Report on Innovation Projects

MS10 - August 2025

Introduction

Spoke 2 has funded 11 Innovation Projects during the three calls opened by the HUB (Jun 2023, December 2023 and February 2024).
Among these, listed in Table 1, only 4 have Spoke 2 as the primary spoke (see those in bold in the same table), and their current status will be described in this report. For more details, one should refer to the full reports delivered as a part of the standard milestone reporting.

Name	Funding Round	Primary Spoke	Industrial part(s)	Academic partners	Secondary Spokes
ML-BLEN	**1st**	**Spoke 2**	**Eni**	**INFN, UNITS**	
PMIP-HAIP	**1st**	**Spoke 2**	**Eni**	**INFN, INAF, UNIBO**	
GDC	1st	Spoke 0	Eni	CINECA, UNIMIB	Spoke 2
AGRI@ INTESA	**1st**	**Spoke 2**	**Intesa Sanpaolo, TASI**	**UNICT, UNISALENTO, UNIBA, UNINA, UNIMIB**	**Spoke 5**
FD	1st	Spoke 10	Intesa	UNIPD, INAF, UNIFI	Spoke 2, Spoke 3
Leonardo IDL	1st	Spoke 3	Leonardo, TASI	INFN, INAF, UNISAPIENZA, POLIMI, UNITOR	Spoke 2

FLF	1st	Spoke 1	Intesa Sanpaolo	UNISALENTO, UNITO	Spoke 2
Hammon	1st	Spoke 3	Leitha', SOGEI	INAF, CMCC, INFN, ENEA, POLIBA, UNISAPIENZA, UNISAL, UNITO, UNITN	Spoke 0, Spoke 1, Spoke 2, Spoke 4, Spoke 5 ...
SAFE	**2nd**	**Spoke2**	**IFAB**	**UNIBO**	
SowStain	2nd	Spoke 9	IFAB	CNR, UNITN, POLIBA, UNINA, UNICAL, UniCal, ENEA, UNIVAQ, UNIBA	Spoke 2, Spoke 4, Spoke 5
ELIO-ENI	3rd	Spoke 7	Eni	CNR, UNITO, UNICAL, UNINA, UNITO, UNITS	Spoke 2
Hammon-eq	3rd	Spoke 3	Unipol	INAF, UNIMIB	Spoke 2

Table 1: projects with Spoke 2 involvement. The lines in **bold** identify the projects where Spoke 2 is the primary Spoke.

In this document we explicitly report only on the projects where Spoke 2 is the primary Spoke; the other projects will be reported by the actual primary Spoke and will include Spoke 2 in the description.

The report focuses on the intermediate report of MS10 as of August 2025. A synthetic description of the project is reported with the advancement and, in some cases, the conclusion of the planned activities.
For further information, including plots and papers, please refer to the full reports by the industrial PIs.

Short description of the projects

ML-BLEN (ENI, INFN and UniTS)

The project *"Blending Machine Learning with advanced numerical simulations: application to the sustainable exploitation of natural resources"* aims to overcome the issues met by Eni's reservoir and production engineers in dealing with complex, multi-scale simulations of the subsurface. This project proposes to explore, implement and apply a bunch of novel techniques known as Physics Informed Machine Learning (PIML) algorithms. These methods represent a novel methodology aiming at building neural networks that are trained to solve supervised learning tasks while respecting any given law of physics described by general nonlinear partial differential equations.

PMIP - HAIP (ENI, INFN, INAF and UniBO)

The project *"Harnessing the power of Artificial Intelligence for predictive maintenance of industrial plants"* aims to demonstrate the usability of modern AI-based techniques in the realization of systems for predictive maintenance and for the modelling of the interdependencies between systems in complex industrial apparatuses. Techniques in use in the academic domain (for example those on the detection of anomalous signals, those parsing log entries from IT systems, those dealing with graphs) will be tested on appropriate data samples coming from Eni productive sites; these data can range from sensors readings in industrial facilities, single apparatus data, complex relationships graphs between multiple
apparatuses.

AGRI@INTESA (Intesa San Paolo and TASI, UNICT, UNISALENTO, UNIBA, UNINA, UNIMIB)

The project Agri@Intesa aims to elaborate a statistical model that leverages satellite and financial data in the context of the agriculture sector to study the possible correlation between observed satellite imageries and financial data of sector's enterprises, with the final aim to further increase the accuracy of bank's credit scoring systems and counterpart assessment in agriculture.
The project is divided into three main workpackages. WP1 focuses on collecting, curating, and structuring satellite image data, ensuring interoperability and alignment with financial datasets and geographical perimeters. WP2 oversees the identification of agricultural

enterprises and the collection of their financial data, defining the scope and type of information to be integrated. WP3 carries out the core scientific work, applying, refining, and validating AI-based algorithms on the combined image and financial data to develop predictive models.

SAFE (IFAB and UniBO)

The SAFE (Secure anomaly detection edge AI system for critical Environments) project consists of the development of an Edge AI system for anomaly prediction in critical industrial processes, with a focus on monitoring stationary rotating components. The system employs sensors and Machine Learning algorithms for the identification of malfunctions in real time. The main steps include problem analysis, data management, and the development of Edge AI models for anomaly detection, followed by extensive te sting and validation.

Scientific validation from the latest available report (from November 2024 to August 2025)

ML-BLEN

The objective in MS10 (specifically for WP3) was to test the prototypes developed in WP2 at a realistic scale, using either synthetic datasets or real data from relevant domains, with possible execution on the ICSC infrastructure.

Activities from M9 continued into M10, with WP3 tests progressing as planned. Task T3.1 is ongoing, and initial results have been documented in reports. In MS11 realistic test cases will be performed under ENI's advice. The possibility of a phase of technology transfer from UniTS and INFN to ENI is envisaged.

The BLEN project is proceeding as expected, with results and papers being published.

The report with the details on the results obtained in the second period of MS10, highlights a transition from traditional and neural-prior methods to an automatic differentiation (AD) log-domain deconvolution pipeline for seismic image filtering and reconstruction, finding that AD offers superior stability and phase consistency on multiple datasets. Classical approaches like Richardson-Lucy, Markov Random Fields, and Expectation-Maximization were investigated but lacked flexibility compared to modern generative methods. Diffusion models, especially physics-informed variants, are seen as promising for future developments, though work is at a preliminary stage. The AD approach, implemented in PyTorch, provided reproducible and reliable results without heavy regularization. It enhanced reflector clarity, continuity, and polarity assessment in practical seismic cases. Future research will focus on understanding the AD method's robustness and refining regularization techniques for further improvement.

PMIP - HAIP

Overall, all work packages advanced as planned, with validated methods and demonstrators providing strong evidence of improved predictive maintenance capabilities.
For Milestone 11, the focus is on completing result interpretation with subject matter experts (WP1) and transferring all codes and results to partners (WP2–WP4). These activities will ensure full knowledge transfer and proper documentation in the final KPI reports.

During this reporting period, WP1 (ENI) integrated new datasets and continued benchmarking with internal predictive tools, providing the report expected as KPI.
WP2 (UniBO) tested multiple neural network architectures, including Autoencoders and transfer learning, producing comparative results to improve anomaly detection. Report as KPI.
WP3 (INAF) finalized model validation and delivered a demonstrator showing that segmentation combined with a heterogeneous ensemble (Random Forest + XGBoost) increased AUC-ROC from 0.86 to 0.98, demonstrating substantial performance gains. Report as KPI.
WP4 (INFN) completed proof-of-concept tests on data streaming platforms and compared partner approaches, confirming the feasibility of deployment. Overall, all work packages advanced as planned, with validated methods and demonstrators providing strong evