodies	- Policy units / commission departments interested in AI (e.g. DGConnect) - Technical parliament committees (e.g. LIBE, IMCO) - AI or technical ministries, groups, and units (e.g. UK DSIT, UK Sovereign AI Unit, EU Expert Group on AI) - Technical or AI thinktanks (e.g. IPPR, Policy Connect)	- Support DBC (through data provision) in developing policy briefs and impact assessment annexes based on ETD-Hub analytics for legislative proposals in T2.5 and T2.6. - Plan to develop workflow frameworks that demonstrate methodologies for collaboration: ETD-Hub provides data; organisations use it in comparative global studies on AI ethics and public perceptions.
Established Industry (AI-intensive sectors and corporate users)	- Digital innovation hubs (e.g. EDIH) - Large corporates investing in AI (e.g. Bosch, Publicis)	- Provide ETD-Hub datasets so consortium partners can tune AutoML constraints, fairness metrics and documentation before deployment. - Explore potential for additional sector pilots (e.g. finance, insurance, energy) using ETD-Hub.
Developers, startups and SMEs	- AI or ethics startups via launchpads (e.g. AI Startups Europe, Rise Europe, Y-Combinator) - SMEs interested in AI	- Explore feasibility of co-hosting community meetups and hackathons using ETD-Hub data to improve bias and transparency in models.
Educational institutions (universities, further education)	- Universities with a strong focus on AI and Ethical AI (e.g. Universidad Complutense de Madrid, Middlesex University, Royal Holloway University of London) - Further education institutions with a focus on citizenship, society, and AI (e.g. Association of Colleges, ETF Foundation)	- Support UAB in using ETD-Hub as a teaching platform in courses on AI ethics, digital democracy, public policy and law (student-run deliberations, simulated consultations). - Plan to develop education modules that integrate ETD-Hub exercises.

2.2.2 KEA2. AutoML platform (Owner: CTL, Co-owner: All)

2.2.2.1 Description

The AutoML platform is designed to translate deliberative ethical outcomes into actionable AI tools and solutions. It consists of two primary layers: an interaction layer relying on generative AI large language models (LLMs) and an AutoML engine. The technology automates the machine learning process, including algorithm selection, hyperparameter tuning, and metadata management, while ensuring models are unbiased and compliant with the EU regulatory framework. It features integrated explainable AI to provide users with insights into model decision-making.

2.2.2.2 Value Proposition

The ALFIE AutoML platform offers a transformative value proposition centred on the democratisation of AI development, enabling individuals and organisations to create machine learning models without requiring deep technical expertise. By providing an interaction layer that utilises LLMs and speech recognition, the platform is allowing non-experts to design and build trustworthy, unbiased, and ethically compliant AI tools through natural language dialogue. A fundamental differentiator is the platform's rigorous alignment with the EU regulatory framework, specifically the European AI Act, ensuring that every generated model adheres to stringent legal and ethical standards from the design phase. Unlike conventional black-box automated tools, the platform features integrated explainable AI (XAI) that delivers transparent and accessible insights into model decision-making, which fosters stakeholder trust and user acceptance across sensitive sectors. The value is further enhanced by an ethics-by-design approach that incorporates a bi-directional feedback loop with the ETD-Hub, effectively translating policy recommendations into actionable technical constraints and bias mitigation strategies, prioritising environmental sustainability.

2.2.2.3 Potential research and market exploitation

The research potential of the ALFIE AutoML platform lies in its pioneering integration of fairness-awareness and environmental sustainability into automated machine learning processes. Strategies include positioning the platform as the lead EU-compliant solution for ethical AI and fostering a collaborative ecosystem where developers and policymakers converge.

Research efforts are directed toward minimising the carbon footprint and energy consumption of AI training through model quantisation, mixed-precision training, and the prioritisation of medium-sized models. Furthermore, the platform serves as a testbed for evolving algorithms that incorporate human-in-the-loop feedback, allowing the system to learn dynamically from user interactions and adapt to changing data environments. Research efforts are also directed toward the ability to handle multi-modal data, integrating text and images, to provide a robust framework for investigating cross-subject variability in high-stakes domains like autonomous vehicle safety. Additionally, the platform enables research into multilingual interaction, exploring how to support both major and minor European languages to democratise AI access across the continent. The exploitation strategy targets high-stakes sectors where there is a critical demand for transparent and unbiased decision-support tools. Commercialisation pathways include offering possible subscription plans for main features, licensing options for industry partners, and tailored consulting services to help integrate with existing systems and building strong partnerships with SMEs, research groups, and industry leaders. This will help

the platform become a trusted source for responsible AI and support the launch of ready-to-use products.

2.2.2.4 Potential customers and engagement activities

Potential end-users include software developers, data scientists, domain experts in sensitive sectors (e.g., healthcare, finance), SMEs, and public authorities seeking to deploy trustworthy AI. These primary users are joined by machine learning engineers who aim to accelerate their development workflows using automated bias detection and ethical compliance tools. Within the industrial sector, compliance officers and managers can utilise the platform for specialised tasks. SMEs and startups could find the tool valuable for its democratisation of technology, allowing those without deep coding expertise to build models that improve their products and services. Furthermore, public sector organisations and regulatory bodies can leverage the platform to validate policy scenarios through generated models, while academic researchers can use it to explore new fairness metrics and benchmark AI safety protocols. The platform is also designed to be inclusive, using speech-to-text to evaluate digital accessibility.

Engagement activities are structured to foster a collaborative and user-centred ecosystem through a variety of diverse outreach channels. This process involves conducting participatory workshops and surveys to capture real-world user needs and requirements. Technology developers and startups can be invited to participate in hackathons and innovation challenges that encourage the creation of creative, ethically aligned AI solutions. To ensure sustained adoption and market trust, ALFIE aims to organise demonstrations in industrial and academic environments, complemented by training programmes and educational guides that empower users with varying levels of technical proficiency to operate the platform effectively.

Table 6. AutoML platform potential customers and activities

Sector/ Community	Potential customers	Activities
Industry	Startups, SMEs, and large-scale industries	- Designing and building unbiased and ethically compliant AI tools to address organisational challenges. - Participating in hackathons, innovation challenges, and consultation services.
Technology and AI professionals	Software developers, data scientists, and machine learning engineers	Integrating bias-free ML modules into proprietary software.
Public sector and governance	Policymakers, regulatory bodies, and public authorities	Utilizing AutoML to validate policy scenarios through generated AI models.
Research and academia	Researchers in AI and ethics, ethics experts, and academic institutions	Exploring new fairness metrics and benchmarking AI safety and exploring automated machine learning techniques for responsible development.
General public	Citizens and members of the general public	Engaging with accessible interfaces and participating in transparency initiatives regarding AI's ethical and legal dimensions.

2.2.3 KEA3. Data warehouse (AutoDW) of ethical data (Owner: CERTH)

2.2.3.1 Description

The AutoDW is ALFIE's central data hub and orchestration layer. It provides secure storage for datasets, experiment metadata, and trained AI models, and provides multi-database access for all system components. In addition to data management, AutoDW integrates an event-driven message-broker mechanism that coordinates workflows across services.

2.2.3.2 Value Proposition

The AutoDW delivers high technical and operational value by serving as the unified backbone for data, models, and workflow coordination within ALFIE. For developers and integrators, it simplifies access to heterogeneous data through a single, secure storage layer and ensures full traceability of experiments, model versions, and service outputs. For system operators, its event-driven orchestration reduces coupling between services, enabling scalable, resilient, and easily maintainable AI pipelines. For end-users and organisations, AutoDW enhances trust and compliance by enforcing structured data governance, reproducibility, and transparent lifecycle tracking of all AI assets.

2.2.3.3 Potential research and market exploitation

The AutoML Data Warehouse (AutoDW) serves as a unified, end-to-end AI operations storage layer that provides significant research and market potential by facilitating reproducible AI benchmarking and orchestrating complex pipelines across research labs, data science groups, and universities. Its integrated event-driven message-broker mechanism coordinates workflows across diverse services, enabling a scalable, resilient infrastructure that manages the entire AI lifecycle, from raw data ingestion to trained model archival. By embedding ethical AI governance into its core design, AutoDW ensures full traceability of experiments and model versions, maintaining clear and auditable data lineage.

Strategically positioned to meet the growing demand for transparent, regulated, and auditable AI infrastructure, especially within the European market aligning with the EU AI Act and GDPR, AutoDW is optimized for on-premise, hybrid, and sovereign AI deployments. Its design reduces vendor lock-in by prioritizing open-source and source-available solutions on its core such as MongoDB and MinIO, promotes extensibility, and tightly integrates with ALFIE's multi-agent intelligence layer. Differentiating itself from hyperscale cloud solutions, AutoDW combines native workflow coordination and experiment management within a single platform. The consortium plans to leverage open-source releases, training workshops, fairness assessment templates, pilot demonstrators, and comprehensive integration support for both public and private adopters, establishing AutoDW as a strategic operational backbone for enterprises and SMEs seeking trustworthy, sustainable, and compliant AI solutions.

2.2.3.4 Potential customers and engagement activities

Potential end-users for the AutoDW include AI developers, data scientists, public administrations, regulators, research institutions, and SMEs who require a unified technical backbone for managing datasets, models, and workflow coordination. These core audiences are complemented by AI research laboratories and universities that utilise the platform for centralised dataset management, model artefact curation, and reproducible AI benchmarking across complex pipelines. In the public sector, government agencies and digital governance offices represent key customers who can employ the technology for the secure storage and

governance of sensitive public data while ensuring that AI development remains transparent and auditable. Furthermore, high-risk industrial users in sectors such as finance, healthcare, and HR tech find significant value in the platform's ability to provide centralised risk data repositories. Additionally, civil society organisations and digital rights NGOs can adopt the AutoDW as a foundational infrastructure component to maintain evidence-based research repositories and monitor AI accountability.

Engagement activities for the AutoDW are structured to build technical credibility and demonstrate its utility as a central hub through a phased validation approach and continuous stakeholder interaction. The plan is to support market entry by providing open-source releases, technical training sessions, and hands-on workshops where potential adopters can interact with fairness assessment templates and pilot demonstrators. Early engagement focuses on conducting pilot programmes and strategic partnerships with research institutions and public agencies to validate complex orchestration workflows and generate impactful reference cases. Stakeholders are also invited to participate in demonstration workshops and conferences to witness the warehouse's ability to coordinate heterogeneous data types and event-driven AI pipelines in real-world environments. This collaborative ecosystem is maintained through a bi-directional feedback loop with the ETD-Hub, ensuring that storage and governance features evolve in response to direct stakeholder input and shifting regulatory requirements defined by the EU AI Act.

Table 7. AutoDW potential customers and activities

Sector/ Community	Potential customers	Activities
AI Professionals and researchers	AI research labs, universities, data science groups	- Centralized dataset management and curation - Model artifact management and experiment tracking - Orchestration of complex AI pipelines - Structured storage for specialized AI data types - Reproducible AI benchmarking
Public sector and policymakers	Government agencies, regulators, digital governance offices	- Secure storage and governance of sensitive public data - Transparent, auditable AI development
Industry (High-risk AI users)	Finance, healthcare, HR tech, public service providers	- For finance could be centralized risk and portfolio data management. - For healthcare could be unified patient and clinical dataset management. - For HR tech could be centralized talent and performance data repository. - For public service providers could be integrated citizen service data hub.
Civil society and ethics communities	Digital rights NGOs, ethics boards, advocacy groups	- Centralized evidence and research data repository. - AI model transparency and auditability. - Event-driven coordination for collaborative projects.

- | | | |
|--|--|--|
| | | - Governance, accountability, and compliance monitoring. |
|--|--|--|

2.2.4 KEA4. Semantic knowledge graph (SemKG) of ethical data (Owner: EHU, Co-owner : CTL)

2.2.4.1 Description

ALFIE's Semantic Knowledge Graph (SemKG) is a processing component that inputs textual data and processes it to produce a graph, whose nodes and edges represent information expressed in the input text. Processing involves the identification of named entities and relations among them, normalisation, duplicate removal and linking to existing ontologies in the domain of application. The resulting graph of interlinked entities carries/conveys information and knowledge from the original input text.

Within ALFIE the input of the knowledge graph is json-formatted discussions on ETD Hub, that consist of individual posts. However, the SemKG can be used as an independent software component and can be adapted/used in other contexts, where it is necessary to populate a knowledge graph from text. Notably, depending on the application domain, ontologies that the semantic knowledge graph considers may have to be reviewed/replaced/enriched by others that cover/match the application domain.

2.2.4.2 Value Proposition

The SemKG software is a system that can convert forum discussions into usable knowledge in the form of a graph. The software can consider multiple ontologies of knowledge areas related to the forum discussion topic, to link concepts and concept relations to well defined and documented knowledge sources. The software can be useful in effectively compiling long discussions into condensed and externally verified knowledge. In turn, this form can be the input of other processing systems, including LLMs. Knowledge graphs can be used to add reliable context to LLMs, using technologies such as Retrieval-Augmented Generation (RAG).

In the context of ALFIE, the software populates a Semantic Knowledge Graph using discussions about AI Ethics in the ETD Hub and several ontologies in the domain of AI and Ethics. The particular SemKG instance can be used further than ALFIE, in applications in the same domain, AI Ethics.

2.2.4.3 Potential research and market exploitation

The SemKG of ethical data presents substantial research potential through its ability to transform unstructured textual information into highly structured, machine-readable representations of entities and their complex relationships. It serves as a vital tool for knowledge engineering by utilizing natural language processing methods, such as Named Entity Recognition and relation extraction, to link abstract ethical concepts to well-documented domain ontologies. Researchers can leverage this technology to provide reliable context for LLMs through Retrieval-Augmented Generation (RAG). Furthermore, the graph is central to investigating how high-level ethical guidelines and regulatory frameworks can be translated into actionable rules for automated code generation, thereby facilitating the creation of AI systems that are inherently aligned with societal values.

In terms of market exploitation, this technology is primarily targeted at Information and Communication Technology (ICT) companies that require sophisticated knowledge management solutions, such as advanced indexing, filtering, and semantic search capabilities. Its value proposition is anchored in its ability to offer a machine-operable layer of ethical awareness, which is increasingly critical for organisations needing to demonstrate compliance with evolving regulations like the EU AI Act. The software component is designed for high adaptability, allowing it to be repurposed across various scientific and technical domains beyond AI ethics, provided that relevant ontologies are available. Exploitation strategies involve a multi-tiered approach, starting with making the software freely available for research and download to encourage early adoption and build a robust community of users. Long-term commercialisation plans include offering possible subscription-based access to curated ethical knowledge bases, professional consultation services for AI experts, and specialised training programs to assist organisations in integrating these ethical frameworks into their own development pipelines.

2.2.4.4 Potential customers and engagement activities

The primary customers for the SemKG are ICT companies that focus on building advanced data warehouses and sophisticated knowledge management systems. These organisations need to follow strict legal obligations that their systems are both effective and fair. Beyond the ICT sector, the technology is targeted at AI development organisations, data scientists, and researchers who are committed to creating responsible AI models that align with specific ethical values and human rights. Additionally, policymakers and regulatory bodies represent a key customer segment, as they can utilise these structured representations of ethical discourse to monitor compliance and inform the creation of new AI governance frameworks. Developers seeking to understand public and expert perceptions of AI also stand to benefit from the actionable insights captured within the graph's interconnected entities and relations.

A central component of the engagement strategy involves making the software component freely available for download and research purposes, which encourages early adoption and fosters a collaborative community of users. Moreover, facilitate professional engagement through specialised training programs and consultative services designed to help organisations integrate ethical reasoning into their existing AI development pipelines. Furthermore, the plan involves the participation in workshops and international conferences to demonstrate its capabilities to policymakers and industry leaders.

Table 8. SemKG potential customers and activities

Sector/ Community	Potential customers	Activities
ICT	Information and Communication Technology (ICT) companies	Developing and utilizing tools for advanced indexing, filtering, and deep semantic searches within ethical datasets.
AI research and development	AI professionals, data scientists, ML engineers, and academic researchers	Engaging in collaborative research. The software will be freely available for download, and conducting comparative evaluations against existing literature.

Legal, regulatory, and policy	Policymakers, government bodies, legal professionals, and regulatory communities	Aligning ethical considerations with legal frameworks and using the graph to translate high-level guidelines into actionable rules for automated AI governance.
Industry and business	Startups, SMEs, and large-scale industries	Consultation services to integrate ethical reasoning into specific industrial development pipelines.
Education and professional training	Universities, educational institutions, and technical training centers	Integrating ethical AI principles into curricula and hosting specialized webinars and training programs for future professionals.
Advanced AI engineering	Developers of LLMs and AI auditors	Utilizing the graph to provide reliable, externally verified context for LLMs through RAG.

2.2.5 KEA5. Agentic Core (Owner: KINIT)

2.2.5.1 Description

The ALFIE Agentic Core is a multilingual, AG-UI-enabled agentic service designed to empower users to create trustworthy AI through an intuitive chat interface and natural language interaction. It functions as the central orchestrator and intelligent backend for the ALFIE platform, managing the end-to-end machine learning workflow from initial requirement gathering to model development. Its core features include user intent recognition, the ability to steer complex conversations, and the seamless execution of machine learning tasks, including bias detection, AutoML, and explainable AI (XAI). The service provides use case assessment and guidance by a specialized knowledge base, which includes insights from the ETD-Hub and the Semantic Knowledge Graph (SemKG) to ensure that AI development remains grounded in ethical and legal standards (to be more precise, a set of standards represented by a set of relevant documents, such as AI Act). Architecturally, the core implements a sophisticated orchestrator-workers design pattern. A manager agent acts as a central coordinator that delegates tasks to specialized worker agents, intelligently combining parallel agent outputs into coherent, actionable guidance for the user. Technically, the Agentic Core is specified as a FastAPI service integrated with connectors for a PostgreSQL database, the Task Manager service, the AutoDW, and a LLM provider service. It relies on AG-UI-based communication with the frontend to standardise interactions and is built to be provider-agnostic, ensuring the flexibility to switch between different LLM models as the technology evolves.

2.2.5.2 Value Proposition

The Agentic Core revolutionises AI accessibility by enabling non-experts to create trustworthy AI through a convenient chat UI. Functioning as the platform's intelligent backend, it uses a multi-agent framework to guide users through the complete machine learning workflow, allowing them to consult specialized agents at each step of the process. By utilizing a manager agent to analyse intent and delegate tasks asynchronously, the core offloads the burden of gathering complex requirements for ML tasks, effectively making the development workflow smoother and more approachable for users regardless of their technical expertise. Furthermore, its robust multilingual support for multiple European languages ensures the tool



is inclusive, thereby increasing the number of potential adopters across diverse linguistic regions. Beyond ease of use, the Agentic Core is designed to make creating AI more aligned with regulatory compliance and ethical standards, specifically the EU AI Act. It achieves this by providing context grounded in a given knowledge base and the SemKG, which retrieves ethical considerations and facilitates the generation of fairness assessments and bias reports. The service's sophisticated architecture follows an orchestrator-workers design pattern that is provider-agnostic, ensuring it can switch between different large language models as needed. This modular and decoupled approach allows the service to be adopted for different AutoML engines, knowledge bases, and AG-UI-compatible frontends, providing the flexibility and scalability required to maintain an innovative edge in the rapidly evolving AI market.

2.2.5.3 Potential research and market exploitation

The Agentic Core developed within the ALFIE project offers significant potential for both scholarly investigation and commercial growth, emphasizing an ethics-first, multi-agent orchestration approach with a natural-language interface. This technology focuses on democratizing machine learning for non-expert users, enabling critical research into performance degradation across various European languages, particularly under-resourced ones, and semantic alignment of complex legal, ethical, and ML domain terminology in diverse linguistic contexts. Integrating LLMs and speech-to-text technologies creates a fertile ground for exploring maintaining model accuracy and ethical compliance.

From a market perspective, the system's comprehensive multilingual support allows ALFIE to target non-English speaking markets where existing tools often lack localized language interfaces. Key application segments include education, state administration, healthcare, finance, and critical infrastructure, industries where ethical risks, fairness, and transparency are paramount. The agentic core acts as a sophisticated orchestrator, coordinating specialized agents for data analysis, search, and ethics using a manager-workers pattern, providing strategic avenues for market entry and high-level research. It supports translating natural language requirements into structured messages for AutoML and enabling human-in-the-loop systems for dynamic algorithm refinement. Exploitation pathways include possible subscription models, custom framework development, consulting services, and enterprise integration into MLOps pipelines. To ensure enterprise readiness and scalability, the system will be containerized as a well-documented Docker image, facilitating easy deployment and validation, especially through early stakeholder engagement, particularly with public sector clients.

2.2.5.4 Potential customers and engagement activities

Researchers, domain experts, software developers, data scientists, ML engineers represent the core user base for the Agentic Core, which serves as the intelligent orchestrator of the machine learning lifecycle. This audience is joined by ethics experts and machine learning researchers who utilise the platform to study the semantic alignment of legal and technical language across various European contexts.

Engagement activities are designed to build a robust ecosystem through a combination of technical validation and community outreach. Scientific and academic engagement is maintained through the publication of research papers, and the delivery of seminars and workshops. To support broader adoption, the project develops comprehensive educational

guides, webinars, and training modules that empower users with varying levels of expertise to navigate the machine learning workflow.

Table 9. Agentic Core potential customers and activities

Sector/ Community	Potential customers	Activities
Scientific and research community	ML scientists, ethics experts	- Publish papers - Present in seminars and workshops - Reach out for collaboration - Organize panel discussions - Share on socials
Private sector	ALFIE stakeholders	- Engage in collaboration - Provide consultations

2.2.6 KEA6. Web interface with XAI features (Owner: CTL, Co-owner: CERTH, KInIT)

2.2.6.1 Description

The web interface is a user-friendly digital platform designed to make unbiased AutoML tools accessible to diverse users with varying skill levels. It serves as the primary interaction layer, functioning as a centralized gateway for users to explore, manage, and discover AI models, datasets, and files. Key features include an intuitive chat-like interface powered by LLMs, allowing for natural language and speech interaction in multiple languages. It enables users to monitor AI workflows, access explainability (XAI) insights and detect bias in AI models and datasets through a clean, modern dashboard. Integrates with ALFIE pipelines and supports human-readable and machine-readable outputs.

2.2.6.2 Value Proposition

The ALFIE web interface fundamentally democratizes access to complex AI technologies by providing an intuitive interaction layer that removes technical barriers for non-experts. By allowing users to interact with the platform through natural language and speech in multiple European languages, it empowers data scientists, developers, and domain experts to create AI models without needing deep machine learning expertise. This inclusive design is specifically tailored to accommodate a wide range of skill levels and physical needs, ensuring that AI-driven innovation is accessible to diverse populations who might otherwise be excluded from participating in the AI landscape. A core element of the interface's value proposition is its rigorous integration of ethical safeguards and legal compliance directly into the user workflow. The platform delivers transparent, bias-aware insights into AI behaviour, supports EU AI Act and GDPR compliance, improves trust, and enables users to identify and mitigate unfair outcomes. This commitment to fairness is bolstered by explainability features that provide transparent, human-readable insights into how models make decisions. The interface fosters a trustworthy environment where users can confidently validate the impartiality and reliability of their AI solutions before deployment. The interface delivers unique value by offering bi-directional feedback loops with the ETD-Hub, ensuring that technical implementation remains continuously refined by societal values and expert deliberations. Furthermore, the platform prioritizes environmental sustainability allowing users to balance high performance with energy efficiency. This holistic approach increases the likelihood that every developed model is not

only technically accurate but also socially responsible, legally compliant, and ready for real-world application.

2.2.6.3 Potential research and market exploitation

The ALFIE web interface's research is rooted in advancing Europe's leadership in AI fairness and establishing new methodologies for technological accountability. A primary research avenue involves the development of sophisticated protocols to detect bias within datasets and pinpointing unfair model behaviours throughout the development lifecycle. This research includes exploring mechanistic interpretability and cognitive-grounded explanation models specifically designed to be understood by human end-users rather than just technical specialists. Furthermore, scientific efforts will focus on the interdisciplinary alignment of ethical principles and legal requirements, translating high-level frameworks like the EU AI Act into actionable technical constraints within the AutoML pipeline. This is significantly enhanced by the integration of SemKG that formalises stakeholder discourse from the ETD-Hub into machine-readable formats, alongside the testing of inclusive design standards to ensure that intuitive interfaces remain accessible to vulnerable populations. The web interface addresses a critical demand for accountable and transparent technology, particularly in high-risk sectors where rigorous compliance with the EU AI Act and GDPR is mandatory. The interface allows non-experts to generate ethically compliant models through natural dialogue, providing a substantial value proposition for SMEs and corporate HR departments requiring bias diagnostics to prevent discriminatory outcomes. Future financial sustainability could be supported through a multi-layered revenue model, including possible subscription-based access as a service, bespoke consulting for fairness audits, and commercial licensing. The platform's primary competitive advantage lies in its human-centric and auditable experience, featuring bi-directional feedback loops for continuous ethical refinement and multilingual support, ensuring a level of transparency and regulatory assurance often absent in proprietary commercial alternatives.

2.2.6.4 Potential customers and engagement activities

Data scientists, AI developers, research institutions, software developers, domain experts and organisations deploying high-risk AI in sectors who require trustworthy AI solutions, public authorities, regulators, governmental bodies, SMEs, and citizens participating in transparency initiatives, AI developers, data scientists, public administrations, regulators, research institutions, SMEs, and organisations deploying high-risk AI represent the primary customer base for the ALFIE web interface. This broad audience is further extended to include corporate HR departments and compliance officers who can utilise the platform to conduct recruitment fairness audits and integrate bias diagnostics into their processes. The interface also serves specialised domains including healthcare providers, and educators who require field-specific dashboards that incorporate explainable AI (XAI) features to interpret complex model outputs. Because the platform is designed to democratise access to AI by removing technical barriers, it is highly applicable to the industry, IT, and government sectors where non-experts seek to design advanced, ethically compliant models through natural language dialogue. Furthermore, the system is specifically tailored for civil society members and vulnerable groups, who rely on inclusive design standards and multilingual speech capabilities to navigate AI resources. By strictly adhering to the requirements of the EU AI Act and GDPR, the web interface establishes a unique market position as a primary interaction layer for organisations that prioritise transparency and regulatory assurance. Engagement activities are designed to build a

collaborative ecosystem where technical development is continuously refined by direct stakeholder feedback. AI professionals and research facilities will be invited to participate in demonstration workshops that leverage the interface's code-generation and model exploration capabilities. Public sector authorities and governmental agencies are engaged through policy roundtables and implementation pilots that demonstrate how technical constraints can be visualised to reflect evolving legal frameworks. To ensure universal accessibility, the project organises participatory co-design workshops and user-testing sessions involving elders and individuals with disabilities to validate that the interface meets rigorous usability standards.

Table 10. Web interface with XAI features potential customers and activities

Sector/ Community	Potential customers	Activities
AI professionals and researchers	AI research labs, universities, data science groups	- Provide reproducible fairness benchmarks and datasets. - Collaborate on research studies and publications. - Deliver workshops on fairness-aware explainability.
AI industry and research	AI Experts, R&D facilities, finance, healthcare, HR tech, public service providers	- Demonstration workshops, collaborative research, and open-source project development. - Participating in hackathons and innovation challenges to build new models using the interface's capabilities. - Collaborating on R&D pilot projects to maintain the innovative edge of the platform and provide testimonials. - Engaging in constrained academic experiments during demonstration workshops to test the interface's code-generation and functionalities - Pilot bias audits using the ALFIE XAI module. - Support integration into existing ML pipelines. - Produce fairness and transparency reports for internal governance. - End-user training sessions and integration of XAI features into field-specific dashboards. - Integration into Business Partner Screening and recruitment fairness audits.
Public sector and policymakers	Governmental agencies, public authorities, government agencies, regulators, digital governance offices	- Consultations on AutoML technologies and policy implementation pilots. - Attending policy roundtables and targeted workshops to observe how technical implementations can reflect the EU AI Act. - Using the interface for comparative policy analysis by visualizing how different ethical constraints affect model outcomes. - Participating in demonstrations to validate the interface as a reference point for trustworthy AI.

		- Demonstrate bias detection workflows. - Support compliance with the EU AI Act. - Prepare policy-oriented fairness assessment summaries.
Legal and advocacy	Compliance Officers, Digital Rights Advocates	- Testing explainability (XAI) features to evaluate if the model decision paths meet "right to explanation" requirements.
Civil society and ethics communities	Vulnerable groups, elders, people with disabilities, digital rights NGOs, ethics boards, advocacy groups	- Participating in participatory co-design workshops and user-testing sessions to ensure the UI meets accessibility standards. - Engaging with "Be A(I)ware" citizen awareness campaigns. - Provide accessible fairness visualisations. - Facilitate public-facing awareness activities. - Co-develop guidelines for responsible AI deployment.

3 Market analysis

3.1 Overview

A market analysis serves as a foundational pillar for any innovation-driven project, providing the insights necessary to navigate complex and competitive industrial landscapes. By assessing the competitive environment, a market analysis provides an overview of the marketplace, enabling the identification of the technological strengths and weaknesses of existing competitors, revealing the underlying motivations and preferences of target customers. By aligning project goals with these user needs, the consortium can guarantee that its technical solutions are not created in isolation but are accurately tailored to address real-world issues and meet market demands.

The market analysis plays a crucial role in identifying and classifying potential risks that could threaten the project's long-term success, and foresee challenges such as market saturation, disruptive technological advancements, or sudden regulatory changes, like new data privacy laws or ethical AI standards. Recognizing these risks early enables the project to shift from a reactive to a proactive approach, crafting mitigation strategies that protect its assets while also seizing emerging opportunities and trends. Additionally, a detailed analysis facilitates the segmentation of the market into distinct niches, allowing the project to focus on high-growth areas where its Key Exploitable Assets (KEAs) can have the greatest impact. This targeted strategy ensures that marketing efforts and product development are directed toward the most promising opportunities for commercial success and societal benefit.

Beyond initial strategic positioning, the insights gained from a comprehensive market study provide a data-driven basis for making complex business decisions and it ensures that the project's value proposition resonates clearly with its intended audience. Crucially, market analysis is not a one-time event but an ongoing, iterative process throughout the entire project lifecycle. This continuous feedback loop allows for agile decision-making and flexible resource allocation, allowing the project to respond effectively to external changes and maintain its relevance. Ultimately, this strategy reduces resource waste and enhances the probability of both commercial and societal success by keeping the project at the cutting edge of a rapidly evolving and highly competitive marketplace.

The market analysis for each Key Exploitable Asset (KEA) within the ALFIE project offers valuable insights into the competitive landscape and serves as a vital foundation for strategic decision-making and the formulation of effective commercialisation plans. By understanding the competitive environment, it is possible to identify unique opportunities for each KEA. Additionally, regulatory considerations are assessed to anticipate compliance requirements, approval processes, and legal limitations, helping to reduce risks associated with market entry. Leveraging these insights enables the development of tailored strategies that position the KEAs effectively within the target market. This includes aligning the KEAs with specific customer needs, refining value propositions, and identifying the most promising channels for commercialisation. The goal is to support sustainable growth by ensuring that each KEA's market entry is strategically aligned and adaptable to the market, thereby maximizing its commercial success and long-term viability.

3.2 Market landscape

The current market research offers a comprehensive analysis of the rapidly changing technological landscape, focusing on both the broader Artificial Intelligence (AI) industry and the emerging Responsible and Ethical AI. By identifying key market drivers, such as the rapid shift toward agentic AI systems and the implementation of the EU AI Act, this analysis sketches a competitive landscape where ALFIE is positioned to provide unique value. Ultimately, the research highlights the significant prospects for the ALFIE project and its individual components within the European and global markets.

The global Artificial Intelligence (AI) market in 2025 has reached a pivotal turning point, transitioning from a phase of speculative hype toward deep operational integration across all major industrial sectors. Financial assessments value the total global AI market is experiencing significant expansion, with current valuations for 2025 ranging between USD 371.71 billion² and USD 757.58 billion³, depending on the inclusion of hardware and specialized services. The market is projected to maintain a compound annual growth rate (CAGR) of approximately 19% to 30% through 2034⁴, eventually exceeding USD 3.6 trillion (Figure 1).




Figure 1. AI Market Size estimates in the period between 2024 and 2034

In terms of specific technologies, deep learning (37.4%) and machine learning (27.7%) continue to represent the largest technological share, providing the backbone for predictive analytics and automated decision-making across the globe.

² <https://www.marketsandmarkets.com/Market-Reports/artificial-intelligence-market-74851580.html>

³ <https://www.precedenceresearch.com/artificial-intelligence-market>

⁴ <https://www.fortunebusinessinsights.com/industry-reports/artificial-intelligence-market-100114>

However, generative AI remains the fastest-growing sub-sector, characterized by a growth rate that outpaces traditional models as organisations deploy multimodal systems capable of generating sophisticated text, code, and media, projected to grow at a CAGR of 22.9%³. A critical market driver in 2025 is the rise of agentic AI, where systems are no longer passive responders but active agents capable of executing multi-step workflows with minimal human intervention. Nearly 75% of new data centres⁵ are being designed specifically for AI workloads. The rise of "AI Factories"⁶, high-performance computing hubs dedicated to model training, is a critical trend in 2025.

The competitive landscape is currently a multi-tiered environment where hyperscale cloud providers like Microsoft, Google, and Amazon compete with frontier model developers such as OpenAI⁷ and Anthropic⁸. Simultaneously, a powerful open-source movement, led by models like Meta's Llama⁹ series, has democratized access to high-performance AI, allowing smaller firms to build sovereign systems on-premise.

While software historically commanded the largest portion of the market share, the current year is seeing a significant surge in the services segment. Within this ecosystem, the market for responsible and ethical AI has matured from a niche concern into a primary business requirement. Enterprises move beyond basic model adoption and require specialized expertise for custom integration, ethical auditing, and regulatory alignment, particularly in response to the stringent requirements of the EU AI Act. The enforcement of the EU AI Act¹⁰ and ISO 42001¹¹ has transformed "Ethical AI" from a corporate social responsibility initiative into a mandatory legal requirement. As consumer trust becomes a measurable financial asset, companies are prioritizing investments in transparency tools, bias mitigation platforms, and explainable AI interfaces. This ensures that as AI becomes more pervasive, it remains aligned with human values and legal standards, effectively bridging the gap between raw technological capability and sustainable societal benefit.

This has created a massive secondary market for AI auditing, bias detection, and compliance automation tools. The global market for Responsible and Ethical AI in 2025 has undergone a fundamental transformation, evolving from a set of voluntary corporate guidelines into a highly structured and rapidly expanding industrial sector. This market is now a critical component of the broader AI landscape, driven by the realisation that trust, transparency, and accountability are the essential "trust layers" required for AI to scale within high-stakes environments.

Whilst the notions Responsible AI and Ethical AI don't coincide¹², they have sufficient overlap in practice for market metrics to be indicative. According to a February 2025 report¹³, the global market Responsible AI was worth USD 910.4 million in 2024. It is projected to reach a value

⁵ <https://www.jll.com/en-us/insights/market-outlook/data-center-outlook>

⁶ <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>

⁷ <https://openai.com/>

⁸ <https://www.anthropic.com/>

⁹ <https://www.llama.com/>

¹⁰ <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689>

¹¹ <https://www.iso.org/standard/42001>

¹² <https://www.techtarget.com/searchenterpriseai/feature/Responsible-AI-vs-ethical-AI-Whats-the-difference>

¹³ <https://market.us/report/responsible-ai-market/>

of USD 47,159 million in 2034, with an estimated CAGR of 48,40% in the period 2025-2034¹⁰. Regarding market segments, software solutions have dominated comfortably, commanding 68.6% of the total market share. Their dominance is expected to continue given that increasingly the exigencies of fairness, auditability and interpretability are increasingly being satisfied through software innovations (Figure 2).




Figure 2. Global Responsible AI market size estimate between 2024-2034

Strategic assessments currently value the global AI governance market at approximately USD 309 million in 2025, with projections indicating a surge to USD 4.83 billion by 2034. This trajectory represents an extraordinary compound annual growth rate of over 35%, reflecting a shift where ethical compliance is no longer viewed as a peripheral cost but as a core business activity that directly influences return on investment and brand equity¹⁴ (Figure 3).

The primary catalyst for this growth is the EU AI Act, which entered a critical enforcement phase in early 2025. This landmark legislation has established mandatory requirements for high-risk AI systems, including rigorous fundamental rights impact assessments, data governance standards, and human oversight protocols. As these regulations become a benchmark for global trade, organisations worldwide are investing in specialized software solutions to automate compliance and auditability. Consequently, the software segment, which includes platforms for bias detection, explainable AI (XAI), and model lifecycle monitoring, now commands roughly 66% of the governance market share. Simultaneously, the services segment is expanding at an even faster rate as companies seek expert consulting to navigate the complex interplay between emerging local laws and international standards.

¹⁴ <https://www.precedenceresearch.com/ai-governance-market>




Figure 3. AI governance market size (2025 - 2034)

The key market drivers of Ethical AI include increasing customer and regulator demand for ethical technologies. On the customer side, more than 70% of customers are concerned about AI bias when it comes to hiring, banking and legal matters, whilst 68% of workers would like their employers to emphasize ethical AI¹⁵. In addition, 60% of investors factor environmental, governance and social considerations into their investment choices. Another leading driver for customer demand for Ethical AI is the complexity of the most common AI models. More than 65% of businesses utilising AI technologies replied, when surveyed by the study, that they do not know how the AI models produce their outputs. This perception of AI, as a “black-box” technology, increases concerns about bias and compliance breakdown and fuels the demand for explainability and transparency to be built into the technology itself. In addition, according to an IBM report stated in a market research study¹⁶, only 35% of people worldwide trust organisations’ usage of AI. Trust in AI is a major driver for its widespread adoption, and a drive which could be easily documented. To this end, the development of Explainable AI (XAI) is a major technological advancement which could increase trust in AI through the opening up of the technological “black-box”. A final key driver for the adoption of ethical AI products is the proliferation of the use of AI in decision-making in sensitive domains such as healthcare diagnostics, banking, financial services and insurance (BFSI) and law enforcement domains¹⁷.

Regionally, North America leads the market with an estimated market value of approximately USD 490 million in 2024 according to one report¹⁷. According to the same report, Europe is the second-largest market with a value of USD 370 million, whilst the Asia Pacific region is the

¹⁵ <https://www.pristinemarketinsights.com/responsible-ai-rai-market-report>

¹⁶ <https://market.us/report/responsible-ai-market/>

¹⁷ <https://growthmarketreports.com/report/ai-ethics-market>

fastest-growing market as it is experiencing a rapid digital transformation along with a significant investment in AI ethics frameworks.

Competitive dynamics in 2025 are characterized by a diverse ecosystem of stakeholders. The market landscape composition consists in a mix of established technology giants, specialised startups, consulting firms, academic institutions and start-ups¹⁵. Major players from all of these categories often collaborate with industry associations, regulators and advocacy groups in order to shape the evolution of regulation and standards to their advantage. In terms of key players, IBM is regarded as a pioneer in ethical and responsible AI, whilst other leading players include software giants Microsoft, Google LLC and DXC Technology, accounting giants Deloitte Touche Tohmatsu Limited, consulting firm Accenture, Siemens, and many others.

A key comprehensive ethical AI tool is IBM's Watson OpenScale platform, which, according to its description, detects and mitigates bias and drift, as well as enables AI models to operate with transparent and explainable outcomes¹⁸. Other major ethical AI tools include Google's What-If tool¹⁹, which enhances transparency and improves fairness, TensorFlow's Responsible AI Toolkit²⁰ which integrates responsible AI practices into ML workflows, and the Organisation for Economic Cooperation and Development (OECD) supported Ethical OS Toolkit²¹, which identifies potential social harm and develops strategies for ethical action.

By 2025, the market analysis shows that ethical issues have become a key part of the intelligent economy. In this evolving landscape, demonstrating that a system is safe, transparent, and fair is just as important as how well the system predicts or performs. Companies are focusing on building trust by showing their AI systems follow ethical standards, which helps them stay competitive and meet new regulations. As AI becomes more used in several areas, making sure systems are ethical will be essential for success and acceptance.

In this competitive field, ALFIE stands out because it is open to everyone. Anyone can participate in the ETD-Hub, and their opinions on ethics can help shape ethical AI models. The second part of ALFIE, the AutoML, uses the ideas from the ETD-Hub to create these ethical AI models, including features like explainability, bias detection, and support for multiple languages, making it a good fit for the EU market.

Current research²² indicates that by the end of 2025, ethical considerations in AI will no longer be viewed as optional, but they will be an essential part of AI development and deployment. This shift makes the ALFIE's Key Exploitable Assets crucial, as this highly regulated environment develops, helping businesses to stay compliant, competitive, and ethically aligned in the evolving AI landscape.

3.3 KEAs market analysis

The market analysis for each project's Key Exploitable Asset (KEA) provides comprehensive insights into the competitive landscape, emerging industry trends, and relevant regulatory considerations. This analysis helps guide strategic decision-making, identify market

¹⁸ <https://www.ibm.com/docs/en/software-hub/5.1.x?topic=services-watson-openscale>

¹⁹ <https://research.google/blog/the-what-if-tool-code-free-probing-of-machine-learning-models/>

²⁰ https://www.tensorflow.org/responsible_ai

²¹ <https://oecd-opsi.org/toolkits/ethical-os-toolkit/>

²² <https://www.deloitte.com/us/en/insights/topics/technology-management/tech-trends/2026/ai-future-it-function.html>

opportunities, and shape effective commercialisation strategies. By leveraging these insights, we aim to position each KEA optimally within its target market, ensuring it aligns with customer needs, complies with evolving regulations, and maximizes its market potential. Ultimately, this approach supports the sustainable growth of each asset and enhances its long-term value in a dynamic and competitive environment.

The European AI market is witnessing unprecedented growth accompanied by significant regulatory shifts, notably the European AI Act, which demands higher standards of transparency, fairness, and accountability. Within this landscape, the **EthiTech Dialogue Hub (ETD-Hub)**, offered by DAS, occupies a strategic niche as an essential tool for navigating the widening gap between technological advancement and ethical alignment. Organisations today face mounting pressure to move beyond simple model automation towards end-to-end ecosystems that integrate human-led policy and deliberative data directly into the development lifecycle. The ETD-Hub addresses this market need by facilitating structured consultations among policymakers, industry leaders, and citizens, thereby transforming raw debate into actionable knowledge for responsible AI governance. Unlike generic discussion forums or standard survey tools that provide only informal debate or quantitative snapshots, the ETD-Hub offers a specialised platform tuned to the specific taxonomies and legal requirements of the AI sector.

Beyond AI, the ETD-Hub has potential application in other market segments. It can be effectively adapted for deliberation on other sensitive emerging technologies such as data-driven health, genomics, biometrics, quantum technologies, and neurotechnology. Additionally, its modular architecture allows for use in general digital policy and platform governance, urban and regional planning for smart cities, and corporate governance or Environmental, Social, and Governance (ESG) initiatives focused on responsible technology use. The platform also serves as a valuable tool for education and capacity-building, providing a sandbox for students and professionals to simulate participatory technology assessments and ethical decision-making processes. The existing competition for the ETD-Hub primarily consists of indirect players that lack the structured, policy-oriented deliberation required for a robust AI ethics dataset. These include generic online discussion platforms like Quora²³, Reddit²⁴, and Stack Overflow²⁵, which facilitate large-scale debate but do not offer domain-specific analytics or integrated governance workflows. Further competition comes from enterprise survey tools such as Qualtrics²⁶ and Survey Monkey²⁷, as well as civic engagement platforms like CONSUL Democracy²⁸, Decidim²⁹, and Go Vocal³⁰, though these are typically focused on quantitative snapshots or general democratic processes rather than specific

²³ <https://www.quora.com/>

²⁴ <https://www.reddit.com/>

²⁵ <https://stackoverflow.com/>

²⁶ <https://www.qualtrics.com/>

²⁷ <https://www.surveymonkey.com/>

²⁸ <https://consuldemocracy.org/>

²⁹ <https://decidim.org/>

³⁰ <https://www.govocal.com/>

technological ethics. Specialised argument-mapping tools like Kialo³¹, Polis³², and Your Priorities³³ represent the closest technical competitors, yet they lack the deep integration with AI development frameworks and European regulatory standards found in the ETD-Hub.

The ETD-Hub possesses significant advantages against competition through its domain specialisation in AI ethics, its strict alignment with EU AI governance/policy, and its unique integration with the AutoML platform. Unlike standard tools ETD-Hub, through SemKG and AutoML, offers the potential to translate deliberative insights into concrete design constraints and evaluation criteria for AI developers; acting as an ethics advisory layer that maps stakeholder expectations to technical configuration. The ETD-Hub can also outmanoeuvre competition by leveraging the considerable academic and policy credibility of the consortium to secure early adoption and by demonstrating impact through high-quality pilot case studies. Furthermore, by providing built-in multilingual support and focusing on a clear niche as the reference platform for European AI deliberation, the ETD-Hub avoids direct competition with broad social platforms while offering a flexible licensing model that meets the diverse needs of public and private sector clients.

Despite the rapid expansion of the EU AI market, organisations face significant obstacles when attempting to scale their AI initiatives, citing challenges ranging from fragmented data infrastructure and interoperability issues to acute skill shortages and governance complexities. The **ALFIE AutoML platform**, presented by CTL with the contribution of all consortium partners, addresses these market gaps by democratising access to advanced machine learning, enabling domain experts, SMEs, and public authorities to design and build trustworthy, unbiased, and ethically compliant AI models through natural language dialogue and speech recognition without requiring deep technical or coding expertise. By integrating sophisticated large language models (LLMs) with a robust engine, the platform addresses the urgent need for accessible tools that align with human-centric values and emerging regulatory standards.

Regarding the competitive environment, existing competition includes standard automated machine learning libraries and established platforms such as Auto-sklearn³⁴, GAMA³⁵, H2O³⁶, and FLAML³⁷. While these tools offer varying degrees of automation, they often function as “black boxes” that lack integrated ethical dialogue features or specific alignment with European values. The ALFIE AutoML platform seeks to expand its reach beyond traditional boundaries. Potential applications include education, manufacturing, energy, logistics, banking, and environmental monitoring. This extends to specialised areas like fairness auditing in finance, healthcare, human resources, smart cities, public-sector decision support, and educational analytics. This breadth of application ensures that the platform remains relevant to any industry requiring compliant data management and reproducible AI pipelines.

³¹ <https://www.kialo-edu.com/>

³² <https://pol.is/home2>

³³ <https://www.yrpri.org/>

³⁴ <https://www.automl.org/automl-for-x/tabular-data/auto-sklearn/>

³⁵ <https://pgijsbers.github.io/gama/master/>

³⁶ <https://h2o.ai/platform/h2o-automl/>

³⁷ <https://microsoft.github.io/FLAML/>

The platform's primary competitive advantage lies in its "ethics-by-design" architecture, which utilizes a bi-directional feedback loop with the ETD-Hub to translate human-led policy deliberations into concrete technical constraints and bias mitigation strategies. This unique capability allows for the creation of "fairness-aware" models that are rigorously aligned with the European AI Act and GDPR, making the platform especially suitable for high-risk sectors. The platform offers unique multilingual interaction capabilities and integrated explainable AI features that deliver transparent insights into model decision-making. Furthermore, the platform prioritises environmental sustainability by considering energy consumption addressing the concern over the carbon footprint of AI training. These features collectively foster stakeholder trust and user acceptance in sensitive, high-risk sectors that conventional "black-box" tools often neglect.

Entrance to the market will be supported through open-source releases, technical workshops, and full-scale demonstrations in industrial environments. By focusing on regulated industries with high compliance needs, the platform will leverage its alignment with the European AI Act to provide a level of regulatory assurance that competitors currently cannot match.

The **AutoML Data Warehouse (AutoDW)**, offered by CERTH, serves as a central data hub designed for AI developers, data scientists, and public administrations that require structured data governance. Beyond the specific pilot use cases in autonomous vehicles, accessibility, and industrial compliance, there is significant potential application in other market segments. Specifically, any industry such as healthcare, finance, or smart cities that requires compliant data management, reproducible AI pipelines, or scalable orchestration of analytic services can adopt the technology as a foundational infrastructure component. When considering the existing competition, the technology faces indirect competition from major cloud storage providers and niche experiment tracking platforms. This includes storage solutions like Amazon S3³⁸, Google Cloud Storage³⁹, Azure Blob Storage⁴⁰, and MinIO⁴¹, alongside AI experiment tracking and model versioning platforms such as MLflow⁴², Weights & Biases⁴³, and DVC⁴⁴. Furthermore, distributed communication technologies like Apache Kafka⁴⁵ or Pulsar⁴⁶ represent another competitive layer. However, these established services often operate in isolation, focusing solely on object storage or event streaming without providing a unified environment for the entire AI lifecycle. The AutoDW offers several distinct advantages against competition by integrating workflow coordination and AI experiment tracking directly into its multi-database backend. While large-scale cloud providers offer nearly unlimited scalability, they often lack native workflow coordination and specialized metadata governance for AI assets. In contrast, the AutoDW provides unified access to heterogeneous data types while ensuring full traceability of experiments, model versions, and service outputs. This event-

³⁸ <https://aws.amazon.com/s3/>

³⁹ <https://cloud.google.com/storage>

⁴⁰ <https://azure.microsoft.com/en-us/products/storage/blobs>

⁴¹ <https://www.min.io/>

⁴² <https://mlflow.org/>

⁴³ <https://wandb.ai/site/>

⁴⁴ <https://www.dvc.org/>

⁴⁵ <https://kafka.apache.org/>

⁴⁶ <https://pulsar.apache.org/>

driven, message-broker-driven orchestration effectively reduces the coupling between services, which enables system operators to maintain more scalable and resilient AI pipelines than those offered by disconnected traditional storage tools.

The strategy to outmaneuver the competition centers on delivering a unified, end-to-end operations backbone that emphasizes transparent, auditable, and ethically designed AI governance. The AutoDW will be optimized for on-premise, hybrid, and sovereign AI deployments, enabling organisations to retain greater control over sensitive data and reduce dependency on public cloud hyperscalers. By minimizing vendor lock-in and encouraging extensibility, the system integrates seamlessly with a multi-agent intelligence layer to deliver high technical value, build trust, and ensure regulatory compliance. This positioning establishes the solution as a vital infrastructure component for responsible AI development across the global market. To support market entry, the plan involves promoting adoption through open-source release, comprehensive trainings, workshops, fairness assessment templates, pilot demonstrators, and dedicated integration support for both public and private sector users, ensuring wide accessibility and ease of deployment.

The Semantic knowledge graph (SemKG) of ethical data, led by EHU with the contribution of CTL, addresses the market demand for responsible AI development by providing a machine-operable layer of awareness that ensures AI systems are fair, transparent, and aligned with societal values. This technology is particularly valuable for Information and Communication Technology (ICT) companies that require advanced knowledge management solutions, such as semantic search, indexing, and filtering within complex datasets.

The core value proposition of the SemKG lies in its ability to convert unstructured textual data, such as extensive forum discussions into usable and condensed knowledge representations. While its current application focuses on AI ethics in the ETD-Hub, the potential application in other market segments further than the ones specified in the project is substantial. The SemKG is designed as an independent processing component that can be adapted to apply to any scientific or technical domain where textual data is available for processing and matching ontologies exist to cover the specific domain. This flexibility allows it to potentially serve diverse sectors such as healthcare and finance by populating domain-specific knowledge graphs from raw text to support complex reasoning and decision-making.

When examining the existing competition, related work encompasses several established methods for knowledge representation, though these typically focus on wider contexts rather than the specific domain of AI ethics. Current state-of-the-art tools include graph databases like Neo4j⁴⁷, ArangoDB⁴⁸, and JanusGraph⁴⁹, which capture interconnectedness through nodes and edges but often face significant limitations regarding horizontal scalability and data consistency when dealing with massive datasets. Furthermore, while AI-driven natural language processing libraries like spaCy⁵⁰ and Stanford NLP⁵¹ facilitate entity extraction, they often function as foundational building blocks rather than integrated text-to-knowledge-graph systems tailored for ethical reasoning. ALFIE's solution distinguishes itself by integrating these

⁴⁷ <https://neo4j.com/>

⁴⁸ <https://arango.ai/>

⁴⁹ <https://janusgraph.org/>

⁵⁰ <https://spacy.io/>

⁵¹ <https://nlp.stanford.edu/>

natural language processing components with a suite of ontologies specifically focused on AI governance, risk, and ethics.

The primary advantages against competition stem from the fact that this is one of the first text-to-knowledge-graph software systems specifically engineered to map ethical discourse to well-defined knowledge sources. By utilising multiple domain-specific ontologies, it provides a structured foundation for tasks that require deep reasoning and inference, such as providing reliable context for large language models through retrieval-augmented generation. To maintain this edge, the plan for how you plan to outmanoeuvre the competition involves conducting extensive comparative evaluations against existing alternatives in the literature to validate effectiveness and build trust. Additionally, the entrance to the market will be supported by making the software available for download and use, thereby fostering early adoption and establishing the project as a reference point for global ethical AI standards.

The ALFIE **Agentic Core**, presented by KInIT, is a multilingual, agentic service that enables diverse users to create trustworthy AI through interaction in natural language and speech. This technology functions as the intelligent backend for the ALFIE's web interface, utilizing a multi-agent framework to guide users through the complete machine learning workflow, from initial data analysis to final model deployment. By automating complex technical steps such as algorithm selection and hyperparameter tuning, the platform democratizes access to AI development while ensuring that AI models are unbiased and compliant with the European regulatory framework. Potential application in other markets includes sectors such as education, state administration. The landscape of existing competition includes various established automated machine learning platforms and research frameworks that offer differing levels of process automation. The AutoML-Agent⁵² serves as a research-focused multi-agent framework for natural language interfaces, yet it lacks the integrated ethics and bias mitigation modules central to the ALFIE project. Mature enterprise platforms such as H2O.ai Driverless AI⁵³ and DataRobot⁵⁴ provide robust lifecycle management and explainability, though they primarily operate through traditional technical dashboards rather than interactive conversational AI. Unlike many commercial black-box tools, the Agentic Core offers unique multilingual and multimodal capabilities, supporting natural language interaction in multiple European languages to enhance accessibility for non-experts. Its integrated explainable AI (XAI) module provides transparent reasoning for model decisions, which fosters greater stakeholder trust than proprietary alternatives. To successfully penetrate the market, the plan is to outmanoeuvre the competition by emphasising its key differentiators as an ethics-first, multi-agent orchestration service with a natural-language interface. By leveraging multilingual support, the platform can capture non-English speaking markets that are currently underserved by established AutoML tools.

To facilitate successful entry into the market, the approach includes early stakeholder validation by engaging with key public-sector partners, demonstrating the application's capabilities through containerisation, packaging it as a comprehensive, well-documented Docker image to ensure seamless deployment and scalability.

⁵² <https://deepauto-ai.github.io/automl-agent/>

⁵³ <https://h2o.ai/resources/product-brief/h2o-driverless-ai/>

⁵⁴ <https://www.ibm.com/products/watsonx-orchestrate>

The project aligns with academic research objectives, offering an innovative system designed to advance understanding in several key areas. The core goal is to develop an ethical, agentic system tailored for multilingual environments, with specific research focused on semantic alignment among legal, ethical, and machine learning (ML) domain language, as well as investigating performance degradation issues within multilingual settings. It encompasses multiple components, including a multi-agent system architecture, utilisation of multilingual models to support diverse languages, an evaluation framework for assessing individual agents and overall system performance, and integration of Retrieval-Augmented Generation (RAG) supporting a Semantic Knowledge Graph. It includes a fast API-based agentic service, a PostgreSQL database for persistent storage, a multi-agent system comprising specialized agents responsible for ML workflows, conversation management, RAG operations, and ethics and compliance. Additionally, it features an AG-UI interface for effective communication and state management, connectors enabling integration with data warehouses, task management systems, and large language model (LLM) service providers. An evaluation framework is incorporated to monitor system performance, and all components are containerized within a Docker image to ensure ease of deployment, reproducibility, and scalability. The Agentic Core offers significant value by enabling advanced research into semantic alignment across various domains such as compliance, ethics, and ML, in multiple languages, and understanding and mitigating performance degradation in multilingual environments. The modular design also allows the service to be adapted for other research initiatives, including workflow management, RAG-based applications, background task execution, human-in-the-loop processes, and multimodal data handling. To reach target audiences, it is intended to publish scholarly articles and technical papers, present findings at seminars, workshops, and conferences, foster collaboration through targeted outreach, organize panel discussions to facilitate knowledge exchange, and share updates via social media platforms. The primary audiences include researchers in machine learning, ethics experts, and domain specialists interested in multilingual AI systems and ethical AI development.

The **ALFIE web interface with XAI features**, offered by CTL, KInIT and CERTH, acts as a user-friendly gateway for non-experts to engage with machine learning workflows. By utilizing natural language and speech interaction powered by large language models, the platform effectively removes technical barriers, enabling a wider range of users to create and manage advanced AI models, targeting the growing demand for transparent and ethically grounded development tools. While current competitors include traditional machine learning platforms and proprietary AI governance tools like Fairlearn⁵⁵, AIF360⁵⁶, and the Google What-If Tool⁵⁷, many existing solutions lack integrated ethical features or fail to offer multilingual support. The ALFIE interface addresses this gap by functioning as a centralised gateway that democratises access to AI for non-experts and domain specialists in high-stakes sectors. The interface is applicable to education, industry, IT, and government sectors. By removing technical barriers through an intuitive interaction layer powered by large language models, the platform positions itself as a uniquely inclusive alternative to technical-heavy or "black-box" AutoML tools that often overlook the societal and legal implications of model deployment. The integration of a bias-aware explainable AI module within the ALFIE web interface addresses these

⁵⁵ <https://fairlearn.org/>

⁵⁶ <https://aif360.readthedocs.io/en/stable/index.html>

⁵⁷ <https://pair-code.github.io/what-if-tool/>

shortcomings by offering explanations, fairness metrics, and data-level interpretability. This ensures that every model generated is transparent and trustworthy, allowing users to understand the internal logic behind predictions and identify potential risks before deployment.

A fundamental differentiator for the web interface is its rigorous alignment with the EU AI Act and GDPR, offering a level of regulatory assurance that many commercial alternatives currently lack. The system's primary competitive advantage lies in its bi-directional feedback loop with the ETD-Hub, which ensures that technical model generation is continuously refined by stakeholder deliberations on societal values. Its capabilities, allowing for natural language and speech interaction in multiple European languages, significantly expand its serviceable market beyond English-centric technical communities. Furthermore, the integration of explainability (XAI) features and knowledge-graph enrichment provides an auditable decision-making process suitable for fairness auditing in finance, healthcare, HR tech, smart cities, public-sector decision support, and education. As the platform transitions from a research prototype to a commercial offering, it is well-positioned to scale into global markets that advocate for trustworthy AI principles.

The approach to facilitating market adoption includes a multi-faceted engagement plan comprising open-source software releases, comprehensive training sessions, hands-on workshops, and the provision of fairness assessment templates. Additionally, the plan involves developing pilot demonstrators to showcase practical applications and offer dedicated integration support to both public institutions and private sector adopters, aiming to lower barriers to adoption and foster trust, and offering cutting-edge academic research focused on enhancing trustworthy AI within the ALFIE framework. The research concentrates on advancing methodologies for data bias detection, developing robust fairness metrics, ensuring model accountability, and creating explainability methods that support transparent decision-making processes. The suite of tools and methodologies is designed to identify and mitigate biases present in datasets, detect unfair behaviors in AI models, and generate clear, human-understandable explanations. The offering encompasses the development of standardized evaluation protocols, fairness indicators, interpretability workflows, and cognitive-grounded explanation models tailored to meet the needs of diverse human users. This work bridges theoretical advancements with practical assessment, enabling trustworthy AI implementation in immersive ALFIE contexts, delivering innovative insights and practical tools for bias detection and transparent model evaluation. The solutions support ethical standards, regulatory compliance, and human-centric AI principles, particularly for high-risk applications. The aim is to reinforce Europe's leadership in responsible AI development. To maximize impact, the plan involves disseminating through multiple channels, including publication of research papers in academic journals and conferences, releasing open-source codebases for community use, organizing workshops focused on bias and fairness in AI, contributing to AI Act standardisation efforts, developing educational materials, and collaborating with academic, industry, and regulatory stakeholders to foster dialogue and adoption.

3.4 SWOT analysis

A SWOT analysis offers a thorough and strategic evaluation of each project's Key Exploitable Assets (KEAs). It provides a detailed overview of the internal strengths and weaknesses, as well as external opportunities and threats that could influence the success and development of ALFIE's products and technologies. By conducting this analysis, we gain valuable insights into the challenges ALFIE's KEAs may face, as well as their competitive advantages in the

marketplace. Furthermore, examining external threats and opportunities enables us to anticipate market trends, identify potential risks, and understand competitive pressures, allowing for proactive strategy formulation. Below, the SWOT analysis for each KEA is presented, offering critical insights into strategic positioning, growth prospects, and long-term potential within the evolving industry landscape.

The initial SWOT analysis of the ALFIE project was provided in deliverable D5.1 (M4). The tables that follow present the detailed SWOT analysis for each Key Exploitable Asset (KEA).

Table 11. SWOT analysis for ETD-Hub

Strengths	Strong domain specialisation: Purpose built for AI governance and ethics offering direct relevance to organisations grappling with the AI Act and trustworthy AI expectations. Integrated deliberation: Combines forums and datasets into one pipeline, a differentiator versus participation workflow only competitor platforms. Academic and policy credibility of consortium: University and research partners plus industrial players lend legitimacy and access to EU and national policy circles, improving trust, pilot opportunities, and impact pathways. Ethics and compliance by design: Strong emphasis on GDPR, responsible data handling, accessibility, and responsible analytics positions ETD-Hub as a low-risk option for public administrations and compliance-conscious enterprises. Connection to AutoML: Translation of deliberative insights into constraints, objectives, and documentation for AutoML/model development via SemKG offers a unique ethics advisory layer that most competitors lack, helping developers operationalise ethics-by-design. Concrete high-value use cases: Three pilot use cases generate real datasets on ethics and acceptability of specific AI systems.
Weaknesses	Early-stage maturity and limited track record: At project end ETD-Hub will likely have only a small number of pilots, compared to mature civic platforms and survey providers with thousands of deployments and well-known reliability. Niche and more complex value proposition: AI and AI ethics deliberation with analytics and AutoML integration but may be harder to explain to non-expert users than simple narratives like run online consultations or create surveys. Higher onboarding and facilitation requirements: Effective use depends on well-trained moderators, facilitators, and policymakers to translate data into action; organisations lacking internal capacity may perceive the platform as resource-intensive versus plug-and-play survey tools. Dependence on consortium commitment for sustainability: Early success hinges on partner collaboration in moderation, AutoML integration, and accessibility auditing. If this is not efficient, we introduce vendor risk. Limited brand recognition: As a new entrant, ETD-Hub lacks the name recognition of established civic tech or research tools, increasing marketing costs and lengthening commercialisation cycles. Limited utility outside AI at present: While technically adaptable to other domains; many features, taxonomies, and narratives are currently only optimised for AI ethics; this will reduce immediate appeal for organisations wanting a single system for all participation topics. Narrower thematic focus: being optimised for AI and AI ethics will limit immediate appeal for organisations seeking a generic participation platform, unless domain

	adaptations are implemented. This limits the initial appeal of ETD-Hub to those outside of its target market. Potentially higher complexity: richer data may require more training and onboarding to extract maximum value compared to simple survey tools or social media. Resource requirements: robust moderation, ethical oversight, and responsible AI analytics may require more organisational resources than “lightweight” engagement tools like surveys.
Opportunities	Regulatory tailwinds: The AI Act will require organisations to conduct risk management, stakeholder engagement, and documentation for high-risk AI systems. ETD-Hub can position itself as a practical tool for compliance. Growing market for responsible AI: Enterprises and public bodies are actively seeking tools for AI governance, ethics impact assessment, and monitoring. ETD-Hub could support model lifecycle tooling and AutoML platforms to become part of dev workflows. First-mover advantage in AI ethics deliberation as a service: Few platforms combine deliberation and sentiment/recommendation datasets for AI. ETD-Hub can establish itself as a reference for best practice in participatory AI governance. Expansion into adjacent domains: DAS’ bespoke architecture can be reused for other sensitive technologies - enabling domain-specific spin-offs with relatively modest reconfiguration. Partnerships with research communities: Collaborations with AI research centres can open new channels for deployment and co-branding. Education and capacity-building market: Universities, professional associations, and training providers need practical tools to teach AI ethics and governance. ETD-Hub can be offered as a teaching sandbox with associated curricula, certifications, and workshops.
Threats	Intensifying competition from generalist and survey platforms: Established tools (Decidim, CONSUL, GoVocal, Qualtrics, SurveyMonkey, Pol.is, Kialo, etc.) may add AI ethics templates, basic sentiment analysis, or compliance features, reducing ETD-Hub’s perceived differentiation. Fragmentation and tool overload: Organisations are already struggling with many tools (risk registers, DPIA tools, fairness dashboards, MLOps, etc.). ETD-Hub risks being seen as one more tool unless integration with existing stacks is very clear. Public sector and academia budget constraints: If ETD-Hub charges for access to the platform or data, economic pressure can delay or reduce adoption, especially if decision-makers believe that existing survey or consultation tools are sufficient. Evolving regulatory and ethical expectations: As AI governance norms evolve, ETD-Hub must keep its taxonomies, analytics, and workflows aligned. Failure to keep up could quickly make features outdated or misaligned with compliance needs. Data protection and trust risks: Any incident or public concern related to privacy, data misuse, or biased analytics could damage trust in ETD-Hub; especially given its focus on ethical AI. This is a reputational risk that must be actively mitigated via strong governance and transparency. Internalisation by large vendors: Big cloud and AI providers (or major survey/civic-tech vendors) might in future embed stakeholder engagement and

	ethics features directly into their AI platforms, leveraging scale, brand, and integration advantages to dominate the market and shut-out ETD-Hub. Brand recognition: existing platforms may benefit from stronger name recognition.
--	---

Table 12. SWOT analysis for AutoML platform

Strengths	Technological integration: The platform combines an AI interaction layer with an advanced AutoML engine, allowing non-experts to build complex models via natural language or speech. Ethical and legal alignment: The platform integrates bias detection, mitigation, and explainable AI (XAI) features to ensure all outputs respect the European legal framework, including the AI Act and GDPR. Inclusivity and accessibility: The system supports multilingual interactions and provides accessible interfaces tailored for diverse user groups, including those with limited technical proficiency.
Weaknesses	Data dependency: The effectiveness of the platform relies heavily on the quality and representativeness of training data, meaning biased or incomplete datasets can still impact model reliability. Computational and resource costs: Training large-scale models requires significant computational power and financial investment, which can be a barrier for some applications. High computational cost of training complex models from scratch, which may lead to resource constraints if not leveraging transfer learning. Technical complexity: Orchestrating a diverse microservices architecture that handles NLP, XAI, and bias detectors involves significant implementation challenges and potential for software bugs. Hardware and performance constraints: Converting high-performance models for deployment on mobile devices (e.g., TFLite) can result in a drop in accuracy or slower inference times. Format restrictions: The system currently has limitations on supported data formats, primarily focusing on CSV.
Opportunities	Rapidly growing market: The European AI market is growing, offering substantial growth potential for ethical AI solutions. Regulatory demand: Increasing global and EU focus on trustworthy AI creates a niche for platforms that automate compliance with emerging legal standards. Democratisation of technology: There is a significant opportunity to empower SMEs and non-technical organisations by reducing the barriers to advanced AI development. Strategic synergies: Collaborating with other EU initiatives (e.g., AI on Demand platform) can enhance the platform's data access and metadata richness.
Threats	Regulatory and legislative lag: The slow pace of legislative processes may cause a lag between technological breakthroughs and current legal frameworks, potentially making recommendations less timely. Technological fragmentation: Differing AI policies across EU member states can hinder the harmonisation of legal frameworks, creating a fragmented market environment.

	Public mistrust: A general lack of public awareness or understanding regarding AI complexities can lead to resistance against AI adoption in daily life. Market pressure: Powerful lobbying from major AI industry entities may prioritise commercial interests over ALFIE's human-centric and ethical objectives. Political and economic instability: Shifts in governmental leadership or funding priorities within member states could impact the long-term sustainability of the project.
--	---

Table 13. SWOT analysis for Data warehouse (AutoDW) of ethical data

Strengths	- Integrated AI lifecycle management - Event-driven architecture - Multi-database flexibility - On-premise and hybrid deployment - Ethical AI and governance by design - Vendor-neutral, extensible design
Weaknesses	- Limited scalability compared to hyperscale cloud platforms: AutoDW provides a unified, purpose-built solution for the ALFIE platform, it may not yet match the scalability, automation, or operational robustness offered by established hyperscale competitors. - Large-scale cloud storage platforms offer virtually unlimited scalability, global distribution, and highly optimized data transfer speeds, capabilities that AutoDW cannot guarantee, especially for high-volume data ingestion over the public internet. - Operational overhead - Performance constraints over public internet - Dependency on ALFIE ecosystem - Lower ecosystem maturity - AutoDW's event-driven coordination relies on co-deployment of other ALFIE services and requires more hands-on system management.
Opportunities	- Growing global demand for transparent, ethical, and regulated AI - Rising need for on-premise, secure AI infrastructure - Partnership with academic and research institutions - Integration with federated learning and privacy-preserving AI
Threats	- Dominance of cloud vendors - Competition from MLOps platforms (e.g. MLflow, Kubeflow, and Databricks) - Infrastructure complexity for small adopters - Rapid changes in AI/ML tooling ecosystems

Table 14. SWOT analysis for Semantic knowledge graph of ethical data

Strengths	Pioneering technology: It is one of the first text-to-knowledge-graph software systems specifically engineered to map ethical discourse to well-defined knowledge sources.
------------------	--

	Foundational ethical layer: The graph acts as a machine-operable layer of ethical awareness, ensuring that ethical considerations are a foundational part of AI design rather than an afterthought. Structured knowledge representation: It transforms unstructured textual information from stakeholder debates into highly structured, machine-readable semantic triplets. Adaptability: The software is designed as an independent processing component that can be adapted to any scientific or technical domain, provided that matching ontologies are available. Advanced reasoning: By utilising domain-specific ontologies, it provides a foundation for tasks requiring deep reasoning and inference, such as guiding automated AI model development.
Weaknesses	Early development stage: The technology is currently a work in progress and has not yet been evaluated extensively. Lack of comparative validation: Due to its early stage, it has not been compared against existing alternatives in the literature to validate its effectiveness. Scalability challenges: Some underlying graph database technologies may face limitations regarding horizontal scalability and query performance when handling massive datasets. Dependency on data quality: The system relies heavily on the quality of raw input data; deficiencies or inconsistencies in the data could significantly impact the reliability of the resulting graph.
Opportunities	High market demand: There is a significant need in the EU for secure AI models and trustworthy platforms that improve ethics in AI applications. Regulatory compliance: New legislation, such as the EU AI Act, creates a strong market demand for tools that ensure AI systems are ethically compliant. LLM integration: The graph can be used to provide reliable, externally verified context for LLMs through RAG. Democratisation of expertise: By integrating ethical reasoning into automated workflows, the platform can democratise access to machine learning expertise for non-experts.
Threats	Rapid technological evolution: The fast-paced evolution of AI technology and ethical standards could affect the platform's long-term market position. Intellectual property risks: Managing IPR within a consortium of both public and private entities might lead to disputes that undermine commercial exploitation. Data and representation bias: Skewed or unrepresentative data sources can introduce systematic distortions into the graph, potentially perpetuating societal inequalities or stereotypes. Privacy and security concerns: Integrating sensitive behavioral or demographic data into the graph raises critical ethical implications regarding individual privacy rights and potential data breaches. Existing competition: While specific to AI ethics, there exists significant competition from broader knowledge representation methods and established graph database providers.

Table 15. SWOT analysis for Agentic Core

Strengths	- Standardized AG-UI interface - Integrated ethics and compliance - Multilingual support - End-to-end service orchestration
Weaknesses	- High computational cost of LLM-oriented architecture - Reliance on LLM accuracy and stability - Limited enterprise maturity - Potential user trust issues - ALFIE Agentic Core relies on worker services to support bias detection, AutoML, XAI, functionalities and is tightly-coupled with ALFIE AutoDW service. - The LLM-based user interface and workflow management may be more costly and introduce security issues like prompt injection. - Competitors are mature, enterprise-grade platforms with proven track records in production, robust deployment options, interpretability, and monitoring.
Opportunities	- Rising demand for responsible, explainable, and ethical AI - Growing adoption of conversational and agentic AI interfaces
Threats	- Rapid evolution of LLMs and agents - Regulatory restrictions on LLM use - Integration and vendor lock-in issues

Table 16. SWOT analysis for web interface with XAI features

Strengths	Unique market positioning: The interface resides within a unique niche in the European market for ethically compliant AI solutions. User-centred and inclusive design: The interface is developed using UCD and inclusive design principles, ensuring that the needs of diverse users are prioritised. Natural language interaction: Users can engage through natural language and speech, making the creation of AI models accessible to non-experts. Multilingual support: The platform is designed to be multilingual, initially supporting English and Spanish before expanding to other European languages. Accessibility standards: The interface targets compliance with accessibility standards. Robust security and authentication: By employing Keycloak as a dedicated Identity and Access Management service, the interface supports secure data handling in line with GDPR. Centralised resource management: Provides a unified dashboard for managing files, datasets, and models, and visualise progress in real-time. Explainability (XAI): A core strength is the integration of Explainable AI features, which provide transparent reasoning behind model decisions to build user trust, bridging a gap not addressed by most XAI tools.
------------------	--

	Regulatory transparency and auditability: Fully transparent, non-proprietary, and easily auditable (aligned with EU AI Act). Adaptive human-centric communication: Human-centric explanations suitable for non-technical users. Modular cross-domain adaptability: Adaptable across domains and datasets due to modular design.
Weaknesses	Implementation complexity: Orchestrating a microservices architecture that integrates NLP, XAI, and bias detector through a single interface presents significant technical challenges and potential for bugs. Simulating complex AI models in real-time may introduce inference latency. Data quality dependencies: Its effectiveness is heavily dependent on the availability of unbiased, high-quality data, which can be difficult to acquire. Bias detection quality depends on data completeness and attribute availability. Requires domain definitions of fairness and depends on data availability. Proof-of-Concept scale: Research-level maturity; not yet optimized for large-scale industry workloads. At its current stage, the deployment may limit the initial representativeness. Absence of universal standards and domain-specific fairness constraints: Requires domain definitions of fairness (no universal fairness metric). Emerging market presence and challenges to commercial acceptance: Limited visibility compared to commercial AI governance tools. Computational cost: The computational cost of the underlying LLMs requires robust infrastructure.
Opportunities	Regulatory momentum: Growing demand for fairness-aware AI tools motivated by the EU AI Act. Global shift toward accountability: There is a rising global trend emphasizing the need for more transparent and fair AI systems, positioning ALFIE to address these widespread concerns. Increased STEM investment: Strategic focus on AI and technology skills worldwide provides a growing user base capable of leveraging these platforms. Mainstreaming algorithmic accountability and bias auditing: Increasing adoption of bias audits in public-sector and industry processes. Strategic integration into AI governance and open ecosystems: Opportunity to integrate into broader AI governance frameworks or open-source ecosystems. Leadership in standardising fairness-aware explainability: Potential to set standards for fairness-aware explanation schemas in European projects.
Threats	Competitive pressure: There is a risk from external markets or companies that may prioritise rapid innovation over ethical considerations. Rapid evolution of commercial AI governance platforms with larger R&D budgets. Lobbying interests: Powerful industry lobbying could influence policies in ways that compromise the human-centric goals of the ALFIE ecosystem. Inexperienced misjudgment and misapplication: Misinterpretation or misuse of fairness outputs by inexperienced users. Regulatory fluidity and legislative lag: Regulatory changes may introduce new requirements beyond current ALFIE scope. Exclusive UX imbalance and market pressure: Competitive pressure from closed-source XAI platforms with polished UX.

3.5 PESTEL analysis

PESTEL analysis serves as essential strategic tools that facilitate the identification, evaluation, and mitigation of external factors which could potentially impact a project's success. By examining these factors, organisations can better understand the broader environment in which they operate and proactively address challenges. Specifically, the PESTEL framework categorizes these external factors into six key areas: **Political** (P), which includes government policies and regulatory changes; **Economic** (E), covering market conditions, inflation rates, and economic growth trends; **Societal** (S), reflecting cultural, demographic, and social attitudes; **Technological** (T), encompassing innovations and technological advancements; **Environmental** (E), considering ecological and sustainability issues; and **Legal** (L), involving laws, regulations, and compliance requirements. Since these external factors directly affect operational decisions, market positioning, and strategic planning, early identification and thorough analysis of these elements equip project teams with vital insights. This enables more effective risk management, informed decision-making, and the development of resilient strategies, thereby increasing the likelihood of successful project execution in a dynamic environment.

The initial PESTEL analysis of the ALFIE project was documented in deliverable D5.1 (M4). Building upon this foundation, the current deliverable provides a comprehensive PESTEL analysis for each Key Exploitable Asset (KEA). These analyses are included in the Appendices (A-F), offering an in-depth assessment of the external factors influencing each KEA. In this section, an overview of these findings is presented, highlighting key insights and implications for the project's strategic planning and risk management, aiming to support proactive decision-making and ensure that each KEA is effectively positioned to adapt to external changes, enhancing the overall success and sustainability of the ALFIE project.

Political factors within the EU are primarily characterized by significant policy fragmentation across member states, where differing governance priorities, regulatory frameworks, and national interests can lead to uneven demand for AI solutions and pose challenges to the harmonisation of legal and operational standards. This fragmentation can create barriers to cross-border collaboration and impede the seamless integration of ALFIE's objectives across different jurisdictions. Additionally, the project is susceptible to risks stemming from political instability, such as upcoming elections, governmental changes, or shifts in policy agendas, which may delay or disrupt participatory initiatives, funding allocations, and regulatory support for AI research and development. Moreover, the influence of powerful industry lobbying groups and the politicisation of AI ethics debates, often framed as a dichotomy between fostering innovation and imposing regulation, could threaten the core human-centric principles of ALFIE. These debates may lead to policy shifts that prioritize economic gains over ethical considerations, potentially undermining the project's commitment to responsible AI development. International collaboration further complicates the political landscape, as aligning KEAs with diverse and sometimes conflicting global standards and regulations requires careful navigation. Differences in data privacy laws, ethical guidelines, and regulatory approval processes across countries can introduce delays, increase compliance costs, and pose risks to the scalability and interoperability of the project's solutions. Overall, navigating the evolving political environment requires proactive engagement with policymakers, continuous monitoring of legislative developments, and strategic alignment with international standards to safeguard the project's objectives and ensure sustainable progress.

ALFIE seeks to foster active engagement with a broad spectrum of stakeholders, including policymakers, legal experts, AI ethicists, industry representatives, and academia. The project aims to facilitate their participation in the ETD-Hub conversations, which focus on critical ethical AI policy topics. Through these collaborative discussions, ALFIE intends to identify areas of convergence and develop comprehensive recommendations that can inform and support the alignment of AI policies across different jurisdictions. It is important to emphasize that a primary objective of ALFIE is to contribute to policy recommendations tailored to the EU. These recommendations are designed to enhance the European Union's strategic position in responsible AI development, address regional ethical and legal challenges, and support the implementation of harmonized policies within member states. However, recognizing the global nature of AI advancement, the project also aspires to contribute insights and frameworks that may be adaptable or influential beyond the EU, potentially guiding international dialogue and cooperation on ethical AI governance. This focus aims to ensure that EU policies are both forward-looking and aligned with emerging global standards, ultimately fostering a cohesive and ethically grounded AI ecosystem within Europe.

Economic barriers pose significant challenges to the successful development and sustainability of the ALFIE ecosystem. These include the substantial financial resources required to establish and maintain comprehensive AI governance infrastructure, such as advanced technology, skilled personnel, ongoing research, and stakeholder engagement. Additionally, the high costs associated with implementing ethical, legal, and compliance standards can disproportionately burden small and medium-sized enterprises (SMEs) with limited resources, risking their exclusion from the ecosystem. Established platforms and commercial providers may bundle competing modules or services, creating further pressure to reduce costs and remain competitive. Additionally, the transitional costs for businesses updating their workflows to adopt ethical AI practices can impede widespread implementation. Despite these challenges, the growing market demand for ethical AI solutions suggests that ALFIE can position itself advantageously by highlighting the long-term economic benefits of ethical practices. Strategic collaborations with major companies could help shape the ethical AI market and mitigate concerns about competitiveness, ensuring that ethical AI remains economically accessible and sustainable for organisations of all sizes.

Social challenges include a general lack of public awareness and understanding of AI's features, capabilities, and legal issues. ALFIE's ETD-Hub aims to make discussions about AI more open and give citizens a chance to participate in shaping AI legislation. However, it faces problems like participation fatigue and the digital divide, which can lead to underrepresentation of marginalized groups. This is common in many tech and governance discussions, as people may lose interest or feel they lack the knowledge to contribute. Differences in digital skills and cultural views on AI ethics also make it harder to reach agreement on regulations. Public trust in AI, concerns about bias, and transparency are ongoing obstacles to wider acceptance. To address this, ALFIE has started a public engagement and education plan, including awareness campaigns, easy-to-understand materials, and events. These efforts aim to encourage more people to get involved and learn about AI.

Technological considerations involve the interoperability of diverse components, such as integrating the ETD-Hub data within AutoML pipelines (AutoDW and SemKG) to ensure a cohesive ecosystem. The complexity of AI development often necessitates advanced technical expertise, potentially limiting accessibility for non-expert users. As the user base grows,



maintaining scalability and performance without compromising security or responsiveness is a primary technical concern.

Environmental factors highlight that the training and development of intensive AI models carry a substantial carbon footprint and significant water consumption. There is increasing societal and regulatory pressure for eco-friendly IT, with public bodies and large enterprises increasingly requiring digital tools to meet strict sustainability reporting standards. Stakeholders may view additional digital platforms as environmentally burdensome, requiring assets like AutoDW and AutoML to optimise resource usage and track energy metrics for transparency.

Legal factors for ALFIE mainly involve the need to comply with complex, changing regulations like the GDPR and the EU AI Act. Non-compliance poses legal risks and could reduce the platform's effectiveness. The slow pace of legislation often means regulations lag behind technological progress, making it necessary to adopt flexible and adaptive legal approaches. A major challenge is ensuring that the project's policy recommendations are timely and relevant, which relates to the Collingridge dilemma: regulating new technology is easier early on, but its full impacts are unknown, and by the time impacts are clear, regulation may be too late because the technology is deeply embedded. To address this, ALFIE should push for agile legal frameworks that can adapt quickly and stay in close contact with legislators, while continuously monitoring technological and societal developments to make sure their policy suggestions stay impactful and appropriate.

4 Exploitation Plan

The ALFIE project is anchored by two primary objectives designed to harmonize human-centric values with advanced technological development through the creation of a collaborative forum known as the ETD-Hub and the AutoML platform. The overarching exploitation strategy is focused on establishing the project as a definitive reference point for ethical AI both within the European Union and globally. To achieve this, ALFIE proactively engages with AI, ethics, and legal professionals to foster a pool of early adopters and experts who can improve the quality of the final outcomes. The long-term success of the ALFIE ecosystem is rooted in its capacity to generate impact far beyond the boundaries of the consortium by maximizing diverse exploitation opportunities and ensuring that all project KEAs are designed for seamless scalability across various sectors including industry, healthcare, education, and government.

The project has identified six KEAs that form the cornerstone of its marketisation strategy, each offering unique value as either integrated components or stand-alone solutions. The first asset is the ETD-Hub, a multilingual digital deliberation platform designed to capture and structure public and expert opinion on AI ethics to inform regulatory standards. The second asset is the AutoML platform, which automates the machine learning process to translate ethical policy recommendations into actionable technical constraints, ensuring models are unbiased and compliant with the EU regulatory framework. Supporting these is the third asset, the AutoDW, which serves as a secure, end-to-end operational backbone for managing the entire AI lifecycle from raw data ingestion to trained model archival. The fourth asset is the SemKG, which processes textual data to map the complex relationships between abstract ethical principles and technical requirements. The fifth asset, the Agentic Core, acts as a multilingual intelligent backend that uses a multi-agent framework to guide users through ML workflows via a natural language chat interface. Finally, the sixth asset is the Web interface with explainable AI features, which provides a user-friendly gateway for non-experts to explore and manage AI models while delivering transparent insights into model decision-making.

The roadmap to market entry for these assets follows a phased approach that prioritizes the European landscape before pursuing global opportunities. Initial efforts focus on anchor pilots primarily in university settings and summer. Following the project's conclusion, ALFIE intends to launch and maintain the ETD-hub and AutoML platform to serve as a live showcase for regulators and policymakers. This will be complemented by a transition toward targeting priority segments such as EU agencies, investors, and industrial players working with European Digital Innovation Hubs (EDIHs). The commercialisation efforts are spearheaded by the consortium's SMEs, which leverage their networks and market knowledge for marketing strategies. Universities within the consortium play a vital role by conducting ongoing research and development to maintain an innovative edge and facilitate training programs for potential users. Research institutions provide continuous validation of the project's effectiveness and its adherence to ethical norms, while industrial partners like Bosch demonstrate the solutions in practical settings to build market credibility.

To ensure financial sustainability, ALFIE employs a multi-layered revenue model that combines recurring income with professional service fees. Possible monetisation mechanisms include subscription-based access, allowing customers to access core features and essential software updates that maintain alignment with evolving regulations. Beyond access fees, ALFIE aims to generate income through bespoke consulting services, such as fairness audits,

explainability assessments, and on-demand framework customisation for domain-specific ethical challenges. Revenue is further bolstered by providing paid training sessions and technical workshops to build professional competencies in AI governance. In the long term, the project may offer stable private instances for organisations seeking to conduct sensitive internal deliberations with higher security and confidentiality needs. The final phase of exploitation involves finalizing model prototypes and conducting validation tests to facilitate a smooth transition from development to market. By aligning rigorously with the EU AI Act and GDPR, ALFIE establishes a unique market position as a provider of regulatory assurance and transparent technology. This commitment to "ethics-by-design" ensures that the future of AI in Europe is deeply rooted in human rights and social equity while providing a blueprint for trustworthy AI development that can be scaled internationally. Success is measured through specific impact indicators including the number of industrial adoptions, the creation of start-ups and spin-offs. Ultimately, this strategy ensures the project's continued relevance and societal impact long after the initial research phase concludes.

4.1 Business Model Canvas

The Business Model Canvas for the ALFIE project serves as a strategic roadmap designed to enhance the impact and sustainability of its KEAs and the broader ethical AI ecosystem. The core value proposition is the democratisation of AI accessibility through a generative AI-powered interaction layer that supports natural language and speech in multiple languages. This approach allows users to communicate their goals effortlessly regardless of technical expertise, ensuring that the resulting AI models meet specific objectives while remaining strictly compliant with European legal standards and ethical AI practices.

The project identifies several primary customer segments, including industrial AI data providers and SMEs or startups that require accessible, trustworthy technology. Additionally, the KEAs are tailored for AI ethics researchers, data scientists, non-profit organisations, and high-level entities such as EU data hubs and policymakers who need to validate policy scenarios through generated models. These customers are reached through various channels, primarily the ALFIE online platform, social media, and active participation in conferences. Engagement is further supported by workshops and events designed to highlight the project's unique role in the responsible AI market.

Customer relationships are built on community engagement and long-term support and training programs. These interactions ensure a transition from transactional use to sustained adoption within target sectors. To ensure financial sustainability, the project employs a diverse set of revenue streams, including asset sales through competitive licensing and subscription models. Revenue is further bolstered by consulting services for industries, framework customisation on-demand, and the generation of specialized reports. Emerging assets like the Agentic Core also offer API integration and standalone software as potential sources of income.

The key activities necessary involve hosting open days and workshops, providing training and education. These activities are supported by vital resources, including the AutoML platform itself, its specialized knowledge base, and the model generation tools. Success also relies on secure data collaboration and rigorous adherence to GDPR and EU AI Act compliance standards.

A collaborative network of key partners forms the foundation of the ALFIE infrastructure. This consortium includes AI data hubs, data providers, research institutions, and government ethics

agencies, as well as industrial partners. The cost structure associated with this ecosystem encompasses fixed costs such as management, administrative expenses, and IPR management, as well as variable costs including R&D personnel, sales commissions, marketing outreach, and technical infrastructure maintenance. By aligning these elements, the ALFIE project positions its KEAs for long-term viability in the rapidly growing European AI market.

4.2 Target stakeholders and engagement

The ALFIE's success is built upon a multidimensional engagement strategy designed to reach a wide array of stakeholders, bridging the gap between advanced technology, regulatory policy, and societal impact. This approach ensures that the project's outcomes, such as the ETD-Hub and the AutoML platform, are grounded in real-world needs and aligned with the European Union's ethical and legal standards. At the core of this strategy is the early engagement of stakeholders, which includes EU and national policymakers and AI researchers. By involving these groups from the outset, ALFIE transforms them into early adopters who help refine the project's technical specifications and ensure that regulatory frameworks are informed by technical feasibility and public consensus.

The scientific and research community remains a critical partner, with universities and technology institutes collaborating on domain taxonomies, fairness benchmarks, and academic publications to advance the state of the art in Responsible AI. Simultaneously, industry stakeholders, including SMEs and large-scale enterprises like Bosch, are engaged through pilot use cases and industrial demonstrations that validate the practical utility of ALFIE's tools in operational environments. These collaborations allow the project to address industrial requirements such as bias mitigation and transparency in high-stakes sectors.

To ensure digital inclusion and social fairness, ALFIE actively targets civil society organisations, advocacy groups, and underrepresented communities. These groups can participate through the ETD-Hub, which serves as an open, multilingual platform where citizens can voice concerns and contribute to the ethical evolution of AI. This grassroots engagement is further supported by citizen awareness campaigns to empower the general public with the knowledge needed to advocate for responsible AI use. Educational institutions also play a key role, as the project integrates ethical AI principles into academic curricula and hosts guest lectures to prepare the next generation of professionals.

Engagement could be facilitated through a variety of channels, perhaps ranging from participatory workshops and hackathons to high-level policy assemblies and international forums. These interactions are supported by a suite of accessible materials, including multilingual videos, podcasts, and newsletters, ensuring that complex technical concepts are understandable to non-specialists. By fostering strategic alliances with standardisation bodies and regulatory agencies, ALFIE aims to play a defining role in shaping global standards for ethical AI, ensuring that its contributions have a sustainable and lasting impact beyond the project's duration, ensuring that the future of AI in Europe remains deeply rooted in human rights and social equity.

4.3 Exploitation impact indicators

To comprehensively evaluate the effectiveness of exploitation efforts and allow for informed adjustments to the project strategy, a structured set of impact indicators is implemented

throughout the duration of the initiative. These indicators function as vital metrics to monitor progress and measure the value generated by various activities. One primary measure of ALFIE uptake involves conducting demonstrations for a diverse range of stakeholders, with target values established for 10 to 15 policy makers or ethics organisations, 15 to 25 AI experts and operators, and 5 to 10 public authorities. Furthermore, the project will track the ALFIE output through the development of two to three open-source projects, which serves as a testament to community involvement and fosters a collaborative ecosystem for innovation.

Regarding technological innovations, the project monitors the introduction of new products and methodologies, aiming for more than 3 distinct products and at least 5 platform adoptions within the AI industry or relevant organisations. The success of commercialisation efforts is measured by the creation of at least two start-ups or spin-offs designed to drive the transition of research into market-ready solutions and create economic value. Finally, the impact of consultations and engagement is documented through a target of 20 consultations regarding the technology and platform, alongside achieving a goal of over 75 participants in workshops and open days to ensure broad outreach and adoption across various sectors. This framework enables timely adaptations of strategies based on insights gained from responsive stakeholder engagement and emerging industry trends.

4.4 Exploitation strategy for ALFIE KEAs

4.4.1 KEA1. ETD-Hub (Owner: DAS)

4.4.1.1 Market Entry and Expansion

The planned market entry strategy for ETD-Hub is phased, aiming to start with Europe and where possible then expand into other jurisdictions with similar AI and data governance frameworks.

1. **Initial anchor pilots:** ETD-Hub will be used by experts primarily in universities and at summer schools to create pilot data for in-project test use cases for Bosch, CTL, and UAB. This will:
 - Produce demo datasets.
 - Allow refinement of value proposition and business case.
 - Identify additional steps necessary in market entry.
2. **Open website deployment (public instance):** We begin full outreach post-ALFIE by launching and maintaining an open, web-accessible deployment of ETD-Hub as a reference instance. This will host public demos and selected consultations providing:
 - A live showcase of functionality.
 - A low-friction entry point for researchers, educators, NGOs and smaller public bodies.
 - A demonstrator for regulators and policymakers exploring tools to support AI Act implementation.
3. **Target priority segments in Europe:** Post-ALFIE, we will then undertake outreach to priority segments leveraging existing relationships within and beyond the consortium. We will look at the following main categories of organisation:



- EU institutions and agencies dealing with AI governance and fundamental rights.
 - Investors working in the AI domain.
 - Industrial players and SMEs working with European Digital Innovation Hubs (EDIHs) and other support structures on responsible AI.
 - Universities and research centres needing a platform for AI ethics research and teaching.
4. **Private and hybrid deployments:** In parallel, and post-ALFIE, we aim to offer private deployments (single-tenant SaaS or on-premises) for organisations seeking to discuss their own technologies internally with higher security, confidentiality or integration needs. This will include:
- Closed consultations (e.g. internal workforce deliberations, sensitive policy questions).
 - Instances more tightly integrated into internal management and AI development pipelines.
5. **Beyond Europe:** After establishing a European footprint and a portfolio of case studies post-ALFIE, we can target regions with comparable regulatory and ethical frameworks. With changes to the AI act it is unclear which regions are relevant as of yet, but they must meet the following criteria - (i) jurisdictions aligning with EU-style data protection and AI governance, and (ii) international organisations and global AI initiatives needing tools for multi-country consultations. Growth for ETD-Hub here is enabled by:
- The open website as a global demonstrator, offering English-language and multi-lingual spaces.
 - Partnerships with universities, networks and ethical AI initiatives outside the project to seed pilots and adapt workflows to local contexts.
 - A modular microservices architecture that allows localisation of legal, linguistic and policy elements while maintaining a common core.

4.4.1.2 Product Differentiation

ETD-Hub's differentiation strategy centres on domain specialisation, deployment flexibility, accessibility, and responsible-by-design architecture. Unlike general-purpose consultation platforms, ETD-Hub is built specifically for AI governance. It embeds AI-ethics taxonomies and risk categories making it immediately relevant for organisations facing AI Act obligations, developers who are conscious of societal impact, researchers or Policymakers seeking to access general opinion, and users looking to influence the future of AI in the EU.

ETD-Hub encourages users to structure debate, producing a substantial volume of detailed raw discussion data with a robust taxonomy. This is necessary to create second-order sentiment and recommendation datasets, which can feed directly into policy impact assessments and regulatory documentation, as well as AI development through SemKG/AutoML pipelines (e.g. constraints, metrics, model cards).

Through this connection to AI development workflows, ETD-Hub acts as an ethics and bias advisory layer for developers. This helps teams seeking concrete design constraints,

evaluation criteria and mitigations with an ethical focus. Value is also offered to non-experts or non-field researchers looking to develop ethically robust AI for novel use cases.

ETD-Hub works as both an open public forum and is forkable for an internal deliberation platform. These versions offer specific benefits. A public open website deployment enables low-barrier access, discovery, and demonstration. Private SaaS and on-premise deployments support organisations with strict confidentiality, security, compliance or integration requirements (e.g. large corporates, regulators, sensitive pilots) looking to canvas opinion and develop rulesets internally.

The platform was conceived with GDPR and responsible AI principles built-in, and this is supported by expert partners on the ALFIE project. This stands out against lighter-weight tools that rely heavily on generic data practices, offering clear consent flows, data minimisation, role-based access control, and transparency for analytics by serving a large, queryable dataset of raw opinion.

ETD-Hub's interfaces and workflows are designed for citizens, practitioners, policymakers, and developers, not only technical users. This lowers the barrier to entry for accessing, conducting, or interpreting complex debates and analytics. Additionally, ETD-Hub's UX approach encourages positive engagement with users given a focus on accessibility; permitting disabled users and users of different ability levels to use the platform. This expands userbase and enables those often underrepresented to have their voices heard. Finally, multilingual support expands the user pool across the EU.

Interoperability and extensibility are also key. Standards-based APIs, a microservices architecture, and modular components allow ETD-Hub to be adapted as required. In future, this could allow ETD-Hub to plug into existing civic participation ecosystems, AI lifecycle tools (MLOps, risk registers, documentation pipelines), and institutional portals and research infrastructures.

This combination of domain-specific depth and technical flexibility ensures ETD-Hub can maintain a competitive edge as AI governance matures.

4.4.1.3 Revenue Generation

We remain agnostic on the exact approach for monetising ETD-Hub until pilot testing for the use-cases is complete. The below options are presented without prejudice, and subject to change:

- **Subscription-Based Access**
 - Public cloud subscription for organisations to access a closed version of ETD-Hub as Software-as-a-Service.
- **Training Services**
 - Paid training programmes (online and on-site) for administrators, moderators, policymakers, legal teams, and developers.
 - Certification tracks (e.g. "ETD-Hub Certified Facilitator" or "ETD-Hub for AI Governance") to build professional communities and standardised competencies.
- **Consultation Services**
 - Tailored, facilitated workshops for sector-specific use cases (e.g. mobility/transport, public sector, accessibility).

- Provision of anonymised, aggregated datasets and analytic reports from ETD-Hub (where ethically and legally appropriate), for research, benchmarking or policy analysis.
- Commissioned bespoke analyses of consultations for investors, regulators, or corporates (e.g. impact on risk profiles, acceptability conditions, scenario comparisons).

4.4.1.4 Partnership Development

Partnerships are central to ETD-Hub's strategy, both within the consortium and with external organisations across academic, technological, research and regulatory spheres.

Project partners already form a strong initial ecosystem:

- **Bosch:** industrial partner for real-world use case deployment in vendor noncompliance.
- **Catalink (CTL):** industry partner on driver alertness and mobility, providing a concrete safety-critical use case.
- **Universitat Autònoma de Barcelona (UAB):** academic partner and accessibility use-case leader, assisting on inclusive design, accessibility standards and education.
- **Centre for Research and Technology Hellas (CERTH):** research partner for AI, analytics, and integration with technical infrastructures.
- **Kempelen Institute of Intelligent Technologies (KINIT):** applied research and ethical AI partner, able to support methods, evaluation and scientific dissemination.
- **Diadikasia Business Consulting (DBC):** business and consulting partner, supporting in compliance and policy planning.
- **University of Brighton (UoB) and Edge Hill University (EHU):** UK academic partners supporting ETD-Hub with SemKG.

Within this group, partnership efforts are currently focussed on co-developing pilot cases (CTL, Bosch, UAB) and creating joint training and education programmes (Bosch, EHU, UAB), i.e. integrating ETD-Hub into curricula and professional courses at the academic partners.

Candidate partnerships outside the consortium could broaden reach and legitimacy. Existing DAS partnerships with UK and Spanish universities (e.g. Royal Holloway University, University of Greenwich, Middlesex University, Universidad Complutense de Madrid) are all strong candidates for running pilots in their networks and city-regions after the ALFIE project, and we also aim to explore the possibility of integrating ETD-Hub into future teaching programmes on AI, law, social sciences and digital governance. We are also considering the possibility of co-leading post-ALFIE research projects and grant applications leveraging ETD-Hub; focused on AI governance, public participation, and digital democracy.

Other partnership development opportunities external to the consortium may be identified and pursued as ETD-Hub is developed further.

4.4.2 KEA2. AutoML platform (Owner: CTL, Co-owners: All)

4.4.2.1 Market Entry and Expansion

The strategy for market entry and expansion for the AutoML platform is designed to first penetrate the European Union market, capitalizing on the platform's unique alignment with the

EU AI Act and the increasing demand for ethically compliant AI solutions. The entry begins with a foundational three-year validation period within the project, during which trial use cases are used to build technical credibility and demonstrate real-world utility.

Upon the validation of the trial use cases, a transition phase will focus on finalizing the prototype of the platform and perform more validation tests.

The expansion strategy involves targeting international regions that maintain similar regulatory interests in human-centric, transparent, and trustworthy AI development. Commercialisation efforts will be spearheaded by the consortium's SMEs, which will leverage their business networks to execute targeted strategies while research institutions provide ongoing validation of ethical compliance to maintain market acceptance. The platform intends to establish strategic alliances with international organisations to help position the AutoML platform as a reference point for the industry. The sustainability of this expansion relies on a continuous market analysis to adapt the platform's features to emerging technological trends and evolving regulations. Progress and success will be measured through specific impact indicators, such as the creation of start-ups and spin-offs and the number of industrial adoptions.

4.4.2.2 Product Differentiation

The AutoML platform establishes clear product differentiation by moving beyond simple model automation to provide an end-to-end ecosystem for ethical and trustworthy AI. A fundamental point of divergence from existing competition is the integration of a bi-directional feedback loop with the ETD-Hub, which ensures that technical development is continuously steered by human-led policy and ethical discourse. This unique architecture allows the platform to translate deliberative ethical outcomes into actionable technical constraints, specifically ensuring compliance with the EU AI Act and global regulatory frameworks. Usability acts as a major differentiator, as the platform democratizes access to advanced machine learning by offering an interaction layer that supports natural language and speech commands in multiple European languages. This enables stakeholders who lack deep coding skills to design and build models that are automatically checked for bias and fairness. Unlike standard black-box AutoML tools, the platform features integrated explainable AI that generates adaptive interpretations tailored to different user levels. Furthermore, the platform incorporates specific measures for environmental sustainability. By combining advanced bias mitigation strategies, multilingual accessibility, and transparent auditing tools, the platform offers a specialized alternative to mainstream machine learning libraries that often overlook the societal and legal implications of AI development.

4.4.2.3 Revenue Generation

Regarding revenue generation, the monetisation approach will be finalized after the pilot testing of the use cases is complete. However, possible mechanisms for generating revenue from the AutoML platform revolves around possible subscription-based models, specifically employing a Platform-as-a-Service (PaaS) approach. This model will allow customers to access the platform's core features and ensure they receive essential software updates to maintain alignment with the evolving EU regulatory landscape.

The monetisation plan could also involve offering specialized consulting services designed to assist organisations with the practical deployment, customisation, and optimisation of AI models within their specific operational contexts. These services extend to providing

implementation assistance and strategic guidance for integrating the platform's machine learning modules into existing industrial or public sector infrastructures.

To enhance user proficiency and ensure the effective adoption of ethical AI tools, revenue could also be generated through the provision of paid training sessions, educational webinars, and technical workshops.

4.4.2.4 Partnership Development

The partnership development strategy for the AutoML platform focuses on establishing a robust network of strategic alliances across academic, technological, research, and regulatory sectors to drive innovation and facilitate broad market adoption.

A cornerstone of this approach is the collaborative ALFIE consortium, where SMEs can lead commercialisation and marketing efforts by leveraging their networks while universities focus on research, development, and the creation of training programmes for potential users. Research institutions can contribute by providing continuous validation of the platform's effectiveness and its adherence to evolving ethical norms, which is essential for establishing market trust. The involvement of industrial partners can allow for full-scale demonstrations in authentic operational environments, providing the necessary credibility for industrial uptake.

Beyond the consortium, collaborations with international standardisation organisations can help define and implement global standards for ethical AI, ensuring the platform serves as a primary reference point in the industry. Engaging directly with regulatory bodies, policymakers, and ethical advocates can ensure that the platform remains compliant with the shifting legal landscape. This strategy is designed to create a sustainable feedback loop between technological implementation and legal discourse, ultimately fostering a transparent and accountable AI development environment that resonates with both commercial and societal requirements.

4.4.3 KEA3. Data warehouse of ethical data (Owner: CERTH)

4.4.3.1 Market Entry and Expansion

The market entry plan for AutoDW focuses its initial efforts on the European market, leveraging the region's mature AI ecosystem and the strong regulatory frameworks established by GDPR and the EU AI Act. This approach responds to a significant market demand for transparent, auditable, and ethically aligned AI solutions that prioritise human-centric design. The plan focuses on early adoption in sectors with high governance and compliance needs, specifically finance, healthcare, public services, and civil society organisations. To build project credibility, the ALFIE's three trial use cases will be used and strategic partnerships across academic and industrial sectors will be developed, while engaging with global experts at major industry conferences.

4.4.3.2 Product Differentiation

AutoDW distinguishes itself by functioning as a centralised, high-performance data hub that unifies secure storage, comprehensive experiment tracking, and systematic model versioning within a single platform. This approach effectively eliminates the operational overhead associated with managing multiple disconnected tools.

Unlike traditional storage systems or basic streaming platforms, AutoDW serves as an orchestration backbone that coordinates complex AI pipelines in real time through an event-driven architecture. This sophisticated design ensures that machine learning operations remain scalable, resilient, and maintainable.

Furthermore, the AutoDW ensures full traceability for every experiment, model, and output. This end-to-end transparency is essential for supporting rigorous regulatory compliance, specifically aligning with the GDPR and the EU AI Act, while fostering ethical AI practices and organisational trust through its seamless integration with other ALFIE's components such as bias detection and explainable AI (XAI).

4.4.3.3 Revenue Generation

The commercialisation and long-term sustainability of the AutoDW may be driven by a dedicated **CERTH's spin-off company**. This entity is established to provide the commercial agility and industry-standard service levels required for professional deployment, acting as the primary commercial vehicle for the platform's market entry.

Revenue generation for the AutoDW will rely on licensing agreements and tiered subscription models. These plans will grant organisations access to AutoDW's core features, including its robust multi-database integration layer for flexible data management. Subscriptions provide access to a commercial full version of the asset, ensuring users benefit from ongoing software updates and versioning capabilities that maintain alignment with evolving EU regulatory frameworks and technical standards.

To complement recurring revenue, tailored deployment, integration, and optimisation support will be provided. These consulting services are designed to help organisations fully leverage AutoDW as a centralized data hub within their existing systems. This includes specialized assistance for system integration and optimizing the data access layer for high-stakes industrial or compliance workflows. By providing expert guidance on framework customisation, organisations can ensure that the AutoDW effectively supports their unique AI development and ethical auditing requirements.

4.4.3.4 Partnership Development

Academic collaborations involve partnering with universities and research institutes to co-develop AI models, validate methodologies, and advance data governance practices. Within this framework, AutoDW serves as the central persistence facility for these research outcomes, storing the trained AI and machine learning models produced through these academic R&D activities.

Regulatory and ethical engagement focuses on working closely with regulatory agencies, ethics boards, and the project's External Ethics Advisor (EEA) to ensure compliance with the EU AI Act and GDPR, anticipate legal changes, and embed responsible AI practices. This engagement ensures that AutoDW's data management protocols, such as anonymisation and purpose limitation, remain strictly aligned with evolving legal and ethical standards.

Finally, participating in industry consortia and research networks allows the project to accelerate adoption, share best practices, and showcase AutoDW's value in real-world deployments. By leveraging industrial partnerships with leaders like Bosch, AutoDW can demonstrate its effectiveness in managing diverse, large-scale datasets while maintaining the transparency and explainability required for broad market and regulatory acceptance.

4.4.4 KEA4. Semantic knowledge graph (SemKG) of ethical data (Owner: EHU, Co-owner: CTL)

4.4.4.1 Market Entry and Expansion

The initial market entry strategy for the SemKG focuses on establishing a strong foothold within the European Information and Communication Technology (ICT) sector, specifically targeting companies that require advanced knowledge management solutions, such as sophisticated indexing, filtering, and semantic search. To drive early adoption and lower entry barriers, the plan is making the software freely available for download and use. This approach is designed to cultivate a robust community who will validate the tool's effectiveness in real-world environments. The SemKG offers a unique context for AI systems, positioning it as a foundational resource for responsible innovation in the expanding European AI market.

Expansion and long-term sustainability are anchored in the technology's inherent adaptability, allowing it to move beyond AI ethics into diverse market segments. Because the software operates as an independent processing component, it can be re-purposed to populate knowledge graphs for any technical domain, provided the relevant ontologies are integrated to cover that specific field. The expansion strategy includes targeting regions outside the European Union that possess similar regulatory landscapes, ensuring the tool remains globally relevant as international ethical standards emerge. Following an initial period of free research use, the transition toward a commercialisation model may be considered.

4.4.4.2 Product Differentiation

The SemKG distinguishes itself as one of the first text-to-knowledge-graph software systems specifically engineered to map complex ethical discourse to well-defined and documented knowledge sources. Unlike traditional knowledge representation methods that focus on wider contexts, this technology is uniquely tailored to the AI ethics domain, utilizing a sophisticated suite of domain-specific ontologies, such as the Data Privacy Vocabulary (DPV), the AI Risk Ontology (AIRO), and the Fairness Metrics Ontology (FMO), to type entities and relations with high precision. The system provides a specialized machine-operable layer of ethics. By grounding AI reasoning in a community-driven body of discourse, the graph acts as a vital bridge between collective human deliberation and automated decision support, ensuring that ethical considerations are treated as a foundational layer of system design rather than a secondary concern.

4.4.4.3 Revenue Generation

While the underlying software component is initially intended to be freely available for research to encourage early adoption, the long-term roadmap involves a transition toward a commercial full version designed for large-scale applications.

Beyond software access, the revenue model also encompasses on-demand framework customisation and the generation of specialized reports, allowing businesses to address domain-specific ethical challenges that may fall outside standard guidelines. Market exploitation is further bolstered by offering framework extensions, such as the development of additional connectors and automated data integration tools that allow the knowledge graph to interface with diverse technical infrastructures. Finally, the provision of paid webinars and specialized training programs offers a scalable entry point for SMEs and researchers to

leverage the graph's reasoning capabilities, ensuring its long-term financial sustainability and market relevance.

4.4.4.4 Partnership Development

Partnership development for the SemKG is fundamentally focused on establishing strategic alliances across academic, technological, research, and regulatory sectors to drive innovation, enhance market reach, and facilitate technology adoption. Within the ALFIE consortium, EHU leads the design and deployment of the knowledge graph, while DAS facilitates the stakeholder discussions through the ETD-Hub that provide raw input data. This internal synergy is augmented by CTL, which validates the accuracy of the graph's linkages, and Bosch, which demonstrates the platform's capabilities in complex industrial environments.

Engaging with experts in AI ethics can offer guidance and ensure the graph remains responsive to the evolving technological landscape. These partnerships are essential for the creation of pilot projects, training programs, and technology transfer initiatives that build trust and marketability. Collaboration with industry stakeholders, developers, and startups allows to gain insights into specific sectoral needs and promotes the integration of ethical guidelines into commercial AI development pipelines.

4.4.5 KEA5. Agentic Core (Owner: KINIT)

4.4.5.1 Market Entry and Expansion

The Agentic Core will target sectors characterized by high regulatory scrutiny and ethical requirements, specifically healthcare, finance, public administration, and critical infrastructure, as these high-stakes environments demand the transparency and trust-building capabilities inherent in the ALFIE platform. To validate its sophisticated workflows and demonstrate real-world impact, the plan involves conducting evaluation runs in collaboration with various stakeholders relevant for this project. A key element of this expansion involves adapting KEA modules to address compliance requirements and leveraging the core's multilingual capabilities to satisfy local language and data regulation needs. Finally, initial lessons learned across various settings will be used to refine workflow templates, documentation, and deployment guides, thereby ensuring a streamlined and faster adoption process in subsequent global markets.

4.4.5.2 Product Differentiation

The Agentic Core distinguishes itself as a unique market solution by serving as the intelligent backend that harmonizes an intuitive AI chat interface with the specialized AG-UI standard for seamless, context-aware natural language interactions. It employs a sophisticated orchestrator-workers design pattern where a manager agent coordinates worker agents that specialize on different aspects that need to be considered when creating ethical AI solutions. This platform is particularly differentiated by its deep integration of ethics and compliance. By acting as the central orchestrator linking the AutoML engine with comprehensive XAI services, the Agentic Core contributes to the model transparency, fairness, and interpretability.

4.4.5.3 Revenue Generation

The revenue generation model for the Agentic Core incorporates specialized consulting services to provide users with dedicated deployment support and extensive workflow customisation. These offerings facilitate the integration with enterprise AutoML and MLOps

pipelines, ensuring that the orchestrator aligns with existing technical environments to enhance model production efficiency.

4.4.5.4 Partnership Development

Partnership development for the Agentic Core focuses on establishing strategic alliances with universities, research labs, and AI ethics centres to co-develop sophisticated workflows and publish demonstrators. These collaborations are essential for validating the Agentic Core within diverse research contexts, ensuring that the intelligent multi-agent backend meets academic standards for ethical AI. Furthermore, the plan involves seeking collaboration with industry partners to implement and customise these AI solutions for specific real-world use cases. This collaborative approach ensures that the orchestrator-workers pattern utilised by the core can be tailored to meet market demands and practical industry constraints.

4.4.6 KEA6. Web interface with XAI features (Owner: CTL, Co-owners: CERTH, KInIT)

4.4.6.1 Market Entry and Expansion

The strategy for entering the market with the ALFIE web interface revolves around a phased approach that prioritizes the European landscape before pursuing global opportunities. Initial entry focuses on penetrating the EU market by capitalizing on the unique alignment of the platform with the European Union's regulatory frameworks. This initial phase leverages ALFIE's unique niche in the ethical AI space, where the web interface serves as a transparent gateway that addresses the growing demand for accountable technology in different sectors.

Expansion will follow to focus on finalizing the first prototype, conducting validation tests, and refining the platform based on feedback from pilot use cases. The vision for growth involves expanding the interface's reach through a coordinated effort, end-user training sessions, and strategic technology transfers.

The plan involves establishing strategic alliances across academic and industrial sectors to facilitate technology adoption in real-world settings. The long-term plan ensures sustainability by transitioning from research prototypes to a commercial product. This transition is supported by a collaborative model where industrial partners demonstrate the tool's credibility in manufacturing and IT, while universities provide the research edge needed to maintain innovation. Revenue generation through possible subscription-based access will provide the financial foundation for continued growth and the scaling of the platform. After establishing its credibility in Europe, expansion will focus on countries with comparable regulatory frameworks. The market strategy relies on open scientific dissemination and institutional partnerships rather than commercial licensing.

4.4.6.2 Product Differentiation

The web interface distinguishes itself from traditional machine learning platforms by moving beyond technical black-box functionalities to provide a transparent, human-centric, and ethically grounded experience. While existing competition often requires deep data science expertise, this interface democratizes access to AI through an intuitive interaction layer powered by large language models integrated into AI agents. This allows users of varying skill levels to create AI solutions through natural language and speech interaction. This approach, which includes support for various European languages, ensures that technical barriers are removed for non-experts, positioning the tool as a unique solution in the global market.

A core pillar of product differentiation is the integration of sophisticated explainability features that foster trust and accountability. Unlike standard AutoML tools that provide results without context, the ALFIE interface employs task-appropriate explainability techniques to interpret model behaviour.

The platform further differentiates itself through its rigorous alignment with European legal and ethical frameworks. Its strengths include EU regulation alignment, multimodal interpretability, knowledge-graph-enhanced explanations, and human-understandable outputs designed for non-technical users.

Its transparent, open scientific architecture provides a competitive edge over proprietary AI governance platforms and gives it a unique position as a fairness-first XAI tool. This ethical focus is maintained through a unique bi-directional feedback loop with the project's ETD-Hub, ensuring that technical model generation is consistently refined by ongoing stakeholder deliberations on societal values. By offering a unified implementation platform that transforms raw data into bias-free, compliant models with clear audit trails, the web interface provides a level of transparency and regulatory assurance that is currently missing from many black-box commercial alternatives.

4.4.6.3 Revenue Generation

The revenue generation strategy for the web interface is built upon a financial model that combines recurring income with professional service fees. The primary foundation for income relies on possible subscription-based access. By offering the platform through a Platform-as-a-Service model, customers gain consistent access to the interface's features, including its intuitive interaction layer and essential software updates. This approach ensures a steady stream of revenue while allowing for scalable growth across various sectors.

Beyond standard access fees, ALFIE will generate income through bespoke consulting and professional services, such as services for fairness audits and explainability assessments. These services are designed to assist organisations with the specific deployment, customisation, and optimisation of the web interface within their existing technical infrastructures. This engagement helps bridge the gap for users who need tailored solutions for their specific industrial or regulatory environments.

Furthermore, revenue is bolstered by conducting paid training sessions and workshops and funded collaborations and pilot projects using the web interface as a methodological base. These educational offerings are essential for enhancing user proficiency and ensuring that the tool is seamlessly integrated into daily workflows.

By leveraging the ALFIE's unique position in the ethical AI market, the consortium aims to establish pricing policies that reflect the high value of unbiased and transparent AI tools. This model allows ALFIE to transition from a research-driven prototype to a commercially viable product that can effectively penetrate the European and global markets.

In the long term, these revenue streams will provide the necessary capital to maintain the platform's innovative edge and ensure it remains aligned with evolving legal and ethical standards.



4.4.6.4 Partnership Development

The partnership development strategy for the ALFIE web interface focuses on establishing strategic alliances across academic, technological and research sectors to enhance market reach. A central objective is to engage different stakeholders as early adopters, which helps to promote the project's results and increase support for the ethical AI market niche.

The consortium will seek to expand the assessment of its tools beyond project partners to evaluate outcomes. To support market acceptance, research institutions like CERTH and KInIT can provide technical proof of compliance and effectiveness, while industrial partners like Bosch demonstrate the interface's functionality in practical settings to build credibility with industrial clients.

Furthermore, building and strengthening relationships with diverse communities is essential, including civil society and advocacy groups focused on digital rights and privacy. Educational institutions could be integrated into this ecosystem to incorporate ethical AI considerations into academic curricula and promote awareness among students and future professionals.

Strategic partnerships will continue to be sought throughout the ALFIE's lifecycle to facilitate the widespread adoption and sustainability of its contributions to the ethical AI landscape. This ecosystem-based partnership model supports long-term adoption and impact.

5 Impact of ALFIE

5.1 Social, economic, and technological impacts of ALFIE across various domains

The ALFIE project is designed to connect ethical ideas with technologies to create responsible AI. At its core, the project seeks to democratise AI by ensuring that advanced machine learning tools are accessible to individuals regardless of their technical proficiency. This is primarily achieved through the synergy between the ETD-Hub and the AutoML platform, which together create a bi-directional feedback loop that allows ethical policy recommendations to directly shape the creation of machine learning models. This means that ethical advice and legal recommendations discussed in the ETD-Hub are used to directly guide how the AI models are built and refined on the AutoML platform.

The impact of ALFIE will begin during the project and continue after it finishes. The project has plans for commercialisation and long-term business models to ensure these tools stay useful and sustainable in the future.

A profound social impact of the ALFIE project is its foundational commitment to digital inclusion and accessibility, which is exemplified through trial use cases. The ETD-Hub serves as a multidisciplinary forum where different stakeholders converge to discuss the ethical implications of AI and provides a baseline of knowledge for stakeholders regardless of their academic background or technical proficiency. Furthermore, ALFIE aims to confront the pervasive issue of algorithmic bias, which has historically resulted in discriminatory outcomes in critical sectors. The AutoML platform incorporates a bias detection to systematically identify and correct discriminatory patterns. By analyzing training data for biases, ALFIE ensures that generated AI models perform consistently and equitably across diverse demographic groups. The integration of explainable AI (XAI) techniques reveals the specific decision paths of machine learning models, making automated results transparent, understandable, and contestable for all members of society. Beyond direct engagement, citizens benefit indirectly as public authorities and governmental bodies may adopt ALFIE's KEAs to improve public safety and social fairness. The adoption of these ethical AI tools by organisations can help ensure that technology prioritises public benefit and adheres to human rights standards to promote citizen welfare.

In terms of economic impacts, ALFIE and its KEAs have the potential to significantly affect the ethical AI market. Tools such as the AutoML platform, the Data Warehouse and Semantic Knowledge Graph, as well as the multilingual tools and XAI features, offer significant innovation and market competitiveness to potential adopters. ALFIE pushes the boundaries of automated machine learning allowing users to interact with the system using natural language and speech, translating human intent and requirements into structured machine learning pipelines without requiring advanced coding skills. The integration of Semantic Knowledge Graphs further enhances this process by linking ethical principles, legal frameworks, and regulatory requirements directly to the model generation phase. This ensures that every AI tool produced is "ethical by design," adhering to the specific constraints of the European AI Act and other legal standards. Another significant technological contribution is the advancement of Explainable AI (XAI). ALFIE provides users with clear justifications for model outputs, identifying feature importance and decision paths to ensure that automated results are both

reliable and auditable. Moreover, the project leads the way in addressing the environmental costs of the AI. By prioritising transfer learning, quantisation, and mixed-precision training, ALFIE reduces the computational power and electricity required for model development. The system aims to maintain a sustainable balance between high performance and environmental responsibility, while providing users with approximate environmental indicators associated with their training activities.

By providing SMEs and startups with user-friendly tools to build their own ethical AI solutions, ALFIE lowers the barriers to entry in the technology sector and encourages the adoption of responsible innovation. In addition, by developing AI infrastructure compliant to EU ethics and data standards, ALFIE reduces dependencies on non-European ecosystems. Due to the transparent and explainable models which are the output of AutoML, its adoption encourages greater market stability and consumer confidence. Moreover, market uptake may be accelerated through the trust-based adoption of AI tools in policymaking, whereas the project may support the creation and adoption of ethical certification schemes which will identify compliant European AI solutions globally.

Ultimately, the ALFIE ecosystem creates a self-sustaining cycle of innovation, oversight, and societal well-being. By continuously updating its central knowledge base with insights from global experts and public stakeholders, the platform remains highly responsive to evolving societal needs and emerging technological risks. This holistic strategy ensures that the future of AI in Europe is not only technologically superior but also deeply rooted in human rights, safety, and social equity, providing a blueprint for trustworthy AI development that can be scaled across various industrial and social domains.

5.2 Standards that influence ALFIE's solutions

The ALFIE project is constructed upon a rigorous foundation of international, European, and industry-specific standards that ensure its technological solutions are not only innovative but also legally compliant and ethically sound. Such standards ensure that ALFIE is aligned with EU and international norms, and that its outputs are trustworthy and interoperable and contribute to the promotion of ethical AI. However, ALFIE aims to go beyond merely complying with the current industry standards; it aims to play a defining role in shaping them and contribute to their evolution, thus ensuring a lasting impact on the collaborative landscape of AI and ML technologies.

At the forefront of these influences is the EU Artificial Intelligence Act (AIA)⁵⁸, which represents the first comprehensive global attempt to regulate AI technologies. ALFIE adopts the AIA's risk-based regulatory framework, which classifies AI systems into four categories: unacceptable, high, limited, and minimal risk. This classification is crucial because it dictates the level of legal safeguards and technical requirements necessary for each application. For high-risk systems, the AIA mandates strict adherence to standards involving comprehensive risk management, data governance, technical documentation, and human oversight. ALFIE's long-term ambition is to ensure all outputs are designed for full compliance should they be placed on the market.

⁵⁸ <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689>

Parallel to the AIA, the General Data Protection Regulation (GDPR)⁵⁹ serves as the primary standard for managing personal data within the ALFIE ecosystem. The project adheres to the seven key principles of GDPR, including lawfulness, purpose limitation, and data minimisation, ensuring that only the data strictly necessary for research objectives is collected and processed. Compliance with GDPR ensures that participants' privacy is protected through technical measures like pseudonymisation and encryption, while also guaranteeing their rights to access, rectification, and erasure.

Ethical excellence in ALFIE is guided by the Ethics Guidelines for Trustworthy AI adopted by the High-Level Expert Group on Artificial Intelligence (HLEG-AI)⁶⁰. These guidelines establish four primary ethical principles: human autonomy, prevention of harm, fairness, and explicability. ALFIE translates these high-level principles into practical requirements for its AutoML platform, ensuring that the generated models are technically robust, transparent, and accountable. ALFIE also adheres to the Assessment List for Trustworthy Artificial Intelligence (ALTAI)⁶¹ to verify that its systems are lawful and ethical across diverse demographic groups. Furthermore, the European Code of Conduct for Research Integrity (ALLEA)⁶² informs the project's scientific conduct, emphasizing reliability, honesty, and respect for research participants.

Technically, ALFIE aligns with the User-Centred Design (UCD) standards ensuring that the platform is developed through iterative feedback from diverse stakeholders, making advanced AI tools accessible to non-experts. To ensure that research outcomes are of the highest quality and reusable, ALFIE implements the FAIR principles, which mandate that all data be Findable, Accessible, Interoperable, and Reusable. This includes the use of persistent identifiers and the adoption of metadata standards for the AutoDW and the Semantic Knowledge Graph. Security and interoperability within the distributed architecture of the AutoML platform are managed through standard protocols like OAuth2 and OpenID Connect for identity management via Keycloak. The platform uses RESTful APIs and open standards such as JSON and XML for structured documents to ensure compatibility across different services and organisational systems. Finally, ALFIE considers emerging environmental sustainability standards. This involves adhering to guidelines that prioritize model efficiency, minimizing energy consumption.

Through this multi-layered adherence to legal, ethical, and technical standards, ALFIE creates a trustworthy ecosystem that bridges the gap between high-level discourse and practical AI implementation.

⁵⁹ <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>

⁶⁰ <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>

⁶¹ <https://digital-strategy.ec.europa.eu/en/library/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment>

⁶² <https://allea.org/code-of-conduct/>

6 Conclusions and Future plans

The ALFIE project is dedicated to harmonising human-centric values with advanced technological development, specifically through the creation of the ETD-Hub and the AutoML platform. The project's success is built upon a multidimensional strategy designed to bridge the gap between advanced technology, regulatory policy, and societal impact, ensuring that its outcomes are grounded in real-world needs and aligned with European ethical and legal standards. By transforming ethical considerations into a distinct competitive advantage, ALFIE aims to position European industries and research institutions at the global forefront of the ethical AI market.

By identifying six KEAs, ALFIE creates a high-value ecosystem that functions as both an integrated suite and a series of stand-alone commercial solutions. These assets are designed to fortify the European Union's position in the global market by transforming ethics into a competitive advantage and a measurable trust layer for high-stakes environments.

ALFIE's future plans involve a strategic transition from research-level prototypes to market-ready solutions. Initial efforts will focus on pilots in university settings, which will be used to refine the project's value proposition. To maintain public and regulatory visibility, the consortium will launch the ETD-Hub and AutoML platform to serve as a showcase for regulators exploring the practical operationalisation of the EU AI Act. Furthermore, technical maintenance will include the integration of human-in-the-loop validation systems to allow experts to review AI-generated summaries and refine model outputs for maximum accuracy.

Market expansion will then target industrial players in sensitive sectors like healthcare and finance, before eventually pursuing global opportunities in jurisdictions with comparable regulatory frameworks.

Appendices

Appendix A: PESTEL analysis for ETD-Hub

Type	Barriers	Mitigation Plan
P	Uncertainty and fragmentation in AI governance priorities: Different EU Member States and agencies may prioritise AI, participation, and ethics differently, leading to uneven demand and long decision cycles.	We will position ETD-Hub explicitly as a support tool for AI Act implementation and related EU digital policies, aligning messaging with EU-level narratives (trustworthy AI, fundamental rights, citizen participation). Marketing will target early adopter countries/agencies first (those already active on AI governance), using their pilots as political proof points.
P	Changing political leadership and agendas: Elections or changes in government can stall or cancel participatory/AI governance initiatives that ETD-Hub benefits from.	Where possible we will target anchor deployments that focus in part on general administrative and regulatory needs that will outlive specific leadership (compliance, risk assessment). We also plan to develop a multi-year roadmap and frame ETD-Hub as neutral infrastructure, not tied to a specific party or government.
P	Politicisation of AI debates and risk of capture: AI and AI ethics can become heavily politicised (e.g. framed as for or against innovation), with powerful stakeholders trying to steer participatory processes to legitimise pre-decided positions.	We have designed ETD-Hub processes with clear governance rules (balanced stakeholder representation, transparent moderation, published methodologies). Marketing and moderation practices will emphasise ETD-Hub's role as neutral infrastructure, separating the platform's credibility from any single political agenda.
E	Price pressure from low-cost and open-source alternatives: Many civic engagement and deliberation tools are free or very low cost (e.g. open-source platforms supported by municipalities or NGOs). If we charge for ETD-Hub, potential customers may struggle to justify paying when "good enough" alternatives appear available at minimal cost.	Any marketing will clearly differentiate ETD-Hub through its AI- and ethics-specific capabilities, data, and integration with AI development/AutoML workflows that generic open-source tools don't provide. In this way, ETD-Hub offers service-centric value rather than competing purely on any potential licence price. Alternatively, ETD-Hub may be free at the point of access.
E	Budget constraints in public sector and academia: Public administrations, universities and NGOs often face limited budgets and procurement hurdles for new digital platforms.	ETD-Hub pricing (if any) remains undecided. If we charge, we will explore tiered pricing, including low-cost or freemium options (especially for research/academic use). We will also examine co-funding models (projects, grants, regional programmes) to subsidise initial deployments. Regardless, we will provide clear cost-benefit arguments (reduced consultation costs, reusable datasets, compliance support) to demonstrate value return.
E	Competition for internal investment in AI tooling: Organisations may prioritise spending on core AI capabilities (MLOps, cloud, analytics) over governance tools like ETD-Hub.	We plan to position ETD-Hub as part of the AI lifecycle, not a separate "nice to have". We can show users how ETD-Hub feeds AutoML, risk management, and documentation. Using project case studies, we also plan to quantify how stakeholder-informed design can reduce later costs (reputational damage, redesign, regulatory issues).
E	Slow adoption and/or procurement processes: Public administrations and larger enterprises often require complex, multi-step procurement	We plan to provide template text for tender specifications and internal justifications (business case, risk analysis), helping champions inside client organisations push adoption through with less effort. In the case of private

	procedures (tenders, legal reviews, IT security checks). Even if ETD-Hub is attractive, these transaction costs and long lead times can delay or derail adoption; especially for smaller contracts.	deployments, where feasible, post-ALFIE we plan to pursue framework agreements or catalogue listings (e.g. with central procurement bodies or partner integrators), so individual departments can procure ETD-Hub with simplified processes.
S	Participation fatigue and low engagement: Citizens and stakeholders may be sceptical that their input on AI will matter, leading to low turnout or superficial contributions.	ALFIE aims to design processes that clearly connect deliberation to decisions (e.g. feedback loops, visible policy changes). We have also designed simple, accessible user interfaces using low-friction participation options (mobile, asynchronous) based on expert advice from UAB. Finally, we are in the process of using trusted intermediaries (NGOs, professional associations) to recruit engaged participants and circumvent this issue for key use case deployments.
S	Digital divides and representativeness concerns: Online deliberation can under-represent less connected or marginalised groups, undermining the legitimacy of insights.	We support accessibility by design (including UAB's accessibility use case as a flagship), multiple languages, and easy onboarding.
S	Low AI and data literacy among participants: Many citizens and even some professionals may not fully understand AI concepts, limitations, and risks, leading to superficial or misinformed contributions.	We plan to expand our existing resource list and embed accessible explainer content (short videos, FAQs, visualisations) directly in ETD-Hub deliberation spaces, likely following the ALFIE project. Additionally, per the code of conduct users must upload guided question formats and structured prompts that help participants reflect on concrete scenarios rather than abstract technology. We also plan to partner with education and outreach organisations following ALFIE to refine the ETD-Hub offering.
S	Polarisation and hostile discourse around AI: Debates on surveillance, automation, or algorithmic discrimination can become polarised and hostile, discouraging participation and undermining constructive deliberation.	We have implemented robust community guidelines and moderation tools (flagging, escalation workflows, expert moderators). We also use interaction designs that reward constructive contributions (reputation score, bronze/silver/gold badging).
T	Limitations and bias in summaries or analytics drawn from ETD-Hub, especially given multilingual support: Analytics by users (e.g. topic modelling, sentiment, stance) may misinterpret nuanced ethical arguments, sarcasm, or minority languages; introducing bias into their derived insights.	We plan to continuously evaluate and document platform performance and limitations, particularly across languages and subgroups. We will also provide to SemKG a dataset which could allow users to drill down from analytics to raw contributions, ensuring transparency and enabling reinterpretation.
T	Integration complexity with existing IT and AI pipelines: Clients may be reluctant to adopt a tool if they believe it requires significant integration (identity management, portals, AI dev pipelines).	ETD-Hub is standalone and will not usually require much integration into a stack except potentially to support data exports assuming this is something we elect to allow. In the case of a future spin-off private internal deployment, we plan to offer implementation support packages and provide reference integrations (e.g. example connectors to AutoML). Generally, we aim to maintain a modular architecture with well-documented APIs and standard protocols.
T	Performance, scalability, and security expectations: As a	We are using scalable deployment methodologies (containerisation, automated monitoring). We also aim to

	participation platform dealing with AI-related topics, ETD-Hub must handle potentially large user loads, ensure uptime, and protect sensitive data.	conduct security and performance testing as necessary in future, and have already adopted other recognised security practices (regular audits, keycloak as SSO, incident response plans). For any future private version, we plan to clearly communicate service levels through SLAs in the hosting offer.
E	Perception of digital tools as environmentally costly: As awareness grows of the carbon footprint of digital infrastructures and AI, some stakeholders may criticise additional platforms as environmentally burdensome.	We have optimised code and infrastructure for resource efficiency (e.g. lightweight data management where possible). We also plan to explore the possibility of serving sustainability reporting (approximate energy/CO ₂ metrics) as part of the platform's transparency.
E	Actual hardware and infrastructure footprint (energy, e-waste): Running additional servers or on-premises instances for private versions can increase energy consumption and, over time, contribute to hardware turnover and e-waste, which may be criticised in sustainability-focused organisations.	We will largely support the cloud-first, CERTH infrastructure web deployment where possible, reducing the need for new hardware. Should any future private instance be required, we plan to provide deployment guidelines that encourage reuse of existing infrastructure and virtualisation instead of dedicated hardware. We plan to also offer best-practice recommendations on green procurement and lifecycle management for clients that require these private installations.
E	Emerging requirements for green IT and sustainability reporting: Public bodies and large enterprises increasingly require digital tools to meet internal sustainability policies (e.g. carbon accounting, green public procurement), which could slow adoption if ETD-Hub cannot provide relevant evidence.	We are exploring the feasibility of documenting and sharing high-level resource use characteristics (e.g. typical CPU/storage footprint). In future, we also aim to prepare template text and indicators that clients can use in their own sustainability reports and procurement justifications, lowering the administrative burden of adopting ETD-Hub; this is likely to occur post-ALFIE.
E	Competing priorities with green agendas: Public bodies may prioritise climate/environment tools and see AI-ethics deliberation as secondary.	Any marketing will exemplify how ETD-Hub can support AI use cases in environmental policy (e.g. AI for mobility, smart grids, climate adaptation), thus combining both agendas. To this end, we target cross-domain pilots (e.g. driver alertness to support smart public transport) to position ETD-Hub as relevant to green transitions.
L	Ownership and reuse rights for deliberation datasets: Questions may arise over who owns and can reuse anonymised sentiment and recommendation datasets (platform operator, organisers, participants, or all of them), potentially blocking exploitation or sharing.	Legal expert partners DBC have defined standard legal terms and licences for data generated via ETD-Hub (e.g. clear clauses in terms of use and project contracts). We requested these experts produce explicit privacy policies explaining how aggregated/anonymised data will be reused and by whom.
L	Strict data protection and privacy requirements (GDPR): ETD-Hub processes potentially sensitive opinions and discussions; missteps in consent, anonymisation, or data handling could create legal exposure.	In ALFIE we have implemented privacy-by-design and by-default: data minimisation, clear consent flows, and robust anonymisation/pseudonymisation; with thorough DMP and DPIA supported by EHU. We have engaged legal and ethics advisors in platform design (KINIT, DBC).
L	Evolving AI and digital regulation: The EU AI Act and related regulations may introduce new obligations (record-keeping, risk assessment, transparency) that affect how ETD-Hub	DAS maintains a regulatory watch and plans to update workflows, policies etc if rules are updated. We have also designed ETD-Hub so that its direct outputs and second-order outputs from the SemKG /AutoML systems (sentiment and recommendation datasets, deliberation records) map directly to compliance needs (e.g. risk logs,



	operates, especially when used in high-risk AI contexts.	stakeholder engagement evidence). Our compliance settings (retention periods, consent options, data localisation) are configurable to adapt to different or changing jurisdictions or sectors.
--	--	--

Appendix B: PESTEL analysis for AutoML platform

Type	Barriers	Mitigation Plan
P	Use cases classified as "High-Risk" under the EU AI Act, inviting heavy regulatory scrutiny.	Use the ETD-Hub as a collaborative space where policymakers and technical partners can align AutoML model generation with human rights and EU values.
E	Training vision models and Large Language Models (LLMs) is computationally expensive, potentially creating a barrier for the intended users.	Prioritise Transfer Learning and fine-tuning existing models rather than training from scratch.
S	Non-expert users may resist AutoML-generated models if they cannot understand the logic behind the decisions.	Integrate XAI to provide explanations, simplified for non-experts, detailed for specialists.
T	LLM-driven components might generate misleading configurations or incorrect code snippets for the AutoML engine.	Implement Human-in-the-Loop (HITL) validation where AI-generated summaries and model outputs are reviewed by human experts for accuracy.
E	Training of AI models brings about a significant amount of water usage as well as a significant carbon footprint.	A range of training optimisation techniques is incorporated into the system to serve sustainability goals.
E	The energy-intensive nature of AutoML and hyperparameter tuning clashes with EU sustainability goals.	Adopt "green" strategies: exclude models with >500 million parameters, use mixed-precision training (16-bit), and implement training time constraints.
E	The development of AI processes is inherently associated with a risk of excessive water and energy resources.	Minimising the carbon footprint and energy consumption of AI training through model quantisation, mixed-precision training, and the prioritisation of medium-sized models. An environmental impact assessment will be considered to be carried.
L	Use cases involve sensitive data processing that must meet strict GDPR and AI Act data governance standards.	Implement Privacy-by-Design; use pseudonymisation, differential privacy, and mandatory informed consent for all participant data.

Appendix C: PESTEL analysis for Data warehouse of ethical data

Type	Barriers	Mitigation Plan
P	Changes in government AI policies or funding priorities could limit adoption in public sectors.	Engage proactively with policymakers and participate in public AI initiatives to align AutoDW with evolving priorities.
E	Economic downturns may reduce organisational budgets for new AI infrastructure.	Offer flexible subscription tiers, pilot programs, and ROI-focused demonstrations to make adoption more cost-effective.
S	Public skepticism about AI transparency and data privacy could slow adoption, especially in civil society contexts.	Emphasize AutoDW's governance, traceability, and compliance features in marketing and outreach campaigns.
T	Rapid advancements in cloud infrastructure and AI platforms could render AutoDW less competitive in scalability.	By utilizing industry-standard containerisation and a microservices-based architecture, the platform ensures that it can leverage the scalability of major cloud providers while maintaining independent control over data governance and ethical auditing.
E	Increasing pressure to reduce energy consumption and carbon footprint of AI workloads.	Optimize AutoDW for efficient resource usage, offer deployment on green cloud platforms, and track energy metrics for transparency.
L	Complex and evolving data protection laws (GDPR, sector-specific regulations) could create compliance challenges.	Implement built-in compliance features, regular audits, and maintain a legal advisory network to ensure adherence to regulations.

Appendix D: PESTEL analysis for Semantic knowledge graph of ethical data

Type	Barriers	Mitigation Plan
P	Varying AI policies across EU member states can lead to fragmentation, hindering the harmonisation of legal frameworks and impeding the effectiveness of the SemKG.	The SemKG will serve as a medium to conduct comparative policy analyses and identify regulatory differences.
E	Stricter ethical and legal requirements represented in the SemKG may increase compliance costs, particularly for SMEs with limited financial resources.	Engage with SMEs to provide guidance and support.
E	Businesses may face significant transitional costs as they update their technologies and workflows to integrate the structured ethical reasoning provided by the SemKG.	Emphasize cost-effectiveness by showing how ethical AI reduces risks of reputational damage and legal fines; offer technical assistance and support programs to ease the transition.
S	Digital literacy disparities among different groups may impact their ability to understand the SemKG's outputs.	Collaborate with organisations focused on digital literacy, provide accessible resources, and offer specialized training programs to ensure broader participation.
T	Integrating the SemKG with other ecosystem components (such as AutoML) may face interoperability challenges, hindering seamless data flow.	Adopt standardized protocols to ensure compatibility and conduct regular testing with technology providers to address technical friction.
E	The increasing proliferation of AI and the large-scale data processing required for the SemKG may contribute to strained global energy demands.	Integrate clean energy production concerns and sustainability parameters directly into the ETD-Hub discussions and the AutoML outputs.
L	Challenges in ensuring that the abstract ethical considerations captured in the SemKG are legally enforceable.	Collaborate with legal experts in AI ethics to bridge the gap between high-level principles and actionable legal frameworks.



Appendix E: PESTEL analysis for Agentic Core

Type	Barriers	Mitigation Plan
P	Cross-border data restrictions or AI export regulations may limit deployment in certain regions.	Use data storage located in the EU.
E	Established platforms may undercut pricing or bundle orchestration modules.	Focus on differentiation through ethics, multilingual support, and automated ML workflow.
S	Skills and Knowledge Gap.	Natural-language guidance can reduce friction.
T	LLM hallucinations may provide false information.	Safeguards such as Judge LLM can reduce the number of incorrect answers.

Appendix F: PESTEL analysis for web interface with XAI features

Type	Barriers	Mitigation Plan
P	Low adoption of AI governance tools in some EU Member States; limited political alignment on fairness standards.	Engage policymakers early; present evidence from ALFIE pilots; contribute to national AI strategies and EU-level discussions.
P	Fragmented AI policies across EU member states may hinder the harmonisation of the interface's legal frameworks.	Use the project as a medium to engage policymakers and conduct comparative analyses to produce actionable recommendations for policy alignment.
E	Budget constraints in public administrations and SMEs may slow adoption.	Provide open-source access; offer low-cost training; demonstrate cost savings through reduced compliance risk.
E	High financial resources are required to maintain a comprehensive ecosystem like the web interface.	Establish a sustainable funding model by exploring public-private partnerships and EU grants.
S	Limited public trust in AI and low awareness of bias/fairness issues.	Create accessible demonstrations and visual explanations; collaborate with civil-society groups to build understanding.
S	Low digital literacy may prevent certain groups from actively using the interface.	Collaborate with digital literacy initiatives and provide training programs and simplified resources to bridge the gap.
T	Fairness metrics require high-quality data, which may be incomplete or inconsistent.	Implement robust data-quality checks; allow configurable fairness definitions; provide fallback interpretability workflows.
T	Technical complexity of AutoML may limit accessibility for users without deep technical backgrounds.	Prioritise user-friendly interfaces, intuitive natural language interaction, and comprehensive documentation for non-experts.
T	Scalability and performance issues as the volume of users and discussions grow.	Conduct regular performance testing to resolve bottlenecks.
E	Increasing scrutiny on energy-intensive AI models.	Emphasize lightweight, efficient explainability methods; benchmark energy use and highlight minimal overhead.
E	High energy demand from the proliferation of AI products and training activities.	Integrate energy consumption parameters into the AutoML platform and include sustainability concerns in the ETD-Hub discussions.
L	Uncertainty around upcoming AI Act conformity requirements and fairness documentation obligations.	Monitor regulatory updates; align KEA outputs with draft compliance documentation; engage with regulatory sandboxes.

L	Non-compliance with privacy laws (GDPR) and emerging AI-specific regulations leading to legal consequences.	Establish clear guidelines for data handling and user consent within the platform.
L	Intellectual Property (IP) disputes regarding the sharing of novel AI models and frameworks.	Clearly define ownership and licensing terms and establish protocols for the fair use and sharing of intellectual property.