

F4E Document**Artificial Intelligence Technology Roadmapping Report v2026**

Artificial Intelligence (AI) is rapidly emerging as a transformative enabler for fusion engineering, offering new opportunities to accelerate design, manufacturing, qualification, operation, and maintenance activities across the fusion lifecycle. To support strategic decision-making and technology prioritization, Fusion for Energy (F4E) conducted a comprehensive AI Technology Mapping exercise under its Technology Development Program (TDP), engaging experts from fusion organizations, research institutions, industry, and AI technology providers. The study identified and assessed more than 70 AI Technology Development Activities (TDAs) across key fusion engineering domains, including design and analysis, manufacturing and quality control, data and knowledge management, systems engineering, diagnostics, predictive maintenance, robotics, digital twins, and project management. Each activity was evaluated in terms of technical maturity, implementation timeline, estimated cost, probability of success, data readiness, and strategic impact. The results indicate that while most AI applications remain at an early to intermediate maturity level, significant opportunities already exist for near-term deployment, particularly in knowledge management, document intelligence, quality assurance, computer vision, predictive analytics, and engineering data exploitation. The assessment also highlights critical dependencies, including data quality, governance, digital infrastructure, human resources, and AI competencies. A strategic roadmap is proposed covering foundational capabilities, deployment-oriented developments, and longer-term ambitions such as integrated digital twins and AI-assisted fusion operations. This report provides the first dedicated technology mapping of AI applications for fusion engineering within the F4E Technology Development Program and offers a structured framework to guide future investments, collaborations, and innovation initiatives. The findings demonstrate that the successful adoption of AI in fusion will depend not only on advances in algorithms, but also on the establishment of robust data ecosystems, governance frameworks, and... (Please see complete abstract on document metadata.)

Approval Process			
	Name	Action	Affiliation
Author	● F4E.DI.FT.PE	23 September 2026:signed	FUTED
Co-Authors			
Reviewers	● F4E.DI.FT.RA		FUTED
Approver	● F4E.DI.FT		FUTED
Read Access	LG: IRIS Team, LG: LG AI experts, LG: Antennas enlarged review group, LG: F4E-ext, AD: IDM_F4E, GG: EC Representative, GG: IAC, project administrator, RO		

Original Document MD5#: 227D230D85A5C33B2E8664CA6DC058E9

** Performed by ** Delegated Reviewer**For Data Protection the names of the actors are omitted and can be consulted in IDM**Printed copies are not controlled. Confirm version status through the F4E document management system (idm@F4E)*

Generated on 23 September 2026

Technology Development Programme



Technology Mapping

2025 Series

Artificial Intelligence in Fusion



**FUSION
FOR
ENERGY**



Contents

Version history	3
Foreword	4
Executive summary.....	5
1 Introduction	6
1.1 Context.....	6
1.2 Artificial Intelligence Technology mapping.....	6
2 Technology mapping process.....	7
2.1 Input report.....	7
2.2 Online workshop	7
2.3 In-person workshop	8
2.4 Final report.....	8
3 Artificial Intelligence Breakdown	9
3.1 Artificial Intelligence overview.....	9
3.2 Technical breakdown of applications.....	9
3.3 Priority Research Opportunities (PROs)	11
3.4 Map of individual applications.....	12
3.4.1 Application of AI in design and simulation.....	13
3.4.2 Application of AI in materials, manufacturing, NDT, and assembly	17
3.4.3 Application of AI in plasma science	22
3.4.4 Application of AI in system engineering, project management, and operations.....	23
4 Summary of the workshop.....	28
5 Outcome: Technology Roadmap.....	29
5.1 Technology dashboards.....	29
5.2 Overview of the AI in fusion landscape in the EU	30
5.2.1 Technology and TDA analysis	30
5.2.2 SWOT analysis	39
5.2.3 Technology Roadmaps	44
6 Conclusion	45
7 Bibliography	47
Appendix 1: Technology Readiness Levels.....	50
Appendix 2: Technology assessment	51
Appendix 3: Technology dashboards	55

Version history

VERSION	DATE	CHANGES
0.0	07/03/2025	First issue: input data for online workshop. Covers: 1. Introduction 2. The mapping process 3. AI technology breakdown (draft) Other sections will be completed after the workshop.
1.1	15/05/2025	After the online workshop, incorporating the changes agreed to the technology map.
2.0	19/09/2025	After the in-person workshop – draft final report made available for comment by participants
2.1	22/09/2026	Final report for publication

Foreword

We are entering a new phase in fusion development, one in which artificial intelligence (AI) is becoming an increasingly important enabler and catalyser of progress. Across research, engineering, manufacturing, and operations, AI can accelerate analysis, improve decision-making, de-risk engineering and manufacturing, and strengthen the pathway from scientific discovery to industrial application. Together with other technological innovations, AI has the potential to shorten the path towards commercial fusion. Realising this potential requires academia, public research organisations, industry, and startups to align around shared priorities, needs, current workstreams, and opportunities, minimising duplication and maximising benefits for the fusion community.

As a European fusion technology hub, Fusion for Energy (F4E) is well placed to help structure this effort. By bringing together stakeholders from across the fusion ecosystem, F4E can identify high-impact research and development needs, promote collaboration, and help channel investment towards the AI capabilities most relevant to future fusion power plants, engineering and operations, while supporting a competitive European supply chain. To this end, F4E organised its first technology mapping workshop dedicated to AI in fusion. The workshop brought together participants from public and private organisations across Europe and non-fusion specialists to assess the current landscape, identify promising applications, and discuss the actions needed to move from early exploration to practical implementation.

This report is the outcome of that collective effort. It provides an overview of the AI applications across the fusion domain, highlights the main opportunities and bottlenecks, and proposes a structured roadmap for future development. It identifies priority gaps and opportunities and provides an initial effort assessment required to address them. Given the nature and rapid pace of AI development, this should be regarded as an initial roadmap rather than a fully comprehensive one. It represents a first collective effort and will need to be reviewed and updated regularly as technologies, priorities, and capabilities evolve. The report is intended as a reference for the European fusion community and a basis for future collaboration, investment, and funding decisions.

The exercise has also established a methodology that can be extended to other technology areas. By combining technical mapping with stakeholder input and strategic prioritisation, this approach can help build a more coordinated and forward-looking innovation framework for fusion in Europe. F4E will contribute through targeted actions that complement ongoing European initiatives, including those pursued through EUROfusion, Fusion Now, Euratom, EIC and other public and private partnerships. Progress will depend on sustained collaboration, shared data and infrastructure, and a common ambition to turn emerging capabilities into practical impact. We invite all stakeholders to take part in that effort.



Dr. Enrique García-Vidorreta
Fusion for Energy
Head of Fusion Technologies and
Engineering Department



María Ortiz-de-Zúñiga López-Chicheri
Fusion for Energy
Head of Project Engineering, CAD and
Data Management Unit

Executive summary

All following analyses are based on the F4E workshops of April and May 2025.

Artificial intelligence is expected to play an increasingly important role in advancing fusion energy, supporting improvements in performance, reliability, and operational efficiency. The 2025 Artificial Intelligence Technology Mapping exercise, conducted under Fusion for Energy's Technology Development Programme, is the first systematic European assessment of AI developments across the fusion engineering landscape. Its results are intended as a reference for the European public and private fusion AI community, and as a basis for the next phase of investment and collaboration.

The exercise was carried out through an online and an in-person workshop held in April and May 2025. 199 participants registered overall, with 90 attending in person and a peak of 137 connected online (section 4). In-person participants represented 31 public and private entities across 12 European countries. The workshop generated 60 candidate developments, of which 58 were carried forward into the final taxonomy. (section 5.2.1). Each was characterised across four primary domains:

- Design and multiphysics simulation
- Materials, manufacturing, Non-Destructive Testing (NDT), and assembly
- Plasma science
- System engineering, project management, and operations

Headline findings. The landscape is in active but early development. 71% of technologies sit at TRL 1–4, 15% at TRL 5–6, and 14% at TRL 7–9 (section 5.2.1). 80% have access to experimental datasets, 17% rely on synthetic data, and 3% report no available data; existing test facilities are available for 95% of technologies reporting facility status. 49 Technology Development Activities (TDAs) were identified across 32 of the 58 technologies (55%), and 51% of TDAs carry a likelihood of success above 80%. These figures should be read as signals of community focus rather than directly comparable measurements across technologies.

Strategic picture. Two findings dominate the strategic reading. First, 88% of the identified applications are general technologies with relevance beyond fusion, so progress in adjacent industries can directly accelerate fusion AI, and several proposed TDAs have established commercial analogues that warrant cross-domain benchmarking before commitment. Second, a small number of enabling TDAs (notably the Fusion Data Platform, the Accessible Database, and the Material Database) underpin many downstream applications and offer disproportionate leverage. The workshop reflects a predominantly public research perspective; private fusion companies operate under different conditions and priorities, and the findings should be read with that distinction in mind (sections 5.2.1 and 6). Data-quality limitations, concentrated non-EU dependency at the foundational AI and cloud layer, and a structural human-resource gap (section 6) shape the remaining risk picture. The implementation roadmap (section 5.2.1) organises the TDAs into four phases: immediate priorities, foundation building, high-TRL cross-sector activities, and strategic long-term developments.

Acknowledgements

F4E acknowledges the editorial and analytical contribution of Novatron Fusion Group (Sweden) to the present report. Thomas Wingate and Niveditha Veem reviewed the workshop materials, analysed the underlying data and dashboards, contributed editorial revisions, and provided methodological observations.

1 Introduction

1.1 Context

In 2024, Fusion for Energy launched a Technology Development Programme (TDP) as part of the implementation actions of its Industrial Policy. This TDP is dedicated to building and reinforcing European Fusion Supply chain capabilities for those technologies that are deemed to be critical for the future of commercial fusion. The programme requires the identification of key technologies to direct R&D contracts to European contractors.

Prioritizing and allocating funding opportunities requires a comprehensive review of the involved technologies on each major fusion technical domain. Doing this exercise in a collaborative way will enable stakeholders to identify which technologies are fundamentally needed (technology mapping) and when are they needed (technology roadmapping). A roadmap built through consensus of key stakeholders in the field can also serve as a powerful argument when seeking additional funding from national and international public and private investors.

To coordinate these efforts, Fusion for Energy has launched a technology mapping initiative uniting academia, research laboratories, industry, start-ups and the ITER Organization to develop a comprehensive technology development roadmap for the application of Artificial Intelligence in Fusion Technologies.

The outcome of this exercise will serve all stakeholders to guide their action in their respective domains, allowing an effective investment of resources. Given the fast evolution of technology, a periodical follow-up of the workshop outcome shall be assured in subsequent technology mapping exercises.

1.2 Artificial Intelligence Technology mapping


The scope of the second Technology Mapping exercise conducted under the TDP and the first dedicated to AI in fusion technologies covers the application of artificial intelligence in design and Multiphysics simulation, material development, manufacturing and non-destructive testing (NDT) including assembly, and plasma science and control, as well as system engineering, project management and operations.

The main associated event is a workshop held in April and May 2025 to generate most of the relevant data and provide an opportunity for participants to network and exchange knowledge.

This document provides a complete overview of the exercise, detailing the process and scope through a comprehensive technology breakdown, summarizing the meetings held and providing the resulting proposed technology development roadmap.

2 Technology mapping process

The technology mapping process consists of four stages.



2.1 Input report

In preparation of the exercise, staff from Fusion for Energy prepare a draft technology breakdown with input from ITER Organization colleagues, listing technologies of interest and grouping them functionally.

This breakdown, together with a brief description of each selected technology, is included in a draft input report (see section 3) for consultation by participants ahead of the first meeting (an online workshop).

2.2 Online workshop

The online workshop provides an opportunity for all participants in the technology mapping exercise to come together. It typically lasts 6 to 7 hours and follows the agenda outlined below:

- Welcome and introductory remarks
- The technology mapping process
- Introductory presentations about the field of interest & state of the art
- Networking opportunity between participants
- Brief overview of technology breakdown
- Joint review of the technology breakdown
- Explanation of the in-person workshop
- Survey feedback and wrap-up

The main output of the online workshop is a comprehensive state-of-the-art list of relevant technologies. This breakdown serves as the foundation for the technology mapping, which is the primary result of the initial workshop exercise. An updated version of the input report, including the revised technology breakdown, is provided to participants prior to the in-person workshop.

2.3 In-person workshop

The in-person workshop aims to provide a detailed characterization of the technologies outlined in the breakdown building on the results and discussions from the online workshop, including their prioritization and corresponding timeline.

The characterization of technologies is carried out in three steps for each technology:

- Agreement on the current Technology Readiness Level (TRL) (see Appendix 1 for definitions)
- Definition of the next step (e.g., analysis, prototype, testing, industrialization plan, etc.)
- Quantification of the technology characteristics (see Appendix 2 for the proposed list of characteristics to be evaluated)

Additionally, a timeline is developed, classifying what is needed and when for the technologies considered in the technology mapping. Typical timelines may span 5, 15, or 30 years, or be categorized into short, medium, and long-term periods.

The workshop is designed to be highly collaborative, with sessions that encourage participants to exchange ideas, build consensus, and provide feedback on specific interests and the mapping process itself. It also offers numerous opportunities for participants to share knowledge and establish partnerships. The workshop typically lasts one and a half days, with designated times for both formal and informal networking.

2.4 Final report

After the in-person workshop, staff from Fusion for Energy will compile the results into a final report, which serves as an evolution of the input report. This report will provide an overview of European capabilities in the field and include a proposed technology roadmap that outlines and prioritizes potential actions for the period leading up to the next review, typically occurring within 2 to 3 years.

Participants are given an opportunity to comment before the final version of the report is published.

3 Artificial Intelligence Breakdown

3.1 Artificial Intelligence overview

Artificial intelligence (AI) is transforming technology by enabling machines to perform tasks that traditionally required human intelligence, such as problem-solving, decision-making, and pattern recognition. From self-driving cars to medical diagnostics and advanced simulations, AI is revolutionizing industries by automating complex processes and improving efficiency.

One of the most powerful approaches within AI is machine learning, which allows computers to learn from data and improve their performance without explicit programming. Unlike traditional rule-based AI systems, machine learning algorithms adapt and evolve through experience, making them highly effective in processing large datasets and identifying patterns. There are different types of machine learning, including supervised learning (trained on labelled data), unsupervised learning (finding hidden structures in unlabelled data), and reinforcement learning (learning through trial and error). Machine learning is a driving force behind many AI applications, such as trend identification, image recognition, natural language processing, autonomous systems, and scientific research, accelerating advancements in automation, prediction, data analysis, and problem-solving.

For the purpose of this technology mapping, the term *AI technology* refers to applications where one or more of the following constitute the core capability: machine-learning models trained on data; statistical or probabilistic methods for inference, classification, or prediction beyond simple lookup; large language models or other generative models used for content generation or transformation; physics-informed or hybrid AI methods combining learned models with explicit physical or rule-based components; or autonomous-agent or reinforcement-learning systems making sequential decisions. Technologies that automate tasks using established non-AI methods, such as rule-based scripting, deterministic optimisation algorithms, or conventional commercial software, are not within scope unless an AI component substantially extends the capability beyond what those methods provide. This boundary is applied consistently throughout the report, but readers should note that several entries in the workshop output describe applications where the AI contribution is incremental relative to existing methods; the per-technology descriptions in section 3.4 and the analysis in section 5.2 reflect the workshop's framings rather than independent re-classification.

3.2 Technical breakdown of applications

Artificial intelligence can be applied at various stages of the development process, from the initial phases of design and simulation to operational applications.

This workshop focuses on four key topics in the development of fusion technologies:

- Design and simulation
- Materials, manufacturing, Non-Destructive Testing (NDT), and assembly
- Plasma science
- System engineering, project management, and operations

The **design and simulation** stage is essential for developing fusion reactor concepts and testing their viability before construction, as well as that of isolated components. Advanced simulations are used to model plasma behaviour, energy dynamics, and material interactions under fusion conditions. Engineers design the reactor components, while simulations help optimize these designs to ensure efficiency, stability, and safety during operation.

Materials, manufacturing, non-destructive testing (NDT), and assembly are all critical stages in developing reliable fusion reactors. Material development focuses on creating and testing advanced substances, such as radiation-resistant alloys and superconductors, that can endure the extreme conditions inside a reactor, including high temperatures, intense radiation, and mechanical stress. Manufacturing then transforms these materials into precise components like vacuum vessel sectors, plasma-facing elements, and superconducting magnets. During and after fabrication, NDT methods such as phased-array ultrasonic testing are used to inspect these complex parts without causing damage, ensuring they meet strict quality and safety standards. Finally, careful assembly brings together all subsystems, requiring high accuracy to achieve proper alignment, integration, and long-term reactor performance.

Plasma science is central to achieving stable fusion reactions. This stage involves managing and confining plasma at extremely high temperatures using magnetic fields. Early prediction of instabilities can allow operators to react promptly and stabilize the plasma, avoiding its phase-out and consequent loss. Researchers use advanced heating techniques to maintain the plasma's energy levels and employ real-time monitoring to adjust control systems and maintain stability.

System engineering, project management, and operations play a key role in turning fusion concepts into working energy systems. This part of the process involves making sure that all the different parts, like magnets, sensors, and heating systems, work well together. By spotting problems early, teams can take quick action to avoid delays or technical issues. Engineers use structured methods to keep the design on track, and project managers rely on real-time information to adjust plans, manage resources, and keep everything running smoothly.

Several cross-cutting challenges should be considered when driving future directions:

- **Data Scarcity and Quality:** Obtaining sufficient high-quality labelled data, especially failure data, for training robust AI models remains a challenge in these high-reliability sectors. Techniques like transfer learning, federated learning, data augmentation, and physics-informed AI are being explored.
- **Validation, Verification, and Qualification:** Rigorous processes are needed to validate AI performance and qualify AI-based systems for use in nuclear and fusion applications, meeting stringent industry standards.
- **Integration and Radiation Hardening:** Integrating AI systems with existing infrastructure and ensuring the reliability of hardware (sensors, processors) in radiation environments are practical hurdles.

The identified applications also span a broad spectrum in terms of the role AI plays. Some represent deep research challenges where machine learning contributes capabilities not achievable by other means. Others describe tasks where AI can accelerate or enhance processes that could also be addressed through conventional automation or general-purpose tools. Recognising this distinction is important for prioritising development effort and investment appropriately.

3.3 Priority Research Opportunities (PROs)

The following PROs are a historical external reference framework (Humphreys, et al. 2020) and serve as an external reference frame rather than a direct output of this workshop. AI advances, along with the urgency of need to bridge key gaps in knowledge for design and operation of reactors such as ITER, have driven planned expansion of efforts in ML/AI around the world, and especially within Europe. This is the aim of this report: to map all AI developments for fusion, to prioritize them according to the developments and needs of fusion energy in Europe and to conclude with a roadmap, to be updated with the required frequency.

Science discovery with Machine Learning

Applying ML to analyse extensive datasets can bridge theoretical gaps, accelerate hypothesis generation, and optimize experimental planning, enhancing understanding in areas like tokamak confinement and plasma-wall interactions.

Machine Learning-Boosted Diagnostics

Utilizing ML to extract maximum information from diagnostics, improve interpretability through data-driven models, integrate multiple data sources, and develop synthetic diagnostics to infer unmeasured quantities, thereby enhancing plasma state identification and regime classification.

Model Extraction and Reduction

Developing data-driven models to elucidate complex plasma behaviours, quantify uncertainties, and accelerate computational algorithms, facilitating faster simulations and improved comprehension of phenomena such as turbulent transport and plasma heating.

Control Augmentation with Machine Learning

Enhancing plasma control by improving control models, creating real-time data analysis algorithms for adaptive regulation, and optimizing plasma discharge trajectories, which is crucial for the effective operation of fusion reactors.

Extreme Data Algorithms

Creating methods for in-situ analysis and reduction of large-scale simulation data, and efficiently managing extensive experimental data, addressing challenges posed by the massive data volumes expected from future fusion experiments and simulations.

Data-Enhanced Protection

Developing algorithms to predict key plasma phenomena and system states, enabling real-time and offline monitoring and fault prediction, vital for preventing disruptions that could damage fusion devices.

Fusion Data Platform for Machine Learning Apps

Establishing a comprehensive system for managing, curating, and accessing fusion experimental and simulation data, supporting scalable application of ML/AI methods to fusion challenges.

3.4 Map of individual applications

The purpose of this section is to provide a brief overview of the applications discussed during the workshop. The application will be briefly explained, detailed information about each application including development needs and characteristics is included in the dashboards.

Design and simulation • AI/ML driven multi-objective optimization frameworks • Categorisation and clustering of CAD clashes • Component tagging for 3D-scanning • Generation of E3D or 3D drawings • Optimisation of magnetic fields using AI-guided inverse design algorithms • P&ID recognition for engineering data extraction and consistency • Quality control of CAD drawings • Active learning for data-efficient experimentation • Digital twins of fusion reactors • GenAI for code template generation • GenAI for synthetic data generation • Hybrid models and Physics-Informed Machine Learning for simulations • ML-driven surrogate modelling	Materials, manufacturing and NDT, and assembly • AI-guided discovery and development of new materials • Generation and curation of materials databases • Material acceptability prediction for a dedicated application • Prediction of behaviour and mechanical properties of new materials • AI-driven acoustic emission for non-metallic material inspection • AI-driven distortion prediction • AI-driven process optimization for advanced manufacturing • Automated digital RT data processing for weld inspection • Automated PAUT data processing for weld inspection • Real-time optimization of manufacturing parameters • Welding success rate prediction • AI-driven extrapolation of operating conditions from local testing • Automated mapping of metrology cloud points to component models • Detect modifications on drawings • Identification of most optimal assembly sequence • Instruments and transport virtual agent • Optimization of component positioning during assembly • Prediction of metrology data
Plasma science • Auto-labelling AI tools for plasma state datasets • Fusion data platform for machine learning applications • ML-accelerated pedestal MHD stability evaluations • Predictive modelling of plasma instabilities and disruptions • Time-resolved, physics-informed neural networks for tokamak total emission reconstruction • Trajectory planning enabled by ML/AI	System engineering, project management and operations • AI-driven validation of requirements • Chatbot for requirements rationalization • LLM-driven check of requirements propagation • AI-powered proposal support tool for project disruptions • Automated contract generation • Automatic status reporting from data sources • Prediction documentation review time • Prediction of contract outcome based on past performance • Prediction of tool usage • Process optimization tool • Risk register forecast prediction • AI-assisted experimental design • AI-driven diagnostic data integration • Automatic generation of EOM report • Dimensionality reduction for streaming data • Fusion specific documentation management • LLM-based prediction of document type • Predictive maintenance of components • Real-time operation of fusion device based on diagnostics/prediction • Recognition of abnormal events in the operation or state of a system

3.4.1 Application of AI in design and simulation

DESIGN

1. AI/ML driven multi-objective optimization frameworks

Multi-objective optimization is a long-established field of engineering analysis with mature commercial tooling (Ansys optiSLang, ESTECO modeFRONTIER, Dassault Isight, HEEDS, among others) deployed across aerospace, automotive, and structural design. Classical methods evaluate each candidate design through a complete simulation run, with metaheuristic search (NSGA-II (Deb, et al. 2002) and similar) exploring the Pareto frontier across thousands of evaluations. The AI/ML contribution is differentiable multi-physics simulation, where the simulator itself is implemented as an end-to-end differentiable function (for example tokamak transport simulators such as TORAX, or stellarator design tools such as DESC) and gradient-based optimisation can move through high-dimensional design space in orders of magnitude fewer evaluations than classical methods require. This is the genuinely recent (2024 onwards) development in the area.

2. Categorization and clustering of CAD clashes

A clash refers to a conflict or interference between components in a digital model, something that would physically collide, overlap or be incompatible. There are geometric (i.e, occupy the same space), clearance (i.e, safety or maintenance violations), functional (i.e, interfere with the functions), and sequencing (i.e, simultaneous when they shouldn't be) clashes. All four clash types can be detected by deterministic algorithms in commercial BIM software: geometric and clearance using bounding volumes around 3D objects or assemblies, functional using rule-based checks against component metadata, and sequencing by attaching schedule items to 3D assemblies (4D BIM). The AI contribution lies in handling the volume of detected clashes, where a typical project produces thousands of candidates the majority of which are irrelevant to safety or constructability. AI is applied in two ways: classifying the severity of each detected clash (critical, tolerable, or irrelevant) and routing it to the appropriate engineering discipline; and clustering related clashes to reduce the analyst review burden by grouping repeated or systematically related issues.

3. Component tagging for 3D-scanning

Identifying and labelling individual components in a 3D model by taking the source information and the schematics and diagrams as reference and matching the real-world scanned components to the reference components using AI. The tagged 3D model can now be connected back to the digital twin or used in maintenance or simulation. Automated component recognition from 3D scan point clouds is a commercially mature capability in industrial plant operations (Leica Geosystems, AVEVA, Hexagon Asset Lifecycle Intelligence among others) and is established research via semantic segmentation methods (PointNet++ and successors) (Qi, Su, et al. 2017) (Qi, Yi, et al. 2017). The fusion-relevant development effort is in training recognition models on fusion-specific component vocabularies, integrating output with fusion plant engineering data management systems, and qualifying tagging accuracy for safety-classified configuration records.

4. Generation of E3D or 3D drawings

Generate 2D or E3D (AVEVA E3D, the plant design platform specified for ITER) drawings from existing 3D models, or extract and apply revision intent from scanned 2D paper drawings including redline markups, using AI-powered tools that support natural language input or image uploads. Automated drawing generation from 3D plant models is a production capability in nuclear and petrochemical industries (AVEVA E3D); AI extraction of markup intent from scanned redlines is commercially deployed in construction and process-plant revision workflows (Datagrid, Pinnacle Infotech). The fusion-relevant development effort is in adapting these capabilities to fusion-specific drawing standards, symbol sets, and ITER/successors configuration-management requirements.

5. Optimisation of magnetic fields using AI-guided inverse design algorithms

AI-guided inverse design algorithms are used to find the global optimum in a complex design space, in order to produce a desired field. Instead of exhaustive forward search, these methods start with the target field configuration and apply gradient-based or surrogate-assisted optimisation to determine the optimal coil geometry or current distribution, accelerating the design of magnetic confinement systems. Open source frameworks for stellarator and tokamak coil optimisation (DESC, SIMSOPT, FOCUS) implement this approach and are active research tools in the fusion community. The AI contribution relative to classical methods is differentiable physics models that make gradient computation tractable across high dimensional coil-geometry spaces.

6. P&ID recognition for engineering data extraction and consistency

P&ID recognition is a mature application of computer vision and deep learning, with deployed commercial offerings (AWS, Microsoft Azure, Hexagon, Bentley, TCS Digitize-PID, Symphony AI and others) covering symbol detection, OCR, topology reconstruction, and output to standardized formats such as DEXPI. The technology is in production across oil and gas, chemicals, and pharmaceutical process industries. The fusion-relevant scope is to apply this established capability to ITER and successor-class plant documentation, integrating with fusion-specific equipment catalogues and connecting recognised P&ID content to plant engineering data management systems. The AI development effort required is in adaptation rather than invention: training symbol-recognition models on fusion-specific symbol sets and integrating the output with fusion plant data models.

7. Quality control of CAD drawings

This technology addresses a well-defined engineering challenge in environments where 2D drawings remain the primary contractual and configuration-management artefacts. Automated AI-based consistency checking between 2D documentation and 3D CAD models can improve design validation, detect geometry or attribute mismatches earlier, and reduce manual review effort, making the reported maturity level (TRL 8) credible in organisations with legacy documentation workflows. However, the relevance of this capability across the broader fusion sector may be limited. Modern engineering practices increasingly rely on Model-Based Definition (MBD) and BIM-oriented workflows, where the 3D model itself is the authoritative source of product information, thereby avoiding many 2D–3D consistency issues by design.

SIMULATION

1. Active Learning for data-efficient experimentation

Active learning uses AI to decide which experiment or simulation to run next, prioritizing the most informative data points at each step. Two distinct applications are relevant in fusion. First, adaptive experimental design: the AI model selects which physical experiment to conduct next based on current uncertainty, reducing the number of machine shots or test runs required to characterise a plasma regime or material behaviour. Second, adaptive simulation sampling: the model selects which simulation configurations to evaluate in order to build a surrogate model efficiently, reducing the total number of high-fidelity simulation runs. Both are established in adjacent domains (Bayesian optimisation for materials discovery, design-space exploration in aerospace and automotive), and both are directly applicable in fusion given the high cost of machine time and compute.

2. Digital twins of fusion reactors

AI-driven virtual replicas of fusion reactors that are continuously updated with real-world data. These digital twins enable real-time simulation and analysis of reactor performance, supporting design optimization, failure prediction, and improved operational efficiency. As a category, digital twins of fusion reactors span a wide range of sub-capabilities, each potentially drawing on distinct AI approaches: sensor fusion and real-time state estimation; ML surrogate models for plasma, thermal-hydraulic, and structural simulations; anomaly detection and predictive maintenance; operational parameter optimisation; and uncertainty quantification over coupled physics models. The TRL 2 workshop assignment reflects the immaturity of the integrated digital-twin concept for full fusion reactors; individual sub-capabilities span a considerably wider maturity range. In adjacent domains AI-integrated digital twins are commercially deployed at scale (e.g. GE Predix in aviation and heavy industry), for fusion, integrated digital twins remain research-grade, as illustrated by the APS-DPP 2024 mini-conference on digital twins for fusion (Schissel, Nazikian en Gibbs 2025).

3. GenAI for code template generation

Use of Generative AI to automatically create code templates for building machine learning-based surrogate models. Rather than operating within the simulation itself, this technology targets the development process. It provides structured starting points that reduce the time and effort needed to implement physics-informed or data-driven simulation approximations. General-purpose AI code generation tools (GitHub Copilot, Amazon Q and others) and domain-specific physics-ML frameworks (NVIDIA PhysicsNeMo, DeepXDE and others) already provide generic surrogate model scaffolding. The fusion-specific contribution is templates that encode fusion-physics structure: boundary conditions and conservation constraints relevant to plasma and thermal-hydraulic simulations, and compatibility with fusion simulation code data formats and interfaces (e.g. PROCESS, SOLPS, JINTRAC). Without this domain encoding, generic code generation tools produce syntactically correct but physically unconstrained starting points.

4. GenAI for synthetic data generation

Application of Generative AI to create synthetic data that enriches existing datasets or fills gaps in training data. This technology supports the simulation pipeline rather than operating within it. It addresses the data scarcity that limits how well AI-based simulation models can be trained, and is particularly valuable in fusion where experimental datasets are scarce and imbalanced. Synthetic data generation is commercially mature in adjacent domains, with platforms such as Gretel, MOSTLY AI, and NVIDIA Omniverse Replicator deployed for computer vision inspection, manufacturing quality control, and autonomous vehicle perception. In fusion, the relevant applications include generating synthetic plasma diagnostic signals to cover measurement regimes not accessible to physical sensors, augmenting limited plasma discharge datasets for disruption prediction training, and generating physics-consistent synthetic material data to supplement sparse irradiation experiment results.

5. Hybrid models and Physics-Informed Machine Learning for simulations

Integration of physical laws with machine learning to build models that are both data-driven and grounded in scientific principles. Here AI is embedded within the simulation framework, replacing computationally expensive components while physical constraints ensure predictions remain physically valid. This makes Physics Informed Machine Learning particularly suited to fusion, where purely data-driven models struggle due to limited training data and difficult extrapolation. Production-grade frameworks for physics-informed ML are commercially available (NVIDIA PhysicsNeMo, MATLAB PINN toolbox) and deployed in oil and gas, energy, and biomedical engineering.

Published fusion studies have demonstrated PINN-based solution and reconstruction approaches for reduced plasma-physics problems, hybrid databases combining experimental discharge data with validated simulation outputs for plasma management training, and physics-informed approaches to multi-diagnostic equilibrium reconstruction (Raissi, Perdikaris en Karniadakis 2018) (Kumar, et al. 2023).

6. ML-driven surrogate modelling

Use of machine learning to create fast, computationally efficient surrogate models that approximate complex physics simulations such as plasma dynamics (Rodriguez-Fernandez, Howard en Candy, Nonlinear gyrokinetic predictions of SPARC burning plasma profiles enabled by surrogate modeling 2022) (Rodriguez-Fernandez, Howard en Saltzman, et al. 2024), turbulence (Clavier, et al. 2025), and material interactions, enabling quicker analysis and design iterations. Unlike Physics Informed Machine Learning, a surrogate substitutes the full physics simulation entirely, trained once on high-fidelity simulation outputs, it then replaces the simulator for subsequent analyses, delivering results orders of magnitude faster. ML surrogate modelling for physics simulation is commercially deployed in aerospace and engineering (Siemens Simcenter, NVIDIA PhysicsNeMo). In fusion, production-quality surrogates for plasma transport and MHD simulation already exist, with speedups of multiple orders of magnitude over the codes they replace. The fusion-specific development challenge is not surrogate methodology, which is well-established, but training data availability, validation against experimental data, and integration with existing tokamak modelling frameworks.

3.4.2 Application of AI in materials, manufacturing, NDT, and assembly

MATERIALS

1. AI-guided discovery and development of new materials

Use of AI and machine learning to accelerate the identification and design of new materials with targeted properties. These approaches enable rapid screening of material candidates and prediction of performance, reducing experimental cycles and development time. In the fusion context, the prioritised properties typically include high-temperature mechanical strength, low neutron activation, irradiation tolerance, and plasma-facing material survivability. AI approaches commonly applied in this area include high-throughput density functional theory calculations, generative models for candidate alloys, machine-learning interatomic potentials, and active learning-driven experimental design. Materials informatics is commercially mature in adjacent domains (e.g. Citrine Informatics in batteries and polymers), and ML interatomic potentials for irradiation-damage modelling are an active research area for nuclear and fusion structural materials. The principal qualification challenge is that computationally proposed materials require extensive experimental validation, including neutron irradiation testing, before they can be accepted under applicable nuclear codes and standards.

2. Generation and curation of materials databases

Development and collection of comprehensive materials databases to support AI-driven discovery, modelling, and simulation. These databases provide structured, high-quality data essential for training predictive models and guiding experimental efforts. In the fusion context, relevant database categories include nuclear activation and transmutation data (e.g. FENDL, TENDL libraries), structural material property databases covering fusion-relevant alloys (EUROFER97, ODS steels, tungsten alloys), irradiation effects datasets, and thermophysical property data. The AI contribution operates at two levels: high-throughput DFT and molecular-dynamics simulation for automated generation of training data, addressing the sparsity of experimental fusion materials data; and ML models trained on the resulting databases for property prediction, candidate material screening, and informed extrapolation with appropriate uncertainty quantification. The general materials informatics infrastructure pattern (FAIR open databases with integrated AI analysis tooling) is well established in adjacent domains, e.g. NOMAD (Draxl et al. 2018) and the Materials Project (Jain, et al. 2013); fusion-specific equivalents remain more siloed and less AI-integrated. Key challenges include experimental validation of computationally generated entries, cross-institutional data standardisation, and qualification of AI-predicted properties within applicable nuclear codes and standards.

3. Material acceptability prediction for a dedicated application

Use of AI and machine learning to predict whether a material meets the acceptance criteria for a specific nuclear application, based on its measured properties and processing history. Nuclear code associations (standards bodies such as ASME and RCC-MRx that define qualification requirements for nuclear components) have begun to explore frameworks for accepting AI as a recognised digital qualification tool, with the potential to reduce the cost and cycle time of material qualification while maintaining the required confidence levels.

4. Prediction of behaviour and mechanical properties of new materials

Use of AI and machine learning to predict material behaviour and mechanical properties, enabling faster screening and less physical testing. In fusion, the focus is on predicting strength, ductility, creep, and fracture toughness of candidate structural and plasma-facing materials (EUROFER97 variants, ODS steels, tungsten alloys) under irradiation and high-temperature conditions where data are sparse. Composition- and process-based property prediction is commercially deployed (e.g., Citrine Informatics, Ansys Granta MI), while fusion-specific irradiation-damage prediction using ML interatomic potentials remains an active research area.

MANUFACTURING & NDT

1. AI-driven acoustic emission for non-metallic material inspection

Application of AI to improve Acoustic Emission (AE) techniques used in pressure testing and in-service inspections of non-metallic materials such as composites and concrete. AI analyses AE signals to predict component integrity and enable early detection of potential failures. AI-augmented AE is an advancing commercial NDT capability deployed in pressure vessel and composite inspection across aerospace and petrochemical industries (Ghaznavi, et al. 2026). The fusion-relevant scope is inspection of SiC/SiCf composite components (breeding blanket and first-wall structures) and bio shielding concrete, where AE methodology from aerospace composite

pressure vessel inspection is transferable but requires qualification for fusion-specific damage modes and irradiation conditions.

2. AI-driven distortion prediction

Manufacturing-induced distortion in welded, cast, and additively manufactured components is conventionally predicted by thermo-mechanical finite-element analysis (FEM). For fusion engineering this is most relevant to welding distortion in large vessel and blanket structures, and to layer-by-layer thermo-mechanical distortion in additively manufactured components. AI adds value over FEM in three ways: (1) machine-learning surrogate models trained on FEM outputs reproduce predictions in seconds to milliseconds, enabling design-space exploration and parameter optimisation that FEM run-times do not support; (2) ML models can ingest design parameters and process specifications directly, making distortion prediction accessible to design engineers without FEM expertise; (3) physics-informed neural operators can run in under 150 ms, enabling closed-loop control of additive manufacturing processes by integrating with in-process thermal and geometric sensors. Each of these is established in the published literature for non-fusion applications and is directly transferable to fusion components.

3. AI-driven process optimization for advanced manufacturing

Application of AI to optimize manufacturing processes such as Hot Isostatic Pressing (HIP) and additive manufacturing. This includes optimizing build strategies, predicting material properties, and detecting defects during production to improve quality and efficiency. The scope of this entry is pre-process and in-process optimisation of the manufacturing recipe: AI models trained on historical process data to recommend HIP cycle parameters (temperature, pressure, dwell time) for a target density and microstructure, and to optimise AM build orientation, support structures, and scan strategies to minimise defects. This is distinct from the real-time sensor-feedback control covered in the adjacent entry on real-time optimisation of manufacturing parameters, and from the post-prediction of distortion covered in the AI-driven distortion prediction entry.

4. Automated digital RT data processing for weld inspection

Automated AI-driven defect detection in digital radiographic testing of welds is mature commercial technology, with deployed offerings from inspection-services vendors and specialised tooling providers (Trueflaw, Yxlon, Visiconsult, North Star Imaging, Bureau Veritas and others) operating in aerospace, automotive, shipbuilding, and conventional nuclear industries. Public benchmark datasets (RIAWELC, GDXray) and standardised industrial metrics including Probability of Detection (POD) support objective performance comparison. The fusion-relevant scope is qualification of these capabilities for safety-critical nuclear-classified welds in vessel, divertor, and breeder-blanket components, where regulatory codes (ASME Section V, ISO 17636, EN 13068, RCC-MRx and other nuclear-specific qualification frameworks) currently treat AI-assisted RT interpretation as an active rather than settled question. Adaptation to fusion-specific component geometries and weld procedures is the development work required.

5. Automated PAUT data processing for weld inspection

AI-driven automated processing of Phased Array Ultrasonic Testing (PAUT) data for welds of components at supplier sites. This enables precise identification of defects, including their exact location and depth. PAUT itself is a mature inspection technique deployed across aerospace, nuclear, and oil-and-gas industries. AI-augmented automated defect recognition on PAUT data is a fast-moving area: a substantial commercial market exists for AI-based weld inspection, but published reviews characterise machine-learning defect classification on PAUT data as advancing rather than fully mature (Na, et al. 2025). The TRL 8 workshop assignment is best read as reflecting the deployed maturity of PAUT itself; the fusion-relevant development effort is in adapting and qualifying defect-detection models for the specific material combinations (EUROFER97, tungsten, copper alloys), weld geometries, and acceptance criteria of fusion components.

6. Real-time optimization of manufacturing parameters

Real-time AI optimisation of manufacturing parameters using sensor feedback (temperature, acoustic emissions, optical monitoring) is an established industrial capability deployed at scale across oil and gas, semiconductors, automotive, and chemicals industries (with production-grade vendors including C3 AI, Imubit, and digital-twin platforms). Fusion relevant processes where this capability would be applied include: precision welding of vacuum vessel and divertor components (overlapping with the welding-specific entries in this section); Hot Isostatic Pressing and additive manufacturing (overlapping with the entry on AI-driven process optimization for advanced manufacturing); large-component machining; and deposition processes for plasma-facing component coatings. The fusion-AI development effort is in adapting commercially mature optimisation methods to the specific physics, materials, and qualification requirements of these processes. Strong IP positions held by industrial vendors are noted in the underlying workshop data as a constraint on technology transfer.

7. Welding success rate prediction

Use of AI to predict the success or failure of Electron Beam (EB) and Tungsten Inert Gas (TIG) welds for components. This helps improve weld quality and reduce defects through early identification of potential issues. ML-based weld quality prediction is an advancing capability in precision manufacturing. A directly relevant fusion-specific implementation has been published: a machine learning model for EB weld quality prognosis applied to actual ITER vacuum vessel manufacturing data (Ortiz de Zuniga, et al. 2026) achieves 83% accuracy for weld conformity prediction and over 90% for defect root cause and length, demonstrating the approach on nuclear-grade components with strict quality requirements.

ASSEMBLY

1. AI-driven extrapolation of operating conditions from local testing

Machine learning models use results from localized or small-scale tests to predict system behaviour under full-scale or varied operating conditions, reducing the need for extensive testing and accelerating design validation. In the fusion assembly context, a representative application is using coupon-level or mock-up test data to predict the response of full-scale first-wall or divertor components under combined plasma-heat and electromagnetic loading where full-scale pre-installation testing is impractical. No commercial product was identified that specifically delivers this capability; multi-fidelity ML for scale bridging in materials applications is an active research area. A known limitation is that ML models trained on small-scale data may not generalise reliably to full-scale regimes outside the training distribution, making uncertainty quantification an essential component of any qualification approach.

2. Automated mapping of metrology cloud points to component models

AI algorithms match 3D point cloud data from metrology scans to digital component models, enabling automated verification of geometry, detection of deviations, and support for quality control during manufacturing and assembly. Automated point cloud registration and deviation

analysis is commercially mature in precision manufacturing metrology (Hexagon PC-DMIS, ZEISS GOM Inspect, Verisurf), deployed across aerospace and automotive production. The fusion-relevant development effort is in adapting these tools to the scale and geometric complexity of fusion components (vessel segments, blanket modules, divertor cassettes) and integrating deviation outputs with fusion configuration management and acceptance workflows.

3. Detect modifications on drawings

AI-mediated feedback of as-built information to design records is well established in the construction sector and is commonly delivered through three layers: (i) digital redlining within CAD/document-management platforms (e.g., AutoCAD, Bluebeam, Autodesk Construction Cloud, Procore); (ii) automated interpretation of redline mark-ups using computer vision combined with large language models; and (iii) scan-to-model change detection, in which laser-scan point clouds are compared directly against as-designed models, bypassing manual mark-ups. The workshop data characterise this application primarily as a digitisation pathway for legacy organisations that still capture changes through paper or PDF redlines. Accordingly, the relevance for Fusion for Energy is institution-specific: it is most applicable where F4E (or its industrial partners) must manage configuration change capture in projects that continue to rely on redlined drawings as the primary as-built record. By contrast, it has limited value for fusion organisations whose construction workflows capture changes natively within a BIM or model-based definition environment.

4. Identification of most optimal assembly sequence

Using AI to select the assembly sequence that maximizes success rates while minimizing costs, improving overall efficiency and reliability of the assembly process. Assembly sequence optimisation is an established capability in digital manufacturing planning, with commercial tooling for sequence simulation and validation in automotive and aerospace (Siemens Tecnomatix Process Simulate, Dassault DELMIA); these platforms use simulation-based optimisation with emerging AI features rather than dedicated ML-driven sequence selection. For fusion, the relevant context is the constrained sequencing of large, heavy, one-of-a-kind components in confined spaces with limited re-work opportunity, where the cost of a mis-sequenced installation is very high. A long-term research direction could be to extend the technology to closed-loop co-optimisation of assembly sequence and component design, where the sequence planner feeds tolerance and accessibility constraints back to the design stage to improve overall assemblability.

5. Instruments and transport virtual agent

Agentic LLM systems for booking, scheduling, and request management workflows are mature commercial capabilities, available through general-purpose enterprise platforms including Microsoft Copilot, Google Gemini, OpenAI ChatGPT, and Anthropic Claude (the workshop's own underlying data names these as the relevant platforms for this entry). Native integration with Microsoft 365 (Teams, Outlook, SharePoint) and Google Workspace is provided directly by the platform vendors. The fusion-relevant scope is deployment of these commercial agents within fusion institutional procurement and operations workflows where regulatory, audit-trail, and data-handling constraints apply. The 75% workshop voter agreement on "regulatory or ethical concerns" as an obstacle reflects this: the technology challenge is in agent-deployment governance and integration rather than in AI capability development.

6. Optimization of component positioning during assembly

Application of Bayesian neural networks to optimize the positioning of components, reducing uncertainties and minimizing the need for on-site adjustments during assembly. This entry is distinct from assembly sequence optimisation: sequence planning determines the order of operations, while positioning optimisation determines the precise spatial placement of each component during installation, explicitly accounting for measurement uncertainty. The

Bayesian approach propagates positional uncertainties through the placement decision, enabling the system to recommend placements that minimise expected fit-up error across a stack of tolerance interfaces. Laser tracker-guided component positioning is deployed in large-structure aerospace assembly (Hexagon, Leica Geosystems); the Bayesian ML layer for uncertainty-aware placement optimisation is research-grade.

7. Prediction of metrology data

Compare the measurement of physical components that are already installed, but are difficult to reach, to the measurements in the CAD models. The measurement of a limited number of measured components should be extrapolated using ML to predict the measurements of the unreachable components. The underlying approach is spatial interpolation: inferring geometry at inaccessible locations from measurements at accessible points, with Gaussian Process Regression (kriging) the established statistical framework offering built-in uncertainty quantification. Analogous applications exist in structural health monitoring and aerospace large-structure assembly metrology. The fusion-relevant context is any fusion device assembly where internal interfaces become physically inaccessible once installation progresses, making predictive metrology a potential qualification tool for alignment verification across both public programmes and private fusion companies.

3.4.3 Application of AI in plasma science

1. Auto-labelling AI tools for plasma state datasets

AI-driven tools that automatically label plasma states in datasets, enhancing labelling accuracy and consistency to improve the quality of training data for machine learning models. General-purpose automated annotation tools are commercially mature but are designed for visual or structured data, not physics-derived label categories. Plasma state labelling requires algorithms that can classify confinement regime, MHD stability, disruption precursor type, and other categories from multi-diagnostic signal data, a domain-specific challenge not covered by general annotation platforms.

2. Fusion data platform for Machine Learning applications

Development of a centralized data platform to collect, manage, and provide fusion-related datasets. This platform supports machine learning applications by enabling efficient data access, integration, and preprocessing to accelerate research and development.

3. ML-accelerated pedestal MHD stability evaluations

Use of surrogate modelling to accelerate predictions of growth rates for magnetohydrodynamic (MHD) equilibria, enabling faster stability assessments of the plasma pedestal. This is a well-developed research area within fusion with multiple published implementations.

4. Predictive modelling of plasma instabilities and disruptions

Development of neural state-space models to predict plasma dynamics during ramp-down phases, enabling better understanding and mitigation of instabilities and disruptions. Disruption prediction is one of the most mature ML application areas in fusion research, with multiple implemented approaches across tokamaks and benchmarking frameworks now available (e.g. DisruptionBench, spanning Alcator C-Mod, DIII-D, and EAST) (Spangher, et al. 2025) (Seo, et al. 2024) (Kates-Harbeck, Svyatkovskiy en Tang 2019).

5. Time-resolved, physics-informed neural networks for tokamak total emission reconstruction

Development of physics-informed neural networks as inverse problem solvers to integrate multi-diagnostic and multi-physics data. This approach enables tomographic reconstruction to accurately assess line-averaged plasma radiation emission in tokamaks. The technique adapts neural network tomographic reconstruction methods established in medical imaging to the specific constraints of fusion bolometric diagnostics and has been demonstrated on JET data for transient radiation events including ELMs, MARFEs, and core radiation anomalies.

6. Trajectory planning enabled by ML/AI

Use of reinforcement learning combined with hybrid physics and machine learning models to safely ramp down plasma current, avoiding disruption-related limits and ensuring stable operation. This is a rapidly maturing research area with demonstrated implementations at TCV, HL-3, and design work targeting SPARC, applying constraint-aware reinforcement learning methods established in robotics and process control to the specific safety requirements of tokamak operation (Degrave, et al. 2022).

3.4.4 Application of AI in system engineering, project management, and operations

SYSTEM ENGINEERING

1. AI-driven validation of requirements

AI models are used to automatically review technical requirements for clarity, consistency, completeness, and feasibility, helping identify ambiguities or conflicts early in the design process. LLM-based requirements quality checking is a commercially available capability (IBM DOORS Next AI features, Jama Connect Advisor, Visure Solutions Quality Analyzer) deployed in aerospace, automotive, and defence programmes. The fusion-relevant development effort is primarily configuration and integration: adapting requirements review models to fusion-specific terminology and system architecture, and qualifying AI-generated outputs within the configuration management and design assurance processes applicable to each organisation, which vary across public programmes, research institutions, and private fusion companies.

2. Chatbot for requirements rationalization

Development and training of a chatbot on project requirements to assist in rationalizing and clarifying them, speeding up the requirements analysis process. This technology could be delivered by the same implementation as the previous technology: a system trained or configured on a requirements corpus with LLM-based query and analysis capabilities provides both quality validation and interactive rationalization.

3. LLM-driven check of requirements propagation

Large Language Models (LLMs) are used to trace and verify the consistent propagation of requirements across system specifications, designs, and documentation, ensuring alignment and reducing the risk of overlooked dependencies. Requirements traceability is a feature of established commercial platforms (IBM DOORS Next, Jama Connect, Visure Solutions); LLM-based automated propagation checking across document hierarchies is an advancing capability with demonstrated applications in aerospace and automotive system engineering. The fusion relevant development effort is the same as for adjacent requirements entries: adapting to fusion-specific terminology, system architecture, and the configuration management processes applicable to each organisation.

PROJECT MANAGEMENT

1. AI-powered proposal support tool for project disruptions

A proactive tool designed to identify optimal project measures and decisions during disruptions by analysing contract types, schedules, costs, and risks of ongoing projects to support effective project management. AI-driven disruption management and decision support is commercially available for complex projects (e.g. nPlan, trained on over 750,000 historical schedules; Nodes & Links for GenAI-based schedule risk analysis), with the strongest adoption in construction and infrastructure megaprojects. The fusion-relevant scope is applying these capabilities to the specific characteristics of fusion programmes: long timelines, complex international supply chains, multi-party contract structures, and regulatory environments.

2. Automated contract generation

AI-mediated contract drafting and generation is a mature commercial capability deployed in enterprise Contract Lifecycle Management platforms (Ironclad, DocuSign CLM, ContractPodAi, Lexion, Spellbook, Harvey AI, Evisort, Sirion). Custom-skill or custom-instruction patterns on general-purpose LLMs reach functional capability with modest implementation effort, and the workshop's underlying data identifies major AI vendors and startups as likely sources of the technology. The fusion-relevant scope is deployment within EU public-sector and international-treaty procurement contexts, where AI-generated contract

terms need traceability of which clauses were inserted by the model, with auditable rationale and qualification under applicable procurement codes. The 60% workshop voter agreement on “lack of explainability” as the dominant obstacle reflects this requirement. The technology challenge is in agent-deployment governance and qualification rather than in AI capability development.

3. Automatic status reporting from data sources

Automatically compiles and summarizes project or system status by extracting key metrics and updates from various data sources, reducing manual reporting effort and improving real-time visibility. Automated AI status reporting is a commercially mature capability, with AI-generated project status summaries now a standard feature in enterprise project management tooling (e.g. Microsoft Copilot for Project). The fusion-relevant development effort is integration with existing programme management systems and qualification of AI-generated outputs within the audit-trail and traceability requirements of regulated fusion programme environments.

4. Prediction documentation review time

This application uses historical document-tracking data (e.g., workflow timestamps, document type, and review-routing metadata) to estimate the time required for future reviews, particularly where multiple reviewers and iterations are involved. The objective is to improve schedule realism without changing reviewer deadlines. At a basic level, this is a regression problem that can be addressed with established non-AI approaches (e.g., three-point estimation, stratified historical averages, and standard project-management analytics). AI adds value when the available data volume and feature complexity justify machine-learning models that capture interactions between document type, reviewer roles, project phase, and contractual or configuration context. One practical constraint is data availability: the longitudinal datasets required to train robust models are typically available to mature organisations such as F4E and ITER but are often not yet present in younger private fusion programmes with limited historical depth.

5. Prediction of contract outcome based on past performance

Identifying several parameters available from 1000+ past contracts to quantify the quality of preparation and negotiations and predict likely contract outcome (e.g. cost overrun) at contract signature. ML-based contract outcome and cost overrun prediction is an established research capability in construction and infrastructure procurement, applied to comparable datasets of thousands of historical contracts.

6. Prediction of tool usage

Machine learning and predictive analytics can help forecast future demand for physical tools, test equipment, facilities, and other constrained resources in large engineering programmes. By analysing historical usage data, programme schedules, and workload patterns, ML models can predict demand more accurately than manual planning methods, reducing bottlenecks and improving resource allocation. In the fusion context, relevant applications include scheduling of shared test facilities, allocation of specialised tooling for assembly operations, and planning of human resource capacity across concurrent workstreams.

7. Process optimization tool

An AI-driven tool designed to help process owners enhance efficiency by gathering data from internal management systems, analysing value-added steps, bottlenecks, and iteration cycles to identify opportunities for process improvement. The general capability has an established commercial baseline in AI-augmented process mining (vendors such as Celonis, UiPath, Software AG ARIS, IBM, Microsoft Power Automate Process Mining), and LLM-augmented features for natural language process queries and AI-suggested optimisations have been added across the category since 2023. The fusion-relevant development effort lies in event-log instrumentation of fusion-institutional process and operations stacks, integration with existing enterprise systems, and qualification within data handling and audit-trail constraints.

8. Risk register forecast prediction

An automated system to consolidate and review existing risks, while identifying potential new risks from multiple sources to enhance overall risk management. In addition, predictive analytics and AI models can be used to forecast the likelihood, timing, and potential impact of risk materialisation based on historical project data, risk trends, and contextual project indicators. AI-automated risk consolidation and identification is a commercially deployed capability in enterprise GRC platforms (IBM OpenPages 9.1, ServiceNow IRM), with Watson NLP-based automated risk classification and issue creation now in production use.

OPERATIONS

1. AI-assisted experimental design

AI is used to propose the most effective experimental setups by analysing prior data, optimizing parameter selection, and reducing the number of required tests to achieve meaningful results efficiently. AI-assisted experimental design using Bayesian optimisation is commercially mature in pharma and materials science (self-driving laboratories at Novartis, Citrine Informatics and others), where it reduces the number of physical experiments needed to reach a target outcome. The fusion-relevant application is to plasma experiment planning: selecting discharge parameters to maximise scientific output per unit of machine time, which is scarce and expensive at major tokamak facilities.

2. AI-driven diagnostic data integration

Plasma state in fusion device operation is inferred from a heterogeneous set of diagnostic instruments operating at different sampling rates, resolutions, and signal-to-noise characteristics. ML approaches enable real-time integration of these diagnostics into a unified plasma state estimate, robust to missing or degraded inputs. Multi-sensor data fusion is a mature discipline in adjacent domains: target-tracking systems in defence and air traffic control have followed the JDL data-fusion reference model since the late 1980s, robotic SLAM routinely fuses inertial, LIDAR, and vision streams, and autonomous-driving stacks (Mobileye, Waymo) integrate camera, radar, and LIDAR at production scale. In fusion, the long-standing reference is the Integrated Data Analysis framework developed at IPP for ASDEX Upgrade (Fischer, et al. 2010), which applies Bayesian inference over multiple diagnostics to a joint plasma-state estimate. ML-based extensions are an active research area, illustrated by (Yang, et al. 2026), which produces a unified embedding across 88 diagnostic channels and supports downstream control and monitoring.

3. Automatic generation of End-of-Manufacturing report

Automatic generation of End-of-Manufacturing reports (EMR) using AI to consolidate manufacturing, inspection, non-conformance, metrology and traceability records from multiple sources into a structured qualification dossier. AI can support information extraction, consistency checking, evidence linking and report drafting, reducing manual effort while improving completeness and traceability.

4. Dimensionality reduction for streaming data

Techniques that reduce data complexity and identify the most relevant features in real-time streaming data, enabling more efficient processing and analysis. Dimensionality reduction for high-rate industrial data streams is an established technique in condition monitoring and industrial IoT, where sub-100 ms latency constraints require significant compression before inference. In fusion, the challenge is the heterogeneous, high-channel diagnostic environment: (Yang, et al. 2026) illustrates the approach, compressing 88 diagnostic channels into a compact embedding while maintaining performance across control and monitoring tasks.

5. Fusion specific documentation management

Implementation of chatbots to enhance the speed and efficiency of searching manuals and databases within documentation management systems. Includes access control measures to

ensure sensitive information is only disclosed to authorized users. The general capability (retrieval-augmented chatbots over an indexed document corpus with role-based access control) has an established commercial baseline in enterprise platforms including Microsoft 365 Copilot, Glean, Atlassian Rovo, Coveo, OpenAI ChatGPT Enterprise, Anthropic Projects, and Google Vertex AI Search. The fusion-relevant development effort lies in deployment within F4E and ITER document repositories, integration with existing access-control regimes, and qualification within data-handling and information-classification constraints.

6. LLM-based prediction of document type

Large Language Models (LLMs) analyse document content to automatically classify and predict document types, improving organization, retrieval, and processing efficiency in document management systems. Document classification has an established commercial baseline (pre-LLM forms in Microsoft SharePoint document understanding, IBM Watson Discovery, Google Document AI, Azure Form Recognizer; LLM-based classification standard in Microsoft Share-Point Premium, M-Files, Hyland, OpenText, Box since 2022-2023). The fusion-relevant development effort lies in adapting classification taxonomies to fusion-engineering document categories (technical specifications, design reviews, quality records, contractual documents, regulatory filings) and deployment within F4E and ITER content repositories.

7. Predictive maintenance of components

AI models trained to detect early signs of anomalies or failures, such as monitoring wall components with infrared cameras to predict overheating from plasma exposure. Additionally, machine learning analyses trends in non-destructive testing (NDT) and operational load data to estimate the remaining useful life of components, supporting aging management and maintenance planning. The entry covers two distinct applications: real-time anomaly detection from live sensor streams (online monitoring) and remaining useful life estimation from accumulated inspection and load history (offline prognostics). Both are commercially mature in adjacent domains including aviation (GE Predix, Rolls-Royce IntelligentEngine) and nuclear fission; the fusion-specific development effort lies in adapting these capabilities to plasma-facing and structural components operating under neutron irradiation and high heat flux conditions for which limited operational lifetime data currently exists.

8. Real-time operation of fusion device based on diagnostics

AI-driven systems use live diagnostic data to monitor and adjust fusion device parameters in real time, optimizing performance, ensuring safety, and preventing disruptions during operation. Real-time feedback control from fused sensor streams is commercially mature in adjacent domains: advanced process control with Kalman-filter state estimators is standard in oil refining, petrochemicals, and power generation (AspenTech DMC3, Honeywell Profit Controller, ABB Predict & Control), and inertial-sensor fusion underpins flight-control law execution in commercial aviation. The fusion-specific contribution is replacing or augmenting classical observers with ML policies that map heterogeneous diagnostic signals directly to actuator commands. The reference demonstrations are Degraeve et al., “Magnetic control of tokamak plasmas through deep reinforcement learning” (Degraeve, et al. 2022), trained and deployed on TCV; and Seo et al., “Avoiding fusion plasma tearing instability with deep reinforcement learning” (Seo, et al. 2024), demonstrated on DIII-D. Both remain research-grade rather than operational practice.

9. Real-time operation of fusion device based on prediction

AI models predict future plasma behaviour and device states in real time, enabling proactive adjustments to maintain optimal performance and prevent disruptions during fusion reactor operation. Prediction-based control using data-driven forward models is commercially established in oil refining and chemical processing (ABB MPC). In fusion, a neural state-space model trained on TCV discharge data predicts ramp down dynamics to generate disruption-avoiding trajectories (Wang, et al. 2025); the HL-3 real-time controller similarly relies on a

data-driven dynamics model as its RL training environment (Wu, et al. 2025).

10. Recognition of abnormal events in the operation or state of a system

By using sensors that regularly or continuously monitor the condition of equipment of products in an experimental setting or in a production environment, a model can be trained to recognize abnormal events. This can help protect hardware and support quality assurance, either by raising an alarm to allow for human intervention, or by automatically activating countermeasures to restore the system to its nominal state. AI-based anomaly detection is commercially mature in manufacturing, energy, and process industries; unsupervised approaches are standard given the rarity of labelled fault events in well-managed facilities.

4 Summary of the workshop

In total, 199 people registered for participation to the Artificial Intelligence Technology Mapping Workshop. The online session of the workshop reached a peak of 137 participants, while 90 participants attended the in-person session. 31 public and private entities were represented, coming from 12 different countries. Fusion for Energy wishes to thank all participants for their inputs during and after the workshop.



Figure 1: Logos of participating entities (Fusion for Energy excluded)




Figure 2: Geographical repartition of the participants of the in-person workshop

Details of the workshop can be found on the [event web page](#). The agenda and outputs including presentations, documents, and recordings are also available there.

5 Outcome: Technology Roadmap

5.1 Technology dashboards

During the in-person workshop and in the process of preparing this report, a lot of valuable data was collected into a database. For every application, the following data is now available:

- TRL
- Dataset characteristics
- Broader market potential
- Major obstacles to the technology
- Existing and additional test facilities
- (Non)-European entities involved
- Technology Development Actions

This data has been arranged into a dashboard for each of the technologies:




Figure 3: Typical technology dashboard

All technology dashboards are available in Appendix 3: Technology dashboards. The dashboards are a view of the database at the time of publishing this document. The database will be updated regularly, and Appendix 3 may be re-published as necessary. We encourage the community to communicate their updates to their Fusion for Energy contact. In the future, we may publish this data for interactive consultation on the Fusion for Energy website.

5.2 Overview of the AI in fusion landscape in the EU

5.2.1 Technology and TDA analysis

The analysis presented in this section covers 58 technologies. Of the 60 initially assessed during the workshop, two were not carried forward into the final taxonomy following the triage process. For technologies where specific data fields were not populated during the workshop, those fields are excluded from the relevant calculations. All percentages therefore reflect the subset of technologies for which data was available. Throughout the report, individual entries are referred to as “applications” when described as specific use cases (see section 3.4) and as “technologies” within the TDA framework used for statistical analysis; the terms are used interchangeably in this context.

The Technology Readiness Level (TRL) values reported throughout this section, based on the scale provided in Appendix 1, reflect the collective assessment of the workshop participants and were established through table-level discussions and voting conducted during the workshop sessions. Each entry was typically assessed by four to seven voters drawn from the participants present at the corresponding table. These values represent the workshop’s perceived-maturity signal for each technology and are useful as a coordinated input from a representative cross-section of the European fusion AI community. They are not the output of independent, evidence-based TRL assessments of the kind used in EU Horizon and ITER procurement contexts, where the assessment is performed against published evidence by a designated reviewer or panel. Where individual entries appear to be at variance with the deployed state of the underlying technology in adjacent industries, the participant-review round is the appropriate venue for raising and reconciling such discrepancies. All assessments reflect the state of development as understood by workshop participants at the time of the workshop (April and May 2025); the gap between workshop and publication should be considered when interpreting TRL assignments for fast-moving technology categories.

A further methodological note concerns the scope of individual entries in the technology mapping. Several applications share a common underlying technology base, for example multiple document management and retrieval entries draw on the same RAG and LLM classification components and should be understood as related capability clusters rather than independent development lines. Beyond this, the 58 technologies span a wide range of abstraction levels, from specific algorithmic methods such as physics-informed neural networks for plasma simulation, to broad capability areas such as AI-driven process optimisation for advanced manufacturing, to enterprise tooling categories such as fusion-specific documentation management. Aggregate statistics in this section, including TRL distribution, dataset characteristics, TDA timelines, and TDA cost structures, combine entries across these abstraction levels. Readers should therefore interpret the aggregates as workshop-level signals of focus and concern rather than as comparable measurements of individual technologies at the same scale. Where finer comparisons are useful, the per-technology dashboards in Appendix 3 provide the underlying data without aggregation.

Similar to TRLs, Technology Development Actions (TDAs) were identified by workshop participants during the in-person characterization process, as described in section 2. For specific technologies, participants defined the recommended next development step. This enabled them to capture the analysis, prototyping, testing or other steps as a TDA. They also came up with estimates on expertise, success probability, implementation time, and cost for each TDA.

TECHNOLOGY

1. Technology Readiness Levels (TRLs)

The overall maturity of AI technologies assessed during the workshop indicates that most are still in early development stages. The data is visualized in Figure 4.

- 71% of technologies are at TRL 1–4, representing early research and conceptual validation.
- 15% fall within TRL 5–6, corresponding to validation and prototype phases.
- Another 14% are at TRL 7–9, indicating technologies close to or ready for deployment.

This distribution highlights the EU's strong research foundation, while also underlining the need for mechanisms to accelerate the transition from research to industrial implementation.




Figure 4: Distribution of TRL levels across technologies

2. Data availability and quality

Dataset availability across surveyed technologies:

- 80% of technologies have experimental datasets; 17% rely on synthetic data; 3% have no dataset
- Of technologies with datasets: 79% are adequate in size, 16% require additional synthesis, 2% are judged inadequate, 3% have no information on this characteristic
- Of technologies with datasets: 16% have clean, ready-to-use data, 51% require preprocessing, 19% need significant cleaning, 14% have no information on this characteristic

Data quality and formatting remain a key bottleneck, reinforcing the need for standardised preprocessing pipelines and shared data platforms.

3. Experimental infrastructure

Existing test facilities are available for 95% of technologies reporting facility status, with only 5% reporting partial access. No technology reported a complete lack of facilities. However,

36% of technologies reporting facility status require additional testing infrastructure, suggesting that while Europe's fusion research ecosystem is well equipped, scaling and accessibility remain challenges.

TECHNOLOGY DEVELOPMENT ACTIONS

1. Parameters

Across the surveyed technologies, 32 (55%) have at least one associated Technology Development Activity (TDA), resulting in a total of 49 distinct TDAs. Most technologies are linked to one or two TDAs, reflecting a balanced distribution of focus areas. This is shown in Figure 5.




Figure 5: Distribution of number of TDAs per technology

The thematic distribution of TDAs across domains in the order of high to low presence:

- Materials, Manufacturing, NDT & Assembly: 39%
- System engineering, Project Management & Operations: 37%
- Plasma Science: 14%
- Design & Simulation: 10%

This distribution, shown in Figure 6, illustrates a well-rounded portfolio, combining foundational, high-TRL engineering applications with longer-term, science-driven developments.




Figure 6: Distribution of TDAs per topic

Overall, 51% of the TDAs report a likelihood of success above 80%, particularly in assembly (80% of assembly-related TDAs have a chance of success of above 80%), simulation (75%), and plasma science (71%).

TDA implementation timelines:

- 17% deliverable in under 6 months (“quick wins”)
- 57% expected to complete within 6 months to 2 years
- 14% require more than 2 years, typically complex, multidisciplinary developments in design, simulation, or plasma science

TDA cost structure:

- 65% require investments above €100,000
- 19% fall in the €50,000–100,000 range
- 4% cost less than €50,000

The most resource-intensive TDAs are in design, simulation, materials research, and system engineering. Assembly and project management tend to be more affordable and faster to implement.

The distribution of the chance of success, implementation time and cost are shown in figure 7 (a), (b) and (c) respectively.




Figure 7: Distribution of chance of success (a), implementation time (b) and cost (c)

2. Priority and high-impact TDAs

The TDAs can be grouped into four strategic tiers according to their role, maturity, and potential impact.

Tier 1 – Foundational Infrastructure:

These TDAs form the backbone of the EU's AI-for-fusion ecosystem:

- **Fusion Data Platform (TRL 2)**
Most critical enabler, referenced as a prerequisite for 15 other technologies. Covers sub-TDAs such as AI-based visualisation tools, platform architecture, and governance for coordinated data user agreements. Success rate is greater than 80%; costs range from <€50k (for governance) to >€100k (platform implementation). Directly addresses the sector's main bottleneck: dispersed, unstructured data.
- **Accessible Database for Simulations and Experiments (TRL 2-4)**
Enables surrogate modelling and digital twins through data standardisation, fast queries, and user-friendly accessibility.
- **Machine-Agnostic AI Tool for Data Annotation (TRL 2)**
A machine-agnostic AI tool for data annotation would support fusion development by improving the curation, validation, and labelling of large and heterogeneous datasets generated across experiments, diagnostics, simulations, and engineering systems. By providing a reusable and interoperable framework across different machines and platforms, it could help address current challenges in fusion data quality, accessibility, and standardisation, enabling more effective AI model development, knowledge sharing, and cross-facility collaboration.

Together, these TDAs represent the essential data infrastructure for the European public fusion AI ecosystem, coordinated by entities such as EUROfusion, ITER, F4E, and several AI laboratories across Europe.

Tier 2 – High-Impact Quick Wins:

Several TDAs are already at high TRL with immediate operational value, some of the examples are:

- **PAUT AI Optimisation (TRL 8):**
The TRL 8 assignment reflects the deployed maturity of PAUT as an inspection

technique. The fusion-specific task is the adaptation and qualification of AI-based defect-recognition models to fusion component materials and acceptance criteria; this layer is less mature than the underlying PAUT technique in adjacent industries.

- **Digital Optimisation of Radiographic Testing (TRL 7):**
Enhances weld inspection verification; strong cross-sector adoption potential, success >80% and close to deployment.
- **Virtual Travel Agent for Metrology (TRL 7):**
AI-based scheduling tool that will automate instrument booking and logistics; low cost to realize (<€100k), high success, immediate efficiency gains. It can be classified as “quick win”.

Tier 3 – Strategic Competitive Advantage:

TDAs that could position the EU as a global leader in fusion-related AI:

- **AI for Discovery of New Materials (TRL 4):**
Leverages European expertise in materials science and HPC to identify new materials with specific physical and material properties. Incurs high cost (>€100k) with success >80%.
- **Material Database Generation (TRL 6):**
Supports codification of material properties into design standards (such as ASME, ITER SDC-IC). Incurs high cost (>€100k) with high success rate >80%.
- **Fast Surrogates for Fusion Scenario Simulations (TRL 2):**
Replaces high-fidelity physics simulations with ML models; highly competitive internationally, with 5+ active EU entities and several non-EU players, underscoring the need for sustained investment and coordination.

Tier 4 – Multi-Phase and Long-Term Developments:

Complex, multi-phase TDAs with longer implementation horizons:

- **AI-Enhanced Requirements Checking (TRL 1):**
Phased approach extending from prototype deployment to full system-wide propagation.
- **Automated Contract Generation (TRL 2):**
Decomposes challenges into sub-TDAs covering framework design, legal mapping, dataset preparation, and proof-of-concept validation.

These activities are particularly valuable for illustrating realistic risk management and incremental progress toward automation of complex administrative and engineering tasks.

3. Interdependencies and cross-sector leverage

Several TDAs serve as enabling components for others. The Fusion Data Platform is foundational for auto-labelling tools, event prediction, digital twins, and plasma machine learning applications. The Accessible Database underpins surrogate modelling and reactor digital twins, while the Governance Framework for Data Sharing facilitates external collaboration and startup integration. Likewise, the Material Database Generation enables downstream activities such as material acceptability prediction and AI-driven materials discovery.

Together, these interdependencies highlight the systemic nature of AI development in fusion: progress in one area can unlock rapid advancements across multiple domains. Prioritising such enabling TDAs will ensure maximum return on investment and strategic coherence across the EU's fusion AI landscape.

It is also worth noting that 88% of the technologies identified in this mapping exercise are classified as general technologies with cross-sector relevance, extending beyond fusion into

other areas. This has two important implications for investment decisions. First, progress made by commercial companies undergoing AI transformation in areas (such as CAD-related AI applications, metrology, welding inspection, and FEM optimisation, etc.) may directly accelerate fusion applications, reducing the need for parallel investment. Second, when prioritising TDAs for maximum return on investment, cross-sector technologies deserve particular attention, since successes achieved in adjacent industries can be leveraged for fusion, while fusion-driven developments in these areas may in turn generate broader industrial value. The fusion AI community is therefore encouraged to actively monitor commercial developments in these domains rather than treating them in isolation.

4. TDA timeline and implementation priorities

The implementation roadmap for the identified TDAs can be structured across four temporal and strategic phases, balancing immediacy, foundational needs, and long-term competitiveness. This timeline provides a pragmatic framework to guide investments, ensure continuity, and maximise returns across short-, medium-, and long-term horizons. The four phases below organise the same TDAs introduced in the strategic tiers above by implementation horizon; some TDAs appear in both groupings.

Immediate Priorities (Next 6 months)

The first group of TDAs focuses on low-cost, high-success-rate actions that deliver tangible value within a short timeframe. These “quick wins” combine high feasibility with direct operational impact and help build early momentum across the ecosystem.

TDA	Cost	Timeline	Success Rate
Governance Framework for Coordinated Data user Agreements	<€50k	<6 months	>80%
Market Survey for Contract Generation	€50-100k	<6 months	>80%
Virtual Travel Agent for Metrology	€50-100k	<6 months	>80%
Identify Potential Areas for Predictive Maintenance	<€50k	<6 months	>80%

Foundation Building (6–24 months)

This second phase targets the creation of shared data infrastructure and tools that underpin over 15 downstream technologies. Each of these activities offers immediate technical value and a clear path to industrial deployment, while also providing spillover benefits for other sectors such as aerospace, manufacturing, and healthcare. Collectively, they represent a bridge between early-stage research and high-TRL implementation.

TDA	Cost	Timeline	Success Rate
Machine-Agnostic AI Tool for Data Annotation	>€100k	6-24 months	>80%
Advance AI Visualisation Tools	€50-100k	6-24 months	>80%
Fusion Data Platform Architecture	>€100k	6-24 months	>80%
Accessible Database for Simulations	>€100k	6-24 months	>80%

High-TRL cross-sector TDAs (6–24 months)

The third group comprises medium-term TDAs with high technical maturity (TRL 6–8) and cross-sector relevance. Each of these activities offers immediate technical value and a clear

path to industrial deployment, while also providing spillover benefits for other sectors such as aerospace, manufacturing, and healthcare. Collectively, they represent a bridge between early-stage research and high-TRL implementation.

TDA	Cost	Timeline	Success Rate
PAUT AI Optimisation	>€100k	6-24 months	>80%
Digital RT Process Optimisation	>€100k	6-24 months	>80%
AI for Materials Discovery	>€100k	6-24 months	>80%
Material Database Generation	>€100k	6-24 months	>80%

Strategic Long-Term Developments (2+ years)

Finally, several TDAs address complex, high-impact areas that will shape the EU's long-term competitive positioning in fusion and AI. While these initiatives require greater investment and longer timelines, their expected returns, both scientific and strategic, are substantial. They advance the EU's position in high-value, fusion-specific AI applications that few global actors currently target.

TDA	Cost	Timeline	Success Rate
Event Prediction, Detection, and Classification	>€100k	>2 years	>80%
AI-enhanced Requirements Checking	>€100k	>2 years	40-80%
Fast Surrogates for Fusion Scenarios	>€100k	>2 years	>80%
Predictive Maintenance Proof-of-Concept	>€100k	>2 years	40-80%

5. Highest-priority TDAs

The table below identifies the TDAs with the strongest return on investment when comparing cost, timeline, and likelihood of success.

TDA	Cost	Timeline	Success Rate	Why High priority
Governance Framework for Data Sharing	<€50k	<6 months	>80%	Lowest cost, fastest implementation, ecosystem-wide impact
Virtual Travel Agent for Metrology	€50-100k	<6 months	>80%	Immediate operational efficiency gains
Material Database Generation	>€100k	6-24 months	>80%	Enables multiple downstream developments and supports standardization
Fusion Data Platform	>€100k	6-24 months	>80%	Underpins 15+ related technologies and forms backbone of EU's fusion data ecosystem
PAUT Optimisation	>€100k	6-24 months	>80%	Close to deployment, cross-sector potential
AI for Materials Discovery	>€100k	6-24 months	>80%	Builds on EU scientific strengths
Auto-Labelling Tools for Data Annotation	>€100k	6-24 months	>80%	Directly addresses data preprocessing bottleneck
Digital RT Process Optimisation	>€100k	6-24 months	>80%	Cross-sector industrial potential

In summary, the most cost-effective strategy for the near term lies in combining Immediate Priority and Foundation Building TDAs. This dual-track approach achieves fast, visible outcomes while laying the groundwork for scalable, high-impact AI applications in fusion.

5.2.2 SWOT analysis

STRENGTHS	WEAKNESSES
▪ Broad technology portfolio ▪ Strong research ecosystem ▪ Established infrastructure ▪ HPC access ▪ High success rate potential ▪ Data availability ▪ Test facilities	▪ Data challenge ▪ Data governance gaps ▪ Regulatory/ethical barriers ▪ Resource constraints ▪ High costs ▪ Gap between research and industrialization ▪ Limited market pull
OPPORTUNITIES	THREATS
▪ Standardization & platformisation ▪ Public-private partnerships ▪ Cross-domain applications ▪ Digital twins & simulation ▪ Talent development ▪ Policy & funding levers ▪ International leadership	▪ Global competition ▪ Data privacy & IP issues ▪ Fragmentation & duplication ▪ High uncertainty ▪ Scalability challenges ▪ Regulatory uncertainty ▪ Talent shortage

1. Strengths

The European ecosystem, encompassing academia, research institutions, and private industry, is actively engaged across all nine subcategories. Numerous EU universities, research centres, and consortia contribute expertise spanning multiple domains. More than five EU entities are actively involved for 36% of technologies. Collaboration within the EU and with international partners is extensive, particularly in data sharing and joint research initiatives. The EU has access to major fusion research facilities (JET, WEST, ASDEX, TJ-II, W7-X, and TCV) as well as to the coordinated framework of the EUROfusion consortium. Computational capabilities can be provided by EuroHPC centres, while EUROfusion, ITER, and F4E offer a coordinated ecosystem for integration and collaboration. A substantial subset of the surveyed technologies also rests on commercially mature foundations in adjacent industries, including NDT inspection (Trueflow, Yxlon, Visiconsult), metrology (Hexagon, ZEISS), materials informatics (Citrine), ML surrogate modelling (Siemens Simcenter, NVIDIA PhysicsNeMo), and broader engineering tooling (Ansys, Dassault Systèmes, AVEVA). For these entries the fusion development effort is one of adaptation and qualification rather than capability creation, allowing EU fusion programmes to draw on tooling and supply chains already deployed in European industry.

Over half of the Technology Development Activities (TDAs) show a likelihood of success above 80%, particularly in the areas of assembly, plasma science, and simulations. Only one TDA reports a success probability below 40%. Experimental datasets are available for most technologies, although additional preprocessing is still required. Among technologies with existing datasets, 79% are of adequate size, while 16% require further data synthesis. Furthermore 95% of technologies reporting facility status have access to a testing facility, and 64% of technologies reporting facility status report that no additional facilities are needed.

2. Weaknesses

The main weakness lies in data quality and format. Only 16% of AI applications identified possess clean, ready-to-use datasets, while 51% have data in unusable formats and 19% have undergone no cleaning at all. Approximately 80% of applications rely on experimental data, 17% on synthetic data, and 3% have no dataset available. The fact that over half of the applications require additional preprocessing or are not cleaned at all highlights a significant bottleneck. Fusion-related data is often confined within unstructured or poorly documented formats, rather than being stored in machine learning-ready, structured databases.

Workshop participants confirmed this as the dominant concern: across all applications assessed, fragmented and spread-out data was cited as the primary obstacle by 40% of respondents, followed by high uncertainty (26%) and dependency on external or proprietary tools (25%). Notably, insufficient computing resources ranked lowest at 11%, indicating that data availability, not compute, is the principal bottleneck for AI development in fusion.

A further limitation is the absence of unified data-sharing frameworks, particularly affecting access for external stakeholders and startups. In most cases, data accessibility is restricted, constraining collaboration and innovation.

Many technologies also face challenges regarding explainability, a critical factor for safety-sensitive applications such as fusion which is cited as an obstacle by 22% of applications assessed. Insufficient transparency can impede trust, regulatory approval, and the adoption of AI-driven tools. Related to this, low model robustness was identified as a concern by 22% of them, a significant issue in a domain where AI outputs may inform safety-critical decisions. It is also to be noted that regulatory and ethical considerations can slow progress, especially in cross-border or public-private collaborations.

Beyond data quality, the fusion AI ecosystem currently lacks systematic frameworks for traceability of AI models and data covering model lineage, dataset provenance, and decision pathways. In a nuclear-regulated environment, this is not a best practice but a qualification requirement, and its absence represents a gap that will need to be addressed before AI tools can be deployed in operational settings.

Limited computational, financial, and human resources present recurring obstacles, particularly for high-cost, long-term initiatives. The human resource constraint warrants particular attention: personnel combining deep fusion-domain knowledge with practical AI and data engineering competence are scarce across both public programmes and private fusion companies alike, and this gap cannot be resolved by procurement alone but requires sustained investment in training, hiring, and knowledge transfer across the fusion AI community. In certain specialised applications, there are currently no active entities engaged and 65% of Technology Development Actions are of high cost.

A substantial share of technologies (71%) remains at early stages of maturity (TRL 1-4) within fusion-institutional contexts, with limited deployment in operational fusion environments. Additionally, several fusion-specific AI solutions exhibit limited potential for broader market application, which may hinder investment and long-term sustainability. The deployed maturity of many of these technologies in adjacent industrial domains differs substantially from the assessments recorded here. For that subset, the binding constraint is qualification, integration, and adoption within fusion programmes rather than technology development, making the weakness as much institutional and procedural as technical.

3. Opportunities

Establishing unified data platforms and standardised processes will significantly accelerate progress and attract new stakeholders to the ecosystem. Developing harmonised data and metadata frameworks will enhance data accessibility, interoperability, and cross-border collaboration both within the EU and globally. Strategic partnerships with private companies and startups can further accelerate innovation, bring in additional resources, and bridge existing expertise gaps. It is also to be noted that 17% of identified TDAs can be delivered within six months at relatively low cost, representing an immediate window for visible progress and stakeholder confidence-building.

The EU venture capital gap for deep-tech AI commercialisation, while currently a threat, also represents a structural opportunity: fusion AI developments that demonstrate clear industrial applicability and cross-sector relevance are well-positioned to attract international investment if the right commercialisation pathways and investor-facing forums are established.

AI technologies developed for fusion applications also hold strong potential for cross-sector deployment, opening new funding opportunities and collaboration pathways. This diversification can reinforce the EU's competitiveness and leadership in global AI markets. 88% of the AI use cases under development have potential applications beyond fusion, including in aerospace, defence, manufacturing, and healthcare. The reverse flow is equally significant: as noted in Strengths and detailed in section 3.4, a substantial share of the technologies mapped here already have mature analogues in adjacent industries. Actively monitoring and transferring these capabilities into fusion contexts represents a faster and more capital-efficient path than parallel development. Practical mechanisms include qualification-focused TDAs that take commercially deployed tools as the starting point, joint development with European industrial vendors, and a systematic cross-domain benchmarking step in future workshop iterations.

Expanding the use of digital twins (understood here as integrated AI systems combining real-time sensor data, surrogate models, and physics simulation) and AI-driven simulations can reduce operational costs, improve safety, and optimise performance. Digital twin methodologies are increasingly mature in adjacent industries such as aviation, energy, and manufacturing. The transfer of these methods into fusion represents a significant near-term opportunity that is likely more immediately valuable than the reverse direction of licensing fusion-developed tools outward. Furthermore, the EU could capitalise on fusion AI expertise by licensing tools to international fusion programmes.

EU's extensive academic and research base can be leveraged to train the next generation of AI and fusion specialists, helping to address critical skill shortages. EU-level initiatives, such as Horizon Europe and Digital Europe, can be mobilised to fill resource gaps and support high-potential TDAs with clear implementation roadmaps and measurable outcomes. Through coordinated action and strategic investment in flagship projects, the EU can consolidate its position as a global leader in AI for fusion, shaping international standards and best practices. With 71% of technologies still at TRL 1–4, the EU is in a position to shape international standards and best practices before the field matures globally: a window that will narrow as US and Chinese programmes accelerate.

4. Threats

The United States, China, the United Kingdom, and major private technology companies (e.g. DeepMind, Microsoft, OpenAI) are advancing rapidly and may outpace EU efforts in several key areas. Both the US and China are making substantial investments in AI for fusion, posing a potential risk to the EU's leadership position.

Dependence on non-EU data sources, language models and AI coding tools, software tools, or cloud infrastructures introduces vulnerabilities and undermines technological autonomy. The exposure is uneven: the foundational AI layer (frontier language models, AI cloud platforms) is heavily US-anchored (OpenAI, Anthropic, Google, Microsoft Azure, AWS), while the industrial engineering tooling on which many entries in section 3.4 build includes substantial EU presence (Siemens, Dassault Systèmes, Hexagon, AVEVA, ZEISS) alongside US vendors (Ansys, Citrine Informatics, Bentley). Sovereignty risk is therefore concentrated at the AI and cloud layer rather than uniformly distributed. Disputes related to data ownership, intellectual property rights, and access restrictions can further hinder collaboration and slow progress.

A related but distinct exposure is the IP position of commercial vendors in adjacent industries on which fusion AI development depends, notably in real-time manufacturing optimisation, materials informatics, and AI-augmented NDT. Where vendor IP prevents adaptation, open-sourcing, or qualification for nuclear use, the EU may need to negotiate dedicated qualification pathways or fund parallel non-proprietary developments to retain control over safety-critical components.

Qualification frameworks will increasingly require demonstrable traceability of AI models and datasets used in safety-relevant applications. The fusion AI community is not yet systematically prepared for this requirement, creating a compliance risk as technologies approach higher TRL levels.

The workshop data reinforces this concern: dependency on external or proprietary tools was identified as an obstacle by 25% of technologies, underlining the strategic importance of building EU-native tooling and open data infrastructure.

A structural threat is the limited capacity within the EU to translate publicly funded AI research into commercial ventures. The Draghi Report (Draghi 2024), citing the European Patent Office's Valorisation of scientific results scoreboard, notes that only about one-third (36%) of patented inventions from European universities and public research organisations are currently commercially exploited, with a further 42% planned for exploitation but not yet realised (EPO, 2020). Many of the AI applications identified in this workshop have cross-sector applicability beyond fusion, yet without dedicated commercialisation pathways, viable developments risk remaining within the public research sphere rather than reaching deployment. This is simultaneously a threat and an opportunity: addressing the EU's commercialisation gap for fusion AI would unlock both economic value and accelerated deployment.

Insufficient coordination within the EU risks leading to duplication of efforts, inefficient use of resources, and missed opportunities for synergy across projects and member states.

Many technologies face considerable uncertainty regarding their success, particularly those requiring long-term investment (over two years) and significant financial commitments (exceeding €100,000).

Heavy reliance on short-term or project-based funding threatens continuity and limits long-term impact. In addition, evolving EU regulations on AI, data protection, and ethics may

introduce delays or compliance burdens, particularly for cross-border collaborations.

Scaling AI solutions from pilot implementations to full operational deployment remains a major challenge, identified as an obstacle by 21% of technologies assessed, particularly for those with high computational or infrastructure demands. Bridging this gap requires not only technical development but sustained investment and commercialisation support of the kind that the EU's current deep-tech AI funding landscape does not consistently provide.

Finally, a shortage of specialised AI and machine learning expertise within the fusion research community could create a critical bottleneck, especially in niche domains requiring deep technical and domain-specific knowledge.


5.2.3 Technology Roadmaps

This section presents some of the Technology Development Actions (TDAs) and technologies in the form of roadmaps. The timings are indicative and may evolve significantly depending on funding available from various sources and associated priorities.

TDAs which are not fundamentally linked to other activities, and which can be executed independently are not included on the roadmaps.


Data Infrastructure & Standardization

Objective: build a foundational data ecosystem to enable downstream AI applications and raise TRLs across domains



Digital Twins & Real-time Operations

Objective: deploy AI-driven digital twins and real-time systems to optimize fusion device performance and safety



6 Conclusion

The first European Artificial Intelligence Technology Mapping exercise for fusion marks an important step towards accelerating the development and deployment of AI technologies for fusion across Europe. It brought together academia, research institutions, startups, and industry to assess current European capabilities and define a clear roadmap for advancing the AI tools and methods needed to support fusion as a viable energy source.

Across four domains: design and simulation; materials and manufacturing; plasma science; and system engineering and operations, 60 AI applications were initially identified, of which 58 were carried forward into the final taxonomy. This analysis produced 49 Technology Development Activities, with 55% of technologies having at least one associated TDA. The landscape reflects a field in active but early development: 71% of technologies sit at TRL 1–4, while 80% have access to experimental datasets and 95% of technologies reporting facility status have existing test infrastructure. Crucially, more than half of all TDAs carry a likelihood of success above 80%, providing a strong foundation for targeted investment. These figures are workshop signals (section 5.2.1) and should be read as indicators of community focus rather than comparable measurements at the same scale. For the substantial subset of technologies whose underlying capability is mature in adjacent industries, the binding constraint within fusion programmes is qualification, integration, and adoption rather than technology development (section 5.2.2).

The findings reveal several cross-cutting themes that shape the strategic picture. First, 88% of the identified AI applications are general technologies with relevance beyond fusion, spanning aerospace, manufacturing, healthcare, and defence. This means that progress in adjacent sectors can directly accelerate fusion AI development, and vice versa. The entries in section 3.4 already incorporate a light cross-domain review by Novatron. Extending this to a deeper systematic benchmarking of proposed TDAs against the state of the art in adjacent industries would reduce redundant investment and increase the return on public funding.

Second, several TDAs serve as enabling components for others (see section 5.2.1): the Fusion Data Platform underpins auto-labelling, event prediction, digital twins, and plasma machine learning applications; the Accessible Database supports surrogate modelling and reactor digital twins; and the Material Database enables material acceptability prediction and AI-driven materials discovery. Prioritising these foundational TDAs delivers more leverage than equivalent investment in individual downstream applications.

Third, while data availability and test infrastructure are strengths, data quality and readiness represent the principal bottleneck identified by workshop participants. This reflects the public research perspective that dominated the workshop. For private fusion companies, whose data and development workflows differ substantially, the binding constraints may differ.

Fourth, the threats facing the EU's fusion AI ecosystem are concentrated rather than uniform (see section 5.2.2). Non-EU dependency sits at the foundational AI and cloud layer, where US-anchored providers dominate, while the industrial engineering tooling layer retains substantial EU presence. Global competition is intensifying, with major investments in fusion AI from the United States, China, the United Kingdom, and private technology companies. A further structural threat is the EU's limited capacity to translate publicly funded AI research into commercial ventures, risking the loss of viable cross-sector developments to the research domain rather than industrial deployment.

Fifth, the human-resource gap noted in section 5.2.2 is structural and cannot be closed by procurement. Sustained training, hiring, and cross-disciplinary knowledge transfer will be required to scale fusion AI capability across the European ecosystem.

Several applications in the System engineering, Project Management and Operations domain, including requirements validation, contract generation, and documentation management, are already achievable today using commercially available LLM models configured for organisational workflows. For these applications, the path to implementation is adoption and integration rather than dedicated R&D. This is not a limitation of the mapping exercise, which rightly captures the full landscape of AI relevance to fusion operations, but it does have implications for how F4E and partners prioritise investment. Fusion-specific R&D effort is better directed at applications where existing tools are insufficient and off-the-shelf solutions are inadequate.

Finally, the workshop applications consistently position AI as a layer that accelerates analysis, screens candidates, and surfaces anomalies for qualified human decision-making, rather than as a replacement for fusion-engineering judgment. This pattern is reinforced by the explainability, traceability, and qualification requirements identified as obstacles for safety-relevant applications (see section 5.2.2).

Several TRL assessments and “quick wins” classifications in this report are likely already out of date at publication, given the pace at which the AI landscape evolves. A lightweight online update activity on a shorter cadence, focused on reassessing TRL assignments for fast-moving entries and identifying newly available commercial analogues, would keep the collective understanding of the landscape and the strategic channelling of funds into TDAs current.

7 Bibliography

- Clavier, B, D Zarzoso, D del-Castillo-Negrete, en E Frénod. 2025. „Generative-machine-learning surrogate model of plasma turbulence.” *Physical Review E* 111: L013202. doi:10.1103/PhysRevE.111.L013202.
- Deb, K, A Pratap, S Agarwal, en T Meyarivan. 2002. „A fast and elitist multiobjective genetic algorithm: NSGA-II.” *IEEE Transactions on Evolutionary Computation* 6 (2): 182-197. doi:10.1109/4235.996017.
- Degrave, Jonas, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco Carpanese, Timo Ewalds, et al. 2022. „Magnetic control of tokamak plasmas through deep reinforcement learning.” *Nature* 602: 414-419. doi:10.1038/s41586-021-04301-9.
- Draghi, Mario. 2024. *The Future of European Competitiveness*. Brussels: European Commission.
- Draxl, C, en M Scheffler. 2018. „NOMAD: The FAIR concept for big data-driven materials science.” *MRS Bulletin* 43 (9): 676-682. doi:10.1557/mrs.2018.208.
- Fischer, R., C.J. Fuchs, B. Kurzan, W. Suttrop, E. Wolfrum, en ASDEX Upgrade Team. 2010. „Integrated Data Analysis of Profile Diagnostics at ASDEX Upgrade.” *Fusion Science and Technology* 58 (2): 675-684. doi:https://doi.org/10.13182/FST10-110.
- Ghaznavi, Ali, Eric Duffner, Emanuel Kästle, en Bartosz Popiela. 2026. „A scalable deep learning approach for robust acoustic emission-based damage classification in composite pressure vessels.” *Results in Engineering* 30: 110668. doi:10.1016/j.rineng.2026.110668.
- Humphreys, David, A Kupresanin, M.D. Boyer, J. Canik, C.S. Chang , E.C. Cyr, R. Granetz, et al. 2020. „Advancing Fusion with Machine Learning Research Needs Workshop Report.” *Journal of Fusion Energy* 39: 123-155. doi:10.1007/s10894-020-00258-1.
- Jain, A, S P Ong, G Hautier, W Chen, W Davidson Richards, S Dacek, S Cholia, et al. 2013. „Commentary: The Materials Project: A materials genome approach to accelerating materials innovation.” *APL Materials* 1 (1): 0110002. doi:10.1063/1.4812323.
- Kates-Harbeck, Julian, Alexey Svyatkovskiy, en William Tang. 2019. „Predicting disruptive instabilities in controlled fusion plasmas through deep learning.” *Nature* 568: 526-531. doi:10.1038/s41586-019-1116-4.
- Kumar, Jai, David Zarzoso, Virginie Grandgirard, Jan Ebert, en Stefan Kesselheim. 2023. „Physics informed Neural Networks applied to the description of wave-particle resonance in kinetic simulations of fusion plasmas.” doi:10.48550/arXiv.2308.12312.
- Na, Yiming, Yunze He, Baoyuan Deng, Xiaoxia Lu, Hongjin Wang, Liwen Wang, en Li Cao. 2025. „Advances of Machine Learning in Phased Array Ultrasonic Non-Destructive Testing: A Review.” *AI* 6: 124. doi:10.3390/ai6060124.
- Ortiz de Zuniga, Maria, Max Febvre, Nawal Prinja, Ana Maria Camacho, en Alvaro Rodriguez-Prieto. 2026. „Machine learning-supported prognosis of electron beam

- welding quality in ITER vacuum vessel manufacturing." *Nuclear Engineering and Technology* 58 (6): 104205. doi:10.1016/j.net.2026.104205.
- Qi, Charles R, Hao Su, Kaichun Mo, en Leonidas J Guibas. 2017. „PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation." *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Honolulu. 77-85. doi:10.1109/CVPR.2017.16.
- Qi, Charles R, Li Yi, Hao Su, en Leonidas J Guibas. 2017. „PointNet++: deep hierarchical feature learning on point sets in a metric space." *NIPS'17: Proceedings of the 31st International Conference on Neural Information Processing Systems*. Long Beach. 5105-5114.
- Raissi, M, P Perdikaris, en G E Karniadakis. 2018. „Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations." *Journal of Computational Physics* 378: 686-707. doi:10.1016/j.jcp.2018.10.045.
- Rodriguez-Fernandez, P, N T Howard, A Saltzman, S Kantamneni, J Candy, C Holland, M Balandat, S Ament, en A E White. 2024. „Enhancing predictive capabilities in fusion burning plasmas through surrogate-based optimization in core transport solvers." doi:10.48550/arXiv.2312.12610.
- Rodriguez-Fernandez, P, N T Howard, en J Candy. 2022. „Nonlinear gyrokinetic predictions of SPARC burning plasma profiles enabled by surrogate modeling." *Nuclear Fusion* 62: 076036. doi:10.1088/1741-4326/ac64b2.
- Schissel, D P, R M Nazikian, en T Gibbs. 2025. „Summary report from the mini-conference on Digital Twins for Fusion Research." *Physics of Plasmas* 32: 050601. doi:10.1063/5.0273586.
- Seo, Jaemin, SangKyeun Kim, Azarakhsh Jalalvand, Rory Conlin, Andrew Rothstein, Joseph Abbate, Keith Erickson, Josiah Wai, Ricardo Shousha, en Egemen Kolemen. 2024. „Avoiding fusion plasma tearing instability with deep reinforcement learning." *Nature* 626: 746-751. doi:10.1038/s41586-024-07024-9.
- Spangher, Lucas, Matteo Bonotto, William Arnold, Druva Chayapathy, Tommaso Galligani, Alexander Spangher, Francesco Cannarile, Daniele Bigoni, Eliana de Marchi, en Cristina Rea. 2025. „DisruptionBench and Complimentary New Models: Two Advancements in Machine Learning Driven Disruption Prediction." *Journal of Fusion Energy* 44. doi:10.1007/s10894-025-00495-2.
- Wang, A M, A Pau, C Rea, O So, C Dawson, O Sauter, M D Boyer, et al. 2025. „Learning plasma dynamics and robust rampdown trajectories with predict-first experiments at TCV." *Nature Communications* 16: 8877. doi:10.1038/s41467-025-63917-x.
- Wu, Niannian, Zongyu Yang, R Li, Ning Wei, Y Chen, Qianyun Dong, J Li, et al. 2025. „High-fidelity data-driven dynamics model for reinforcement learning-based control in HL-3 tokamak." *Communications Physics* 8: 393. doi:10.1038/s42005-025-02302-y.
- Yang, Zongyu, Zhenghao Yang, Wenjing Tian, Jiyuan Li, Xiang Sun, Guohui Zheng, Songfen Liu, et al. 2026. „FusionMAE, a self-supervised pretrained model to optimize and

simplify diagnostic and control of fusion plasma.” *Communications Physics* 9.
doi:10.1038/s42005-026-02626-3.

Appendix 1: Technology Readiness Levels

For this workshop, a TRL scale from 1 to 9 will be used¹.

Environment	Goal	Product / Evaluation	Outputs	TRL	Description
Laboratory	Research	Proof of concept	Scientific articles published on the principles of the new technology	TRL 1	Basic principles observed
			Publications or references highlighting the applications of the new technology.	TRL 2	Technology concept formulated
			Measurement of parameters in the laboratory	TRL 3	Experimental proof of concept
Simulation	Development	Prototype	Results of tests carried out in the laboratory.	TRL 4	Technology validated in lab
			Components validated in a relevant environment.	TRL 5	Technology validated in relevant environment
			Results of tests carried out at the prototype in a relevant environment.	TRL 6	Technology demonstrated in relevant environment
Operational	Implementation	Commercial product/service (certified)	Result of the prototype level tests carried out in the operating environment.	TRL 7	System prototype demonstration in operational environment
			Results of system tests in final configuration.	TRL 8	System complete and qualified
		Deployment	Final reports in working condition or actual mission.	TRL 9	Actual system proven in operational environment

¹ Martínez-Plumed, F., Gómez, E. and Hernández-Orallo, J., *AI watch, assessing technology readiness levels for artificial intelligence*, Publications Office, 2020, <https://data.europa.eu/doi/10.2760/15025>

Appendix 2: Technology assessment

The following list outlines the information displayed on the dashboard and the corresponding possible responses.

Remark: A “/” on the dashboard indicates that the question is not applicable or that no data were collected.

1. DATASET CHARACTERISTICS

PARAMETER	QUESTION	ANSWERS
AVAILABILITY	Is there an existing dataset available?	▪ Synthetic dataset ▪ Experimental dataset ▪ No
SIZE AND DIVERSITY	Is the available dataset large and diverse enough to develop this AI solution effectively?	▪ Yes ▪ No ▪ Synthesize more data ▪ More experimental data expected in 3-6 months
CLEANING	Is the available dataset clean and well-structured?	▪ Yes ▪ No ▪ Needs preprocessing (in PDF's, unlabelled images, ...)
OWNER OF DATASET + OPEN	Who owns the dataset and is it open?	Free text
BROADER MARKET POTENTIAL	Is the AI technology specific to fusion, or can it be applied in other domains as well?	Free text

2. MAJOR OBSTACLES

This section highlights the major obstacles to implementing the AI technology. During the workshop, participants could vote for multiple obstacles. The data is currently presented as a percentage, calculated by dividing the number of votes for each obstacle by the total number of votes.

The obstacles are:

NAME ON DASHBOARD	FULL NAMING
NONE	none
SPREAD-OUT DATA	spread-out data
LOW EXPLAINABILITY	lack of explainability
LOW MODEL ROBUSTNESS	low model robustness
REGULATORY/ETHICAL	regulatory or ethical concerns
SCALABILITY ISSUES	scalability issues
HIGH UNCERTAINTY	high uncertainty
TOOL DEPENDENCY	dependency on external or proprietary tools

3. FACILITIES

Under facilities, three major questions are answered:

PARAMETER	QUESTION	ANSWERS
EXISTING TEST FACILITIES	Are there existing test environments or facilities that can be used to develop or validate/benchmark this AI solution?	▪ Yes ▪ No ▪ Partially
ADDITIONAL TEST FACILITIES	Are additional test facilities or validation environments needed?	▪ Yes ▪ No
COMPUTATIONAL & INFRASTRUCTURE REQUIREMENTS	What are the computational and infrastructural requirements for developing and deploying the AI system?	Free text

4. ENTITIES

PARAMETER	QUESTION	ANSWERS
ENTITIES IN THE EU	How many entities are currently working on this AI application in the EU?	▪ 0 ▪ 1 ▪ < 5 ▪ > 5
ENTITIES OUTSIDE THE EU	How many entities are currently working on this AI application outside of the EU?	▪ 0 ▪ 1 ▪ < 5 ▪ > 5
ENTITIES – LIST	List known institutions, organisations, companies, or research groups	Free text
ENTITIES LEADING	If few or no EU entities are involved, who currently holds the leading expertise in this application?	Free text
BROADER MARKET POTENTIAL	Is the AI technology specific to fusion, or can it be applied in other domains as well?	Free text

5. TECHNOLOGY DEVELOPMENT ACTIONS

PARAMETER	QUESTION	ANSWERS
TDA NAME		Free text
TDA DETAILS		Free text
TDA EXPERTISE	Is there expertise in Europe for doing the TDA?	▪ Yes ▪ No
TDA EXPERTISE DETAILS	List known institutions, organisations, companies, or research groups	Free text
CHANCE OF SUCCESS	What is the probability of successful outcome?	▪ < 40% ▪ 40-80% ▪ > 80%
IMPLEMENTATION TIME	What is the time needed to implement the TDA?	▪ < 6 months ▪ 6 months to 2 years ▪ > 2 years
COST	What is the cost for the next step?	▪ < 50k ▪ 50k to 100k ▪ > 100k

Appendix 3: Technology dashboards

Design and simulation

- AI/ML driven multi-objective optimization frameworks
- Categorisation and clustering of CAD clashes
- Component tagging for 3D-scanning
- Generation of E3D or 3D drawings
- Optimisation of magnetic fields using AI-guided inverse design algorithms
- P&ID recognition for engineering data extraction and consistency
- Quality control of CAD drawings
- Active learning for data-efficient experimentation
- Digital twins of fusion reactors
- GenAI for code template generation
- GenAI for synthetic data generation
- Hybrid models and Physics-Informed Machine Learning for simulations
- ML-driven surrogate modelling

Materials, manufacturing and NDT, and assembly

- AI-guided discovery and development of new materials
- Generation and curation of materials databases
- Material acceptability prediction for a dedicated application
- Prediction of behaviour and mechanical properties of new materials
- AI-driven acoustic emission for non-metallic material inspection
- AI-driven distortion prediction
- AI-driven process optimization for advanced manufacturing
- Automated digital RT data processing for weld inspection
- Automated PAUT data processing for weld inspection
- Real-time optimization of manufacturing parameters
- Welding success rate prediction
- AI-driven extrapolation of operating conditions from local testing
- Automated mapping of metrology cloud points to component models
- Detect modifications on drawings
- Identification of most optimal assembly sequence
- Instruments and transport virtual agent
- Optimization of component positioning during assembly
- Prediction of metrology data

Plasma science


- Auto-labelling AI tools for plasma state datasets
- Fusion data platform for machine learning applications
- ML-accelerated pedestal MHD stability evaluations
- Predictive modelling of plasma instabilities and disruptions
- Time-resolved, physics-informed neural networks for tokamak total emission reconstruction
- Trajectory planning enabled by ML/AI

System engineering, project management and operations

- AI-driven validation of requirements
- Chatbot for requirements rationalization
- LLM-driven check of requirements propagation
- AI-powered proposal support tool for project disruptions
- Automated contract generation
- Automatic status reporting from data sources
- Prediction documentation review time
- Prediction of contract outcome based on past performance
- Prediction of tool usage
- Process optimization tool
- Risk register forecast prediction
- AI-assisted experimental design
- AI-driven diagnostic data integration
- Automatic generation of EOM report
- Dimensionality reduction for streaming data
- Fusion specific documentation management
- LLM-based prediction of document type
- Predictive maintenance of components
- Real-time operation of fusion device based on diagnostics/prediction
- Recognition of abnormal events in the operation or state of a system



AI/ML driven multi-objective optimization frameworks




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
TORAX, GOOSE		Specific to fusion

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	All EU experimental fusion facilities
Additional Test Facilities	Additional Test Facilities - List
Yes	Infrastructure providers

Computational & Infrastructure Requirements

GPUs
Cloud
HPC

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	EUROFUSION ITER EUROHPC CENTERS	US ENTITIES ASIA (i.e., China & South Korea) UKAEA (e.g., Digilab) Private Companies (e.g., DeepMind (Google), Microsoft Research, AWS, OpenAI)

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
AI/ML driven multi-objective optimization frameworks		Yes	European universities, research centers, EUROFUSION, ITER, F4E, private companies	40 to 80%	>2 years	>100k



Categorization and clustering of CAD clashes




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
ITER + not open		General technology, can be applied to other domains

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER, F4E
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

large memory for CAD, data pipelines for 3D geometry + metadata, GPU

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	EUROfusion, ITER, UK Atomic Energy Authority, CEA, Max Planck Institute for Plasma Physics, General Atomics, Bechtel, Autodesk, Siemens	Autodesk and Siemens (BIM/CAD ecosystems), Bechtel (large-scale infrastructure execution)

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------



Component tagging for 3D-scanning



CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
ITER		General technology, can be applied to other domains

MAJOR OBSTACLES

none	33%
spread-out data	33%
low explainability	0%
limited resources	33%
low model robustness	67%
regulatory/ethical	0%
scalability issues	0%
high uncertainty	0%
tool dependency	67%

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER, F4E
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

High GPU usage for point-cloud deep learning, large storage for 3D scans, preprocessing pipelines for point clouds/meshes

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
--------------------	-------------------------	-----------------	------------------

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------



Generation of E3D or 3D drawings



CHARACTERISTICS

DATASET


Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
ITER		General technology, can be applied to other domains

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER, F4E

Additional Test Facilities	Additional Test Facilities - List
No	/

MAJOR OBSTACLES



Computational & Infrastructure Requirements

GPU workstations, CAD integration worklines

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
--------------------	-------------------------	-----------------	------------------

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------



Optimisation of magnetic fields using AI-guided inverse design algorithms




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Synthetic dataset	Yes	Yes
Owner of dataset + open		Broader market potential
/		Specific to fusion

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	PPPL (US) CIEMAT (SPAIN) and PROXIMA FUSION (GERMANY) MAYBE have databases and or tools about that

Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

/

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
/	/	/	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------



P&ID recognition for engineering data extraction and consistency




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
ITER		General technology, can be applied to other domains

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER, F4E
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

GPU workstations

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	AWS, Microsoft Azure, Hexagon, Bentley	

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------



Quality control of CAD drawings




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
ITER + not open		General technology, can be applied to other domains

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER, F4E

Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements
/

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
/	/	/	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------



Active learning for data-efficient experimentation




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
open literature, EUROfusion, labs, research institutions		General technology, can be applied to other domains: chemistry, biology, engineering fields, useful in problems with costly evaluations, applicable in all domains with a potential data drift

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	Tokamak, HPC centers
Additional Test Facilities	Additional Test Facilities - List
Yes	Additional test facilities would be desirable

Computational & Infrastructure Requirements

Data storage

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	Tecnia, specific research institutes	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------



Digital twins of fusion reactors




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Synthesize more data	No
Owner of dataset + open		Broader market potential
research institutes, EUROfusion		General technology, can be applied to other domains: aerospace, defense, transportation, machinery D&E, health monitoring

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	All fusion devices/facilities in Europe
Additional Test Facilities	Additional Test Facilities - List
Yes	IFMIF-DONES, long pulse tokamaks, particle accelerators that share data for the fusion energy field

Computational & Infrastructure Requirements

More HPCs would be desirable. A public data platform should be available to share, maintain, and manage data within the community

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	Every institute is currently working on this.	Some EU players, but expertise generally could be Deep Mind

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Accessible database to store and share data for simulation and experiments	Applied to fusion Standarization is key Fast queries Easily accessible, not necessarily public, which would foster its use	Yes	Some experts are available in Europe but the expertise is broad in this domain, not rocket science, so the expertise here is not a limiting factor for this TDA	>80%	6 months to 2 years	>100k
Creation of a complete Digital Twin for Fusion		Yes	Most labs in Europe with expertise on AI and Fusion	40 to 80%	>2 years	>100k



GenAI for code template generation



CHARACTERISTICS


DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
/		General technology, can be applied to other domains

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	EUROfusion IM workflows, HPC clusters
Additional Test Facilities	Additional Test Facilities - List
Yes	Benchmark repositories, surrogate-model validation sandboxes, GPU testbeds

MAJOR OBSTACLES



Computational & Infrastructure Requirements

GPU infrastructure, HPC integration, code repositories, CI/CD pipelines, storage for simulation datasets

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	EUROfusion, ITER, Max Planck Institute for Plasma Physics, UK Atomic Energy Authority, NVIDIA, GitHub, OpenAI	NVIDIA, EUROfusion, UK Atomic Energy Authority

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------



GenAI for synthetic data generation




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
various data owners, research institutions		General technology, can be applied to other domains

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	HPC systems
Additional Test Facilities	Additional Test Facilities - List
Yes	GPU and TPU centers

Computational & Infrastructure Requirements

GPU and TPU centers

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	Tecnia, specific research entities	Fusion specific: none Generally: Meta

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------



Hybrid models and Physics-Informed Machine Learning for simulations



CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Synthetic dataset	Synthesize more data	Yes


Owner of dataset + open

data owners, research institutions, EUROfusion

Broader market potential

General technology, can be applied to other domains: machinery, transportation, fluid dynamics, astrophysics, geophysics, optics, medicine, aerospace

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities

Yes

Existing Test Facilities - List

Fusion devices

Additional Test Facilities

Yes

Additional Test Facilities - List

HPC systems

Computational & Infrastructure Requirements

HPC resources

ENTITIES

Entities in the EU

>5

Entities Outside the EU

>5

Entities - List

Every research entity


Entities Leading

No one in specific

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------

ML-driven surrogate modelling



CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Synthetic dataset	Synthesize more data	Needs preprocessing (in PDFs, unlabeled images...)


Owner of dataset + open

Euratom, data owners, research institutions, EUROfusion

Broader market potential

General technology, can be applied to other domains: transportation, machinery, health monitoring, D&E, geochemistry, hydrogeology

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities

Yes

Existing Test Facilities - List

Tokamak devices

Additional Test Facilities

Yes

Additional Test Facilities - List

There are already test facilities to develop this technology, but more would be better

Computational & Infrastructure Requirements

More data storage and availability
Need for standardized preprocessing of data

ENTITIES

Entities in the EU

>5

Entities Outside the EU

>5

Entities - List

Every research entity or university

Entities Leading


Not a leading company or entity.
Outside fusion, it would be Europe in general.

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Accessible database to store and share data for simulation and experiments	Applied to fusion Standarization is key Fast queries Easily accessible, not necessarily public, which would foster its use	Yes	Some experts are available in Europe, but since the required expertise is broad and not highly specialized, it's not a limiting factor for this TDA	>80%	6 months to 2 years	>100k
Use of surrogate modelling in environments where we have data/expertise	Surrogate modelling is already well established in aeronautics, and a similar approach could be applied in fusion. Using this methodology would also help promote its adoption by demonstrating parallels with other sectors. (Estimated TRL 3–4 after the TDA, potentially reaching TRL 6 later)	Yes	Specialized labs on this topic	>80%	6 months to 2 years	>100k



AI-guided discovery and development of new materials




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Yes
Owner of dataset + open		Broader market potential
literature, academia, HI IBERIA platform, materials simulation repositories, labs, enterprises, industry + open		General technology, can be applied to other domains: fission and accelerator target materials, defense, medical, automotive, space

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	SCK CEN
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

/

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	/	/	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
AI for the discovery of new materials	AI to discover new materials with dedicated properties, which are demanded by design and manufacturing constrains. The task is focused on proposing a short list of new materials to be tested and validated by physical testing. No manufacturing or testing included on this TDA (possibly on a different one).	Yes	DAES SA, SCK CEN, HI IBERIA, TECNALIA, BARCELONA SUPERCOMPUTING CENTER	>80%	6 months to 2 years	>100k



Generation and curation of materials databases




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Synthetic dataset	Yes	No
Owner of dataset + open		Broader market potential
F4E, labs, academia, industry, databases + open (or open on request)		General technology, can be applied to other domains: fission, research on materials, defense, industry

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	- Qualified or accredited laboratories for material characterization, for example: all members of the EUROFUSION consortium, and others. - Irradiation facilities, for example: Material Test Reactors such as fast and thermal nuclear reactors, Ion Beam, spallation sources, fusion spectrum and others. - High Heat Flux Test Facilities - Hot Cell Test Facilities
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

High Performance Computing

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	TECNALIA and other similars.	Joint Research Center (JRC); KIT (Germany) and Petten (Netherlands); ITER organization; EUROFUSION; PSI - SINQ (Switzerland), ESA (European Space Agency), SCK CEN / BR2 (Belgium), CEA (France) among others.

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Material Database Generation	Material Database Generation with potential for codification for introduction into the design standards (e.g. ASME, ITER SDC-IC);	Yes	Entities with expertise on the need of identification of properties of material for introduction on design codes. Also experience on generating and managing Databases and mastering AI tools.	>80%	6 months to 2 years	>100k



Material acceptability prediction
for a dedicated application



CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Synthetic dataset	Yes	No


Owner of dataset + open

CEA, CERN (sgobba), Il barabash, Oakridge test reactor, Paul Scherrer Institute (PSI), academia, material manufacturers, F4E, labs + not open or open upon request

Broader market potential

General technology, can be applied to other domains: all application using materials, fission, accelerators, space, defense, research

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities

Yes

Existing Test Facilities - List

- Qualified or accredited laboratories for material characterization, for example: all members of the EUROFUSION consortium, and others.
- Irradiation facilities, for example: Material Test Reactors such as fast and thermal nuclear reactors, Ion Beam, spallation sources, fusion spectrum and others.
- High Heat Flux Test Facilities
- Hot Cell Test Facilities

Additional Test Facilities

Yes

Additional Test Facilities - List

IFMIF-DONES

Computational & Infrastructure Requirements

Standard computing facilities

ENTITIES

Entities in the EU

0

Entities Outside the EU

0

Entities - List

/

Entities Leading


/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Usage of AI tools to enhance the introduction of new materials on (nuclear) design codes.	Use of AI tools to support the introduction of new materials into (nuclear) design codes. Requires deep knowledge of material qualification standards and AI expertise, with successful completion of the "Generation and curation of materials database" as a prerequisite.	Yes	DAES SA, Tecnalia, HI Iberia, SCK CEN, Barcelona Supercomputing Center	40 to 80%	>2 years	>100k



Prediction of behaviour and mechanical properties of new materials




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	No
Owner of dataset + open		Broader market potential
research labs, private companies, academia + synthetic is open, experimental is not		General technology, can be applied to other domains: space, fission, defense, industry

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	- Qualified or accredited laboratories for material characterization, for example: all members of the EUROFUSION consortium, and others. - Irradiation facilities, for example: Material Test Reactors such as fast and thermal nuclear reactors, Ion Beam, spallation sources, fusion spectrum and others. - High Heat Flux Test Facilities - Hot Cell Test Facilities
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

/

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
/	/	/	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
AI tool to predict physical and mechanical properties based on partially available data	Some material tests (e.g., cyclic loading, time-dependent properties, irradiation ageing) are lengthy and resource-intensive. This AI tool aims to predict such properties using limited available data. Example TDA sub-tasks include: - Unguided classification of existing materials - Structuring and organizing available data	Yes	DAES SA, Tecnalia, HI Iberia (IFMIF-DONES), Barcelona Supercomputing Center, SCK CEN	40 to 80%	6 months to 2 years	>100k



AI-driven acoustic emission for non-metallic material inspection




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Synthesize more data	Yes
Owner of dataset + open		Broader market potential
labs, universities		General technology, can be applied to other domains: construction, maintenance, quality control, structural health moniting, traditional NPPs

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	Several facilities across a wide range of fields
Additional Test Facilities	Additional Test Facilities - List
Yes	/

Computational & Infrastructure Requirements

Moderate computational needs

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	Already with EU Horizon OpenAE - sharing between entities.	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Acoustic Emission optimization	There is already ongoing work with the EU for example					



AI-driven distortion prediction




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
No	No	No
Owner of dataset + open		Broader market potential
/		General technology, can be applied to other sectors, cross-domain technology

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Partially	/
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

/

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
/	/	/	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Pilot case (TBD) for relevant application to use as starting platform	Needs development from scratch. Standardize methodology to "teach AI".					



AI-driven process optimization
for advanced manufacturing




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Yes
Owner of dataset + open		Broader market potential
research centers, universities, machine/hardware manufacturers, software developers, private companies		General technology, can be applied to other sectors

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	Present in almost every country through AM workshops and development, but with few actual manufacturers of AM machines
Additional Test Facilities	Additional Test Facilities - List
Yes	Labs, universities: high variety of PhD work globally

Computational & Infrastructure Requirements

Accessible and available on local computers.
Lots of data processed locally as AM needs programming for each step (micrometer by micrometer). Straight forward.

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	Tecnia, ...	Can be any individual or consortia - versatily

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
AI driven optimization of Additive Manufacturing	Develop tool to process and utilize datasets during manufacturing. Adopt design for AM with automatic support by AI.	Yes		>80%	6 months to 2 years	> 100k
AI optimization of manufacturing processes supported by acoustic sensor data	- Identifying potential failures - Optimizing costs by defining when components meet "good enough" criteria - In-situ validation of manufacturing processes (e.g., monitoring sound waves to assess material maturity in concrete, glues, metals, polymers) - Verifying success of heat or manufacturing treatments by measuring sound waves, calibrating with reference data, and enabling AI-driven compensation	Yes	Exists in different fields with different TRLs and maturities	>80%	6 months to 2 years	> 100k



Automated digital RT data processing for weld inspection



CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Yes
Owner of dataset + open		Broader market potential
research institutions, manufacturers		General technology, can be applied to other domains

MAJOR OBSTACLES

none	50%
spread-out data	0%
low explainability	0%
limited resources	0%
low model robustness	0%
regulatory/ethical	50%
scalability issues	0%
high uncertainty	0%
tool dependency	17%

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	Versatile use in many technology fields
Additional Test Facilities	Additional Test Facilities - List
Yes	Supported by labs and universities

Computational & Infrastructure Requirements

Not limiting

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	/	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Digital optimization of RT Processes		Yes	Develop tool to complement inspectors and use as verification step	>80%	6 months to 2 years	>100k



Automated PAUT data processing for weld inspection




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Yes
Owner of dataset + open		Broader market potential
ITER, software developers, manufacturers		General technology, can be applied to other sectors

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER, manufacturers on PAUT and from different fields

Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements
/

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
/	>5	UT inspection is done in many fields continuously generating lots of data that is normally available without IP issues	EU, labs and inspection related companies as Bureau Veritas

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
PAUT AI Optimization	Develop defect and anomaly detection during PAUT with AI. Supported by manual inspection findings (manual inspection is mandatory in many technologies and will make the implementation phase quick and reliable)	Yes	Know-how available in many areas	>80%	6 months to 2 years	>100k



Real-time optimization of manufacturing parameters




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Yes
Owner of dataset + open		Broader market potential
research institutions, manufacturers		General technology, can be applied to other domains

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	Manufacturing process optimization is ongoing in many places, but IP is strong and limiting
Additional Test Facilities	Additional Test Facilities - List
/	/

Computational & Infrastructure Requirements

High-end computer with storage

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	Very high number of companies and organizations working on it, but in very diverse fields. Labs as Tecnalía can share data if not linked to specific funded projects restricted by clients (available data should be enough to get the technology development started)	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Process Optimization for Distortion Prediction for Additive Manufacturing	Predict influence of manufacturing process on geometrical deviations. - Pre-compensate via AI supported predictions. - Compensate in-situ. - Real time Metrology feedback to AI.	Yes	- Universities and labs develop new, publicly accessible process algorithms - Companies develop processes tied to IP, availability varies case by case - Existing processes are material-specific; not all materials covered - Requires selection of both material and end-user need, as well as machine/tool - Rapidly growing global know-how - Each project generates lessons learned applicable across manufacturing methods	>80%	6 months to 2 years	>100k
Real-Time Manufacturing optimization	Real-time monitoring, correction and "tweaking" of parameters for best outcome	Yes		>80%	6 months to 2 years	>100k



Welding success rate prediction



CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Synthesize more data	Yes
Owner of dataset + open	Broader market potential	
universities, manufacturers	General technology, can be applied to other domains: welding, automotive	

MAJOR OBSTACLES

none	33%
spread-out data	17%
low explainability	0%
limited resources	0%
low model robustness	0%
regulatory/ethical	50%
scalability issues	0%
high uncertainty	0%
tool dependency	0%

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	EU, labs, universities and companies in many fields
Additional Test Facilities	Additional Test Facilities - List
/	/

Computational & Infrastructure Requirements

Moderate need of high end computer


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	Tecnalia	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Welding optimization and in-situ compensation	Brainstorming ideas: - In-situ automated compensation of welding parameters - Challenge: parameters are validated/specified, not the process itself → limits AI's ability to adjust inputs - Potential for large-scale, repetitive, and reproducible welding processes - First step: pilot project, though implementation may be challenging - In-situ measurement tools for defects (NDT, metrology) and distortions → explore predictive models to train AI - Capture lessons learned after each process to improve the next	Yes	Top Level welding institute for example in collaboration with industry	40 to 80%	6 months to 2 years	>100k
Welding Process Optimization	Select a welding process with reproducible production as starting point, not complex state of the art welding					

AI-driven extrapolation of
operating conditions from local
testing



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	No
Owner of dataset + open		Broader market potential
suppliers, industry, ITER, F4E + not open		General technology, can be applied to other sectors

MAJOR OBSTACLES

- none
- spread-out data
- low explainability
- limited resources
- low model robustness
- regulatory/ethical
- scalability issues
- high uncertainty
- tool dependency

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	Any existing testing facility
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

GPU, storage, digital twins, virtual commissioning

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
<5	<5	Airbus	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------



Automated mapping of metrology cloud points to component models




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
suppliers, industry, ITER, F4E + not open		General technology, can be applied to all metrology applications

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER, entities doing scanning
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

GPU (a lot), data storage,

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
<5	<5	Software providers (Hexagon, Innovmetric,...)	Software providers (Hexagon, Innovmetric,...)

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Model to automatize the process of alignment cloud to CAD	Challenge: input data may be too large - Clean input data (without the CAD) and filter useful data (select primitive geometries) - Pre align with CAD - Finalise data cleaning - Final alignment - Output: transformation matrix (pointcloud)	Yes	Hexagon, Innovametric, any other software that deals with pointcloud	>80%	6 months to 2 years	>100k



Detect modifications on drawings



CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	No	No
Owner of dataset + open		Broader market potential
F4E, ITER, suppliers + not open		General technology, can be applied to other domains

MAJOR OBSTACLES

none	0%
spread-out data	33%
low explainability	33%
limited resources	0%
low model robustness	33%
regulatory/ethical	0%
scalability issues	0%
high uncertainty	33%
tool dependency	33%

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	Environment is existing, normal PCs could manage it. No special infrastructure is needed.
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

GPUs, Storage


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	Any industry with over 20–30 years of experience that is now seeking to digitalize its data	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Agent to identify, compare and provide reasoning behind modifications on DWGs based on AI algorithm	- Prepare data: set up drawings with modifications to feed in the model. Operators might be needed to be trained on a mark-up code - Prepare a report on the modifications on DWGs and root cause of the changes (if possible) - Be able to interact with user - Be able to track result/accuracy of the results. Time to implement depends on the complexity of the Agent to provide reasonable (sound) rationals of the modifications.	Yes	Any AI agent developer	>80%	<6 months	50k to 100k

Identification of most optimal assembly sequence




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Synthetic dataset	Synthesize more data	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open airbus, F4E, ITER, suppliers + not open		Broader market potential General technology, can be applied to other domains impacted by assembly sequences

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER, mock-ups, digital twins, other big complex projects
Additional Test Facilities	Additional Test Facilities - List
Yes	If we use a digital twin, a virtual environment would be needed

Computational & Infrastructure Requirements

GPU, storage, real-time processing


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
<5	<5	Dassault Systèmes	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------

Instruments and transport
virtual agent




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	No
Owner of dataset + open		Broader market potential
Copilot agent		General technology, can be applied to other domains

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	Digital AI tools (Copilot, Gemini...)

Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

Connection to the net


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	OpenAI, Google, Anthropic, Deepseek	/

TECHNOLOGY DEVELOPMENT ACTIONS


TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Virtual travel agent for metrology	An agent that is able to organise instrument booking and logistics for metrology instruments, integrated with teams/outlook or other relevant tools, that allows booking of instruments for different metrology actions, and booking of shipping or travel needs. - Booking of instrument - Recording start and endpoints for logistics - Contacting logistic companies for quotations - Reserving shipping - Recording the instrument position in a calendar Payment implementation could be difficult.	Yes	Any AI developer	>80%	<6 months	50k to 100k

Optimization of component positioning during assembly



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
F4E, ITER, suppliers + not open		General technology, can be applied to other domains



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER or other complex projects (DTT in Italy, DONES, JT60, CFS)
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements
Data storage, real-time processing, GPU


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
<5	<5	/	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Assembly procedure optimizer	This TDA is also related to "Identification of most optimal assembly procedure" AI model capable of providing a visual assembly sequence of the assembly of (complex) systems, being able to simulate the different steps taking into account (constraints parameters, dimensional compliance, tolerance, schedule, cost, manufacturing techniques, health and safety constraints, non conformities...) - Alert operators of high risk steps - Compliance with assembly tolerance - Adherence with 3D model - Feed the model with design data, and then feed with as-built dimensions - Feed with the time to due assembly, number of resources, - Feed with assembly procedure tools	Yes	Input data to be gathered. CT Engineering Group, or other engineering companies in consortium with AI developers	40 to 80%	>2 years	>100k

Prediction of metrology data




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
F4E, ITER, private Companies, metrology service providers + not open		General technology, can be applied to other domains

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	Workshops

Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

Data storage and GPU

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
<5	<5	/	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Build an AI model to predict dimensional metrology data	Approach: - Consolidate all available data (both accessible and hidden) - Split datasets into training and validation sets - Begin with a defined component/manufacturing process, then extend to cassette bodies (evaluate hidden parts when closed), Vacuum Vessel, upper launcher (mirror positioning after final assembly) Implementation timeline depends on data availability.	Yes	Geatop, CT Engineering Group, Hexagon	>80%	6 months to 2 years	50k to 100k



Auto-labelling AI tools for plasma state datasets




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Synthetic dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
/		Specific to fusion

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	JET
Additional Test Facilities	Additional Test Facilities - List
Yes	Data needs to be labelled and more data is needed from existing facilities

Computational & Infrastructure Requirements

Same infrastructure as for "Fusion data platform for Machine Learning Applications" (this being a pre-requisite for labelling), data experts are required

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
1	<5	EU: DEFUSE (EUROfusion/EPFL) Outside EU: DECAF (closed source, not accessible, might require a collaboration agreement)	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Machine agnostic AI Tool for Data Annotation	Core functions: data curation, validation, and annotation Common architecture: token-based access and shared framework Specialized tools (e.g., labeling for fusion) Potential reuse of and adaptation from existing tools Tools can also be embedded in data platforms (secondary objective) Development time depends on tool reuse, regulatory constraints, and platform progress	Yes	Savantic AI Lab EUROfusion affiliated labs DEFUSE (some modules could be leveraged for annotation)	>80%	6 months to 2 years	>100k



Fusion data platform for ML applications




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Synthetic dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
Individual git repositories, IAEA Data Lake, ITER/IMAS, GA-FDP, ITER (rudimentary)		Specific to fusion

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	Referring to IMAS, we have access to WEST as a machine to produce data. Availability will need to be checked for members not belonging to ITER or EUROfusion networks. Effort to be put to share all the existing data to entities working in this domain. Question about the type of agreements that need to be put in place to assure overall access to data.
Additional Test Facilities	Additional Test Facilities - List
Yes	Extension to fully cover the available test facilities (WEST). IMAS is just a building block of the data platform. There exist several scattered efforts that need to be coordinated and/or integrated. There is an overall need of standardization in model curation processes, metadata tracking, exchange of information, etc.

Computational & Infrastructure Requirements

Strongly dependent on the structure defined for the data, etc. Need for physical space for hosting the data and provide the access to the community, based on an industry standard for data sets.

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
0	1	General Atomics in collaboration with Hewlett Packard Enterprise	General Atomics in collaboration with Hewlett Packard Enterprise

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Advanced AI visualization tools	To visualize modelization and simulation outputs to non-specialized public; to allow for benchmark analysis. Could be part of the data platform or a separate stand alone tool. Data compression for better visualization and clustering.	Yes	EPFL for instance.	>80%	6 months to 2 years	50k to 100k
Fusion Data Platform	Comprehensive architecture enabling ML applications and development, in line with IMAS (standardization, curation, access to metadata information, provenance of data), clear and well-defined platform architecture based on industry standard tools, open access, version controlled, feature store, evolutions and traceability (by the community) A standardisation policy inside EUROfusion members to ease the access to data, on top of any MoU for external parties.	Yes	Existing tools available, e.g. common metadata frameworks (Hewlett Packard). Savantic AI, VTT	>80%	6 months to 2 years	>100k
Governance on coordinated data user agreements	Process to assure data accessibility not only within EUROfusion partners, but extensive to startups and globally via bilateral agreements, MoU or another formulas.	Yes	EUROfusion, F4E, ITER IO, IAEA cooperation framework	40 to 80%	<6 months	<50k



Machine learning-accelerated pedestal MHD stability evaluations



CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Synthetic dataset	Synthesize more data	Yes

Owner of dataset + open

EUROfusion, individual researchers, research groups + both open and not open

Broader market potential

Specific to fusion

MAJOR OBSTACLES

none	0%
spread-out data	0%
low explainability	0%
limited resources	67%
low model robustness	0%
regulatory/ethical	0%
scalability issues	100%
high uncertainty	0%
tool dependency	0%

FACILITIES

Existing Test Facilities

Partially

Existing Test Facilities - List

Existing integrated simulations

Additional Test Facilities

Yes

Additional Test Facilities - List

Might depend on the use cases and on the validation definition and objectives

Computational & Infrastructure Requirements

Sufficient HPC capacity

ENTITIES

Entities in the EU

>5

Entities Outside the EU

>5

Entities - List

EUROfusion and all the other entities listed for PM technology

Entities Leading


/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Fast surrogates for fusion scenario simulations	To replace high-fidelity simulations by machine learning models: pedestal simulations, MHD, turbulence, exhaust, neutrals, sources, etc.	Yes	Various labs working already in individual models, some coordination might be sought but they might have different settings e.g. DIFFER, VTT, UKAEA, CALDERA AI	>80%	>2 years	>100k



Predictive modelling of plasma instabilities and disruptions




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
some diagnostic groups might have ownership of data, ITPA, EUROfusion Disruption Database + open to members		General technology, can be applied to other domains

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	The issue is not about the test facilities but about curation and exploitation of existing data
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

Need of additional HPC computing means to be dedicated to researchers

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	>5	EUROfusion, VTT, CIEMAT, CEA, Max Planck, DIFFER, MIT, UKAEA Princeton, Columbia, University of Rome, University of Cagliari	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
AI Enabled Scientific Discovery	Leverage on AI algorithms to allow for scientific automated discovery					
Event Prediction, Detection and Classification	Tool that is able to catch more than disruptive events like for instance anomalies, scenario changes, transition confinement events, and any other interesting events for optimization (to get alerts) (Also connection to auto-labelling TDA, same family models. Outcome of auto-labeler TDA will be relevant to this TDA and viceversa). Some of the applications might require real time.	Yes	EUROfusion groups (several) working on predition, classification, etc. Savantic	>80%	6 months to 2 years	>100k



Time-resolved, PINNS for tokamak total emission reconstruction



CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
both open and not open		Specific to fusion

MAJOR OBSTACLES

none	0%
spread-out data	0%
low explainability	0%
limited resources	100%
low model robustness	100%
regulatory/ethical	0%
scalability issues	0%
high uncertainty	0%
tool dependency	0%

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	Referring to IMAS, we have access to WEST as a machine to produce data. Availability will need to be checked for members not belonging to ITER or EUROfusion networks. Effort to be put to share all the existing data to entities working in this domain. Question about the type of agreements that need to be put in place to assure overall access to data.
Additional Test Facilities	Additional Test Facilities - List
Yes	Extension to fully cover the available test facilities (WEST). IMAS is just a building block of the data platform. There exist several scattered efforts that need to be coordinated and/or integrated. There is an overall need of standardization in model curation processes, metadata tracking, exchange of information, etc.

Computational & Infrastructure Requirements

Strongly dependent on the structure defined for the data. Need for physical space for hosting the data and provide the access to the community, based on an industry standard for data sets.

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
0	1	General Atomics in collaboration with Hewlett Packard Enterprise	General Atomics in collaboration with Hewlett Packard Enterprise

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------



Trajectory planning enabled by ML/AI




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Synthetic dataset	Synthesize more data	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
individual researchers or groups		Specific to fusion

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	Limited validation efforts are currently underway for this technology, but future work will be needed to integrate experimental results
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

Existing resources in term of HPC are considered sufficient


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
<5	<5	DeepMind, EPFL, DIFFER Some private companies and startups (Next Step Fusion) MIT, Princeton University, CFS	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------


AI-driven validation of requirements



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open	Broader market potential	
ITER, domestic agencies	General technology, can be applied to other domains	

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER, JT-60SA
Additional Test Facilities	Additional Test Facilities - List
Yes	IFMIF-DONES in near future, and several industrial projects

Computational & Infrastructure Requirements

Already available in ITER

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
<5	<5	Automaticity Solutions, Capgemini	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
AI-enhanced requirements checking	PHASE 1: A prototype will be developed to check the propagation of a chosen sSRD to its PA (from F4E). The % of correctly identified requirements will be the accuracy parameter. --- Data set: sSRD, PA and Applicable Documents (all data available in ITER). PHASE 2 (TRL 4): - the developed prototype is implemented in the other sSRDs (so the same level of propagation but different systems). PHASE 3 (TRL 5): - go one level up(or down), apply same methodology and check if changes have to be applied in the process. PHASE 4 (TRL 6): - do full propagation of a given system PHASE 5 (TRL 7): - do full propagation of all systems	Yes		40 to 80%	6 months to 2 years	>100k



Chatbot for requirements rationalization




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
ITER, EURATOM RTD, IPP, CEA, EPFL, CCFE		General technology, can be applied to other domains: fission, aerospace, construction (also originates from other domains)

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER, JT-60SA, IFMIF Dones
Additional Test Facilities	Additional Test Facilities - List
Yes	Complex industries

Computational & Infrastructure Requirements

Already available in ITER


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
<5	<5	Automaticity solutions	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Chatbot for requirements	Chatbot to assist in clarifying requirements and speed up analysis could be developed in parallel with the other solutions, nevertheless rationalization of requirements shall be tackled in a different solution after implementing requirement propagation and validation.	Yes	Chatbot done for design review input package validation in ITER project	40 to 80%	6 months to 2 years	>100k


LLM-driven check of requirements propagation



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	No
Owner of dataset + open		Broader market potential
ITER, domestic agencies + open to members		General technology, can be applied to other domains

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	/

Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

Large knowledge graph stored in links, information, ...


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
<5	<5	Automaticity Solutions, Capgemini, Cognizant (USA)	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
LLM-driven check of requirements propagation		Yes		40 to 80%	6 months to 2 years	>100k

AI-powered proposal support tool for project disruptions



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset		
Owner of dataset + open	Broader market potential	
not open	General technology, can be applied to other sectors	

MAJOR OBSTACLES

none	40%
spread-out data	40%
low explainability	0%
limited resources	0%
low model robustness	0%
regulatory/ethical	0%
scalability issues	0%
high uncertainty	40%
tool dependency	0%

FACILITIES

Existing Test Facilities	Existing Test Facilities - List

Additional Test Facilities	Additional Test Facilities - List

Computational & Infrastructure Requirements


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost


Automated contract generation



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
ITER + not open		General technology, can be applied to other domains

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	F4E and ITER organization
Additional Test Facilities	Additional Test Facilities - List
Yes	To move the model to private contracts, private companies are needed for training and validating the model

Computational & Infrastructure Requirements

Limited computational power necessary, LLMs


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
/	/	ITER organization have tried to develop this technology	Major AI companies like Microsoft, as well as startups, may already be developing such models, which could also be integrated into ERP systems

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Design a framework for sharing contract data	Define a framework, where companies can be motivated to share their contract data.	Yes	legal experts, procurement experts	<40%	>2 years	>100k
Identify the relevant legal framework	Identify legal framework, which can serve as a basis for generation new contracts	Yes	legal experts, procurement experts	>80%	6 months to 2 years	50k to 100k
Market survey	Define the scope of the AI model. Conduct a market survey, if there are already models available and if yes, what is there status.	Yes	Procurement specialists, legal specialists, IT	>80%	<6 months	50k to 100k
Preparation of a Dataset for a first Proof of Concept	Use IO or F4E dataset for a first proof of concept. Test with exisiting systems like AI powered contract review, legittai, aline, legito, genAI, amto, typli etc.	Yes	IO / F4E	40 to 80%	<6 months	50k to 100k
Validate tool after development	Validate tool in different companies and different types of contract courses	Yes	legal, procurement, IT	40 to 80%	6 months to 2 years	>100k

Automatic status reporting from data sources



CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
Copilot, can be achieved without AI, open and proprietary tools		General technology, can be applied to other domains

MAJOR OBSTACLES

none	100%
spread-out data	0%
low explainability	0%
limited resources	0%
low model robustness	0%
regulatory/ethical	0%
scalability issues	0%
high uncertainty	0%
tool dependency	0%

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	F4E and ITER organization
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

LLMs, Agents, custom model, standard computational power


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
/	/	Some organizations are reportedly pursuing implementation. In the context of ERP systems, such AI capabilities should ideally be offered as an integrated feature	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Market survey	Define scope of AI model. Support the reporting and finding of hidden data, or other complicated to find data. Collect data from different systems. Create dashboards of daily current status. Conduct market survey on the current status of available tools for such applications.	Yes	Procurement experts	>80%	<6 months	50k to 100k
Test available models / Proof of concept	Select data relevant to different cases to test do model on these different use cases. Test available models on F4E/ IO data.	Yes	F4E/IO	40 to 80%	<6 months	50k to 100k

Prediction documentation review time



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	No
Owner of dataset + open		Broader market potential
ITER		General technology, can be applied to other domains

MAJOR OBSTACLES

none	50%
spread-out data	25%
low explainability	50%
limited resources	0%
low model robustness	0%
regulatory/ethical	0%
scalability issues	0%
high uncertainty	0%
tool dependency	0%

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	F4E & IO, every company with document management

Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

Limited amount of computational power

ENTITIES


Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
/	/	/	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost



Prediction of contract outcome based on past performance




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
F4E, ITER + not open		General technology, can be applied to other domains

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	F4E and ITER organization
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

Pretrained LLMs, limited computational requirements


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
/	/	/	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------

Prediction of tool usage



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
no conventional software tools available		General technology, can be applied to other domains

MAJOR OBSTACLES	
none	100%
spread-out data	0%
low explainability	0%
limited resources	0%
low model robustness	0%
regulatory/ethical	0%
scalability issues	0%
high uncertainty	0%
tool dependency	0%

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	F4E and ITER

Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements
moderate compute (CPU-heavy with optional GPU)


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
		SAP and Oracle	

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost

Process optimization tool



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
/		General technology, can be applied to other domains

MAJOR OBSTACLES

none	40%
spread-out data	60%
low explainability	20%
limited resources	0%
low model robustness	0%
regulatory/ethical	0%
scalability issues	0%
high uncertainty	0%
tool dependency	0%

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	F4E and ITER

Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

moderate-to-high CPU compute, large structured operational datasets


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
		Celonis, SAP, IBM	

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost

Risk register forecast prediction



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
/		General technology, can be applied to other domains

MAJOR OBSTACLES

none	0%
spread-out data	0%
low explainability	25%
limited resources	0%
low model robustness	0%
regulatory/ethical	0%
scalability issues	0%
high uncertainty	0%
tool dependency	75%

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	F4E and ITER
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

moderate compute (CPU + optional GPU for NLP), strong governance and traceability layers


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
			IBM and ServiceNow

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost


AI-assisted experimental design



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Yes
Owner of dataset + open		Broader market potential
WEST, LLNL		General technology, can be applied to other domains: defense, research, robotics, aerospace, ...

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	/
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

/


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
/	>5	SEED E-science Research , ELI Beamlines	Google (medical experiments)

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------


AI-driven diagnostic data integration



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Yes
Owner of dataset + open		Broader market potential
WEST, AUG, JET, JT60, ASDEX, West, Jet, TJ-II and JET		General technology, can be applied to other domains: automation, robotics, aerospace, defence; large scale experiments, weather predictions, big science, industrial plants

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	NBTF, ASDEX, TJII, JET, W7X, WEST
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

Data storage


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	/	RFX, IPP, CIEMAT, JET, W7X, CEA, SEED E-science Research	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------

Automatic generation of End-of-Manufacturing report



CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset		
Owner of dataset + open	Broader market potential	
all fusion supply chain actors	General technology, can be applied to manufacturing industry	

MAJOR OBSTACLES

none	0%
spread-out data	57%
low explainability	14%
limited resources	14%
low model robustness	14%
regulatory/ethical	0%
scalability issues	14%
high uncertainty	14%
tool dependency	14%

FACILITIES

Existing Test Facilities	Existing Test Facilities - List

Additional Test Facilities	Additional Test Facilities - List

Computational & Infrastructure Requirements


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost


Dimensionality reduction for streaming data



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Yes
Owner of dataset + open		Broader market potential
TJ-II and JET		General technology, can be applied to other domains

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	/
Additional Test Facilities	Additional Test Facilities - List
/	/

Computational & Infrastructure Requirements

Computational power


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
1	/	W7X (solution exists, needs more development)	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Implement and validate developed algorithm on operating machines		Yes	W7X	40 to 80%	6 months to 2 years	>100k


Fusion specific documentation management



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Yes
Owner of dataset + open		Broader market potential
ITER, F4E		General technology, can be applied to other domains with a need for documentation: aerospace, fission, automotive

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER

Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements


/

ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
/	/	/	/


TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Build on ITER AI assistant by giving it access to more data						
Develop MCP for communication between AI assistants of "subscribers" (or open source)		Yes	SAVANTIC, CIEMAT, ITER	>80%	6 months to 2 years	>100k
Model for automatic creation of doc metadata						



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Yes
Owner of dataset + open		Broader market potential
TJ-II, F4E, WEST, JET, ITER		General technology, can be applied to other domains with a need for documentation: aerospace, fission, automotive



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER and F4E
Additional Test Facilities	Additional Test Facilities - List
No	/


Computational & Infrastructure Requirements
Not limiting

ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	/	ITER, F4E, UKAEA, CIEMAT	/


TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open		Broader market potential
Consortio rfx, FZ Juelich, West		General technology, can be applied to other domains



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	JET, ITER, ASDEX, TJII, NBTF
Additional Test Facilities	Additional Test Facilities - List
No	/

Computational & Infrastructure Requirements

Real-time access to data, specific sensors, homogenous and comprehensive data acquisition and analysis system, large data storage over long timeframe


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	/	/	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Develop proof of concept on fusion specific system	Building on experience from industrial systems	Yes	SAVANTIC, U of Ghent, NBTF	40 to 80%	6 months to 2 years	>100k
List potential areas of application and associated development needs	To prioritize fusion specific items	Yes	All fusion device operators	>80%	<6 months	<50k
Validate proof of concept for predictive maintenance		Yes	Machine operators	40 to 80%	6 months to 2 years	>100k

Real-time operation of fusion device based on diagnostics



CHARACTERISTICS


DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Yes
Owner of dataset + open		Broader market potential
MITech, General Atomics, JET, MIT		General technology, can be applied to other domains: defense

FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	TCV

Additional Test Facilities	Additional Test Facilities - List
/	/

MAJOR OBSTACLES



Computational & Infrastructure Requirements

Computing power and storage


ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
1	1	EPFL	General Atomics

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
----------	-------------	---------------	-----------------------	--------------------	---------------------	------

Real-time operation of fusion device based on prediction



CHARACTERISTICS

DATASET		
Availability	Size and diversity	Cleaning
Experimental dataset	Yes	Yes
Owner of dataset + open		Broader market potential
AUG, JET		General technology, can be applied to other domains


FACILITIES

Existing Test Facilities	Existing Test Facilities - List

Additional Test Facilities	Additional Test Facilities - List

Computational & Infrastructure Requirements

MAJOR OBSTACLES




ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost

Recognition of abnormal events
in the operation or state of a
system




CHARACTERISTICS

DATASET

Availability	Size and diversity	Cleaning
Experimental dataset	Synthesize more data	Needs preprocessing (in PDFs, unlabeled images...)
Owner of dataset + open	Broader market potential	
FZJUELICH, JET database, MIT	General technology, can be applied to other domains: aerospace, conventional energy generation, medicine, manufacturing	

MAJOR OBSTACLES



FACILITIES

Existing Test Facilities	Existing Test Facilities - List
Yes	ITER, JET, RFX, ASDEX, TJII
Additional Test Facilities	Additional Test Facilities - List
Yes	Simulator/digital twin

Computational & Infrastructure Requirements

Additional specific sensors for anomaly detection

ENTITIES

Entities in the EU	Entities Outside the EU	Entities - List	Entities Leading
>5	/	ITER, UKAEA, RFX, CIEMAT, SIEMENS, SAVANTIC, IPP (TBC)	/

TECHNOLOGY DEVELOPMENT ACTIONS

TDA Name	TDA Details	TDA Expertise	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Validate existing algorithms on operating machines	Get expert feedback during operation to gain and label training data	Yes	ITER, WEST, CIEMAT	40 to 80%	>2 years	>100k

TECHNOLOGY DEVELOPMENT ACTIONS (1)

Technologies	TDA Name	TDA Details	TDA Expertise Details	Chances of Success	Implementation Time	Cost
▲						
AI/ML driven multi-objective optimization frameworks	AI/ML driven multi-objective optimization frameworks		European universities, research centers, EUROFUSION, ITER, F4E, private companies	40 to 80%	>2 years	>100k
AI-driven acoustic emission for non-metallic material inspection	Acoustic Emission optimization	There is already ongoing work with the EU for example				
AI-driven distortion prediction	Pilot case (TBD) for relevant application to use as starting platform	Needs development from scratch. Standardize methodology to "teach AI".				
AI-driven process optimization for advanced manufacturing	AI driven optimization of Additive Manufacturing	Develop tool to process and utilize datasets during manufacturing. Adopt design for AM with automatic support by AI.		>80%	6 months to 2 years	>100k
AI-driven process optimization for advanced manufacturing	AI optimization of manufacturing processes supported by acoustic sensor data	- Identifying potential failures - Optimizing costs by defining when components meet "good enough" criteria - In-situ validation of manufacturing processes (e.g., monitoring sound waves to assess material maturity in concrete, glues, metals, polymers) - Verifying success of heat or manufacturing treatments by measuring sound waves, calibrating with reference data, and enabling AI-driven compensation	Exists in different fields with different TRLs and maturities	>80%	6 months to 2 years	>100k
AI-driven validation of requirements	AI-enhanced requirements checking	PHASE 1: A prototype will be developed to check the propagation of a chosen sSRD to its PA (from F4E). The % of correctly identified requirements will be the accuracy parameter. --- Data set: sSRD, PA and Applicable Documents (all data available in ITER). PHASE 2 (TRL 4): - the developed prototype is implemented in the other sSRDs (so the same level of propagation but different systems). PHASE 3 (TRL 5): - go one level up(or down), apply same methodology and check if changes have to be applied in the process. PHASE 4 (TRL 6): - do full propagation of a given system PHASE 5 (TRL 7): - do full propagation of all systems		40 to 80%	6 months to 2 years	>100k
AI-guided discovery and development of new materials	AI for the discovery of new materials	AI to discover new materials with dedicated properties, which are demanded by design and manufacturing constrains. The task is focused on proposing a short list of new materials to be tested and validated by physical testing. No manufacturing or testing included on this TDA (possibly on a different one).	DAES SA, SCK CEN, HI IBERIA, TECNALIA, BARCELONA SUPERCOMPUTING CENTER	>80%	6 months to 2 years	>100k
Auto-labelling AI tools for plasma state datasets	Machine agnostic AI Tool for Data Annotation	Core functions: data curation, validation, and annotation Common architecture: token-based access and shared framework Specialized tools (e.g., labeling for fusion) Potential reuse of and adaptation from existing tools Tools can also be embedded in data platforms (secondary objective) Development time depends on tool reuse, regulatory constraints, and platform progress	Savantic AI Lab EUROfusion affiliated labs DEFUSE (some modules could be leveraged for annotation)	>80%	6 months to 2 years	>100k

TECHNOLOGY DEVELOPMENT ACTIONS (2)

Technologies	TDA Name	TDA Details	TDA Expertise Details	Chances of Success	Implementation Time	Cost
Automated contract generation	Design a framework for sharing contract data	Define a framework, where companies can be motivated to share their contract data.	legal experts, procurement experts	<40%	>2 years	>100k
Automated contract generation	Identify the relevant legal framework	Identify legal framework, which can serve as a basis for generation new contracts	legal experts, procurement experts	>80%	6 months to 2 years	50k to 100k
Automated contract generation	Market survey	Define the scope of the AI model. Conduct a market survey, if there are already models available and if yes, what is there status.	Procurement specialists, legal specialists, IT	>80%	<6 months	50k to 100k
Automated contract generation	Preparation of a Dataset for a first Proof of Concept	Use IO or F4E dataset for a first proof of concept. Test with exisiting systems like AI powered contract review, legittai, aline, legito, genAI, amto, typli etc.	IO / F4E	40 to 80%	<6 months	50k to 100k
Automated contract generation	Validate tool after development	Validate tool in different companies and different types of contract courses	legal, procurement, IT	40 to 80%	6 months to 2 years	>100k
Automated digital RT data processing for weld inspection	Digital optimization of RT Processes		Develop tool to complement inspectors and use as verification step	>80%	6 months to 2 years	>100k
Automated mapping of metrology cloud points to component models	Model to automatize the process of alignment cloud to CAD	Challenge: input data may be too large 1. Clean input data (without the CAD) and filter useful data (select primitive geometries) 2. Pre align with CAD 3. Finalise data cleaning 4. Final alignment 5. Output: transformation matrix (pointcloud)	Hexagon, Innovametric, any other software that deals with pointcloud	>80%	6 months to 2 years	>100k
Automated PAUT data processing for weld inspection	PAUT AI Optimization	Develop defect and anomaly detection during PAUT with AI. Supported by manual inspection findings (manual inspection is mandatory in many technologies and will make the implementation phase quick and reliable)	Know-how available in many areas	>80%	6 months to 2 years	>100k
Automatic status reporting from data sources	Market survey	Define scope of AI model. Support the reporting and finding of hidden data, or other complicated to find data. Collect data from different systems. Create dashboards of daily current status. Conduct market survey on the current status of available tools for such applications.	Procurement experts	>80%	<6 months	50k to 100k
Automatic status reporting from data sources	Test available models / Proof of concept	Select data relevant to different cases to test do model on these different use cases. Test available models on F4E/ IO data.	F4E/IO	40 to 80%	<6 months	50k to 100k
Chatbot for requirements rationalization	Chatbot for requirements	Chatbot to assist in clarifying requirements and speed up analysis could be developed in parallel with the other solutions, nevertheless rationalization of requirements shall be tackled in a different solution after implementing requirement propagation and validation.	Chatbot done for design review input package validation in ITER project	40 to 80%	6 months to 2 years	>100k
Detect modifications on drawings	Agent to identify, compare and provide reasoning behind modifications on DWGs based on AI algorithm	- Prepare data: set up drawings with modifications to feed in the model. Operators might be needed to be trained on a mark-up code - Prepare a report on the modifications on DWGs and root cause of the changes (if possible) - Be able to interact with user - Be able to track result/accuracy of the results. Time to implement depends on the complexity of the Agent to provide reasonable (sound) rationals of the modifications.	Any AI agent developer	>80%	<6 months	50k to 100k

TECHNOLOGY DEVELOPMENT ACTIONS (3)

Technologies	TDA Name	TDA Details	TDA Expertise Details	Chances of Success	Implementation Time	Cost
▲						
Digital twins of fusion reactors	Accessible database to store and share data for simulation and experiments	Applied to fusion Standarization is key Fast queries Easily accessible, not necessarily public, which would foster its use	Some experts are available in Europe but the expertise is broad in this domain, not rocket science, so the expertise here is not a limiting factor for this TDA	>80%	6 months to 2 years	>100k
Digital twins of fusion reactors	Creation of a complete Digital Twin for Fusion		Most labs in Europe with expertise on AI and Fusion	40 to 80%	>2 years	>100k
Dimensionality reduction for streaming data	Implement and validate developed algorithm on operating machines		W7X	40 to 80%	6 months to 2 years	>100k
Fusion data platform for ML applications	Advanced AI visualization tools	To visualize modelization and simulation outputs to non-specialized public; to allow for benchmark analysis. Could be part of the data platform or a separate stand alone tool. Data compression for better visualization and clustering.	EPFL for instance.	>80%	6 months to 2 years	50k to 100k
Fusion data platform for ML applications	Fusion Data Platform	Comprehensive architecture enabling ML applications and development, in line with IMAS (standardization, curation, access to metadata information, provenance of data), clear and well-defined platform architecture based on industry standard tools, open access, version controlled, feature store, evolutions and traceability (by the community) A standardisation policy inside EUROfusion members to ease the access to data, on top of any MoU for external parties.	Existing tools available, e.g. common metadata frameworks (Hewlett Packard). Savantic AI, VTT	>80%	6 months to 2 years	>100k
Fusion data platform for ML applications	Governance on coordinated data user agreements	Process to assure data accessibility not only within EUROfusion partners, but extensive to startups and globally via bilateral agreements, MoU or another formulas.	EUROfusion, F4E, ITER IO, IAEA cooperation framework	40 to 80%	<6 months	<50k
Fusion specific documentation management	Build on ITER AI assistant by giving it access to more data					
Fusion specific documentation management	Develop MCP for communication between AI assistants of "subscribers" (or open source)		SAVANTIC, CIEMAT, ITER	>80%	6 months to 2 years	>100k
Fusion specific documentation management	Model for automatic creation of doc metadata					
Generation and curation of materials databases	Material Database Generation	Material Database Generation with potential for codification for introduction into the design standards (e.g. ASME, ITER SDC-IC);	Entities with expertise on the need of identification of properties of material for introduction on design codes. Also experience on generating and managing Databases and mastering AI tools.	>80%	6 months to 2 years	>100k
Instruments and transport virtual agent	Virtual travel agent for metrology	An agent that is able to organise instrument booking and logistics for metrology instruments, integrated with teams/outlook or other relevant tools, that allows booking of instruments for different metrology actions, and booking of shipping or travel needs. - Booking of instrument - Recording start and endpoints for logistics - Contacting logistic companies for quotations - Reserving shipping - Recording the instrument position in a calendar Payment implementation could be difficult.	Any AI developer	>80%	<6 months	50k to 100k

TECHNOLOGY DEVELOPMENT ACTIONS (4)

Technologies	TDA Name	TDA Details	TDA Expertise Details	Chances of Success	Implementation Time	Cost
▲						
LLM-driven check of requirements propagation	LLM-driven check of requirements propagation			40 to 80%	6 months to 2 years	>100k
Machine learning-accelerated pedestal MHD stability evaluations	Fast surrogates for fusion scenario simulations	To replace high-fidelity simulations by machine learning models: pedestal simulations, MHD, turbulence, exhaust, neutrals, sources, etc.	Various labs working already in individual models, some coordination might be sought but they might have different settings e.g. DIFFER, VTT, UKAEA, CALDERA AI	>80%	>2 years	>100k
Material acceptability prediction for a dedicated application	Usage of AI tools to enhance the introduction of new materials on (nuclear) design codes.	Use of AI tools to support the introduction of new materials into (nuclear) design codes. Requires deep knowledge of material qualification standards and AI expertise, with successful completion of the “Generation and curation of materials database” as a prerequisite.	DAES SA, Tecnalia, HI Iberia, SCK CEN, Barcelona Supercomputing Center	40 to 80%	>2 years	>100k
ML-driven surrogate modelling	Accessible database to store and share data for simulation and experiments	Applied to fusion Standarization is key Fast queries Easily accessible, not necessarily public, which would foster its use	Some experts are available in Europe, but since the required expertise is broad and not highly specialized, it's not a limiting factor for this TDA	>80%	6 months to 2 years	>100k
ML-driven surrogate modelling	Use of surrogate modelling in environments where we have data/expertise	Surrogate modelling is already well established in aeronautics, and a similar approach could be applied in fusion. Using this methodology would also help promote its adoption by demonstrating parallels with other sectors. (Estimated TRL 3–4 after the TDA, potentially reaching TRL 6 later)	Specialized labs on this topic	>80%	6 months to 2 years	>100k
Optimization of component positioning during assembly	Assembly procedure optimizer	This TDA is also related to "Identification of most optimal assembly procedure" AI model capable of providing a visual assembly sequence of the assembly of (complex) systems, being able to simulate the different steps taking into account (constraints parameters, dimensional compliance, tolerance, schedule, cost, manufacturing techniques, health and safety constraints, non conformities...) - Alert operators of high risk steps - Compliance with assembly tolerance - Adherence with 3D model - Feed the model with design data, and then feed with as-built dimensions - Feed with the time to due assembly, number of resources, - Feed with assembly procedure, tools	Input data to be gathered. CT Engineering Group, or other engineering companies in consortium with AI developers	40 to 80%	>2 years	>100k
Prediction of behaviour and mechanical properties of new materials	AI tool to predict physical and mechanical properties based on partially available data	Some material tests (e.g., cyclic loading, time-dependent properties, irradiation ageing) are lengthy and resource-intensive. This AI tool aims to predict such properties using limited available data. Example TDA sub-tasks include: - Unguided classification of existing materials - Structuring and organizing available data	DAES SA, Tecnalia, HI Iberia (IFMIF-DONES), Barcelona Supercomputing Center, SCK CEN	40 to 80%	6 months to 2 years	>100k
Prediction of metrology data	Build an AI model to predict dimensional metrology data	Approach: - Consolidate all available data (both accessible and hidden) - Split datasets into training and validation sets - Begin with a defined component/manufacturing process, then extend to cassette bodies (evaluate hidden parts when closed), Vacuum Vessel, upper launcher (mirror positioning after final assembly) Implementation timeline depends on data availability.	Geatop, CT Engineering Group, Hexagon	>80%	6 months to 2 years	50k to 100k

TECHNOLOGY DEVELOPMENT ACTIONS (5)

Technologies	TDA Name	TDA Details	TDA Expertise Details	Chances of Success	Implementation Time	Cost
▲						
Predictive maintenance of components	Develop proof of concept on fusion specific system	Building on experience from industrial systems	SAVANTIC, U of Ghent, NBTF	40 to 80%	6 months to 2 years	>100k
Predictive maintenance of components	List potential areas of application and associated development needs	To prioritize fusion specific items	All fusion device operators	>80%	<6 months	<50k
Predictive maintenance of components	Validate proof of concept for predictive maintenance		Machine operators	40 to 80%	6 months to 2 years	>100k
Predictive modelling of plasma instabilities and disruptions	AI Enabled Scientific Discovery	Leverage on AI algorithms to allow for scientific automated discovery				
Predictive modelling of plasma instabilities and disruptions	Event Prediction, Detection and Classification	Tool that is able to catch more than disruptive events like for instance anomalies, scenario changes, transition confinement events, and any other interesting events for optimization (to get alerts) (Also connection to auto-labelling TDA, same family models. Outcome of auto-labeler TDA will be relevant to this TDA and viceversa). Some of the applications might require real time.	EUROfusion groups (several) working on prediction, classification, etc. Savantic	>80%	6 months to 2 years	>100k
Real-time optimization of manufacturing parameters	Process Optimization for Distortion Prediction for Additive Manufacturing	Predict influence of manufacturing process on geometrical deviations. - Pre-compensate via AI supported predictions. - Compensate in-situ. - Real time Metrology feedback to AI.	- Universities and labs develop new, publicly accessible process algorithms - Companies develop processes tied to IP, availability varies case by case - Existing processes are material-specific; not all materials covered - Requires selection of both material and end-user need, as well as machine/tool - Rapidly growing global know-how - Each project generates lessons learned applicable across manufacturing methods	>80%	6 months to 2 years	>100k
Real-time optimization of manufacturing parameters	Real-Time Manufacturing optimization	Real-time monitoring, correction and "tweaking" of parameters for best outcome		>80%	6 months to 2 years	>100k
Recognition of abnormal events in the operation or state of a system	Validate existing algorithms on operating machines	Get expert feedback during operation to gain and label training data	ITER, WEST, CIEMAT	40 to 80%	>2 years	>100k
Welding success rate prediction	Welding optimization and in-situ compensation	Brainstorming ideas: - In-situ automated compensation of welding parameters - Challenge: parameters are validated/specified, not the process itself → limits AI's ability to adjust inputs - Potential for large-scale, repetitive, and reproducible welding processes - First step: pilot project, though implementation may be challenging - In-situ measurement tools for defects (NDT, metrology) and distortions → explore predictive models to train AI - Capture lessons learned after each process to improve the next	Top Level welding institute for example in collaboration with industry	40 to 80%	6 months to 2 years	>100k
Welding success rate prediction	Welding Process Optimization	Select a welding process with reproducible production as starting point, not complex state of the art welding				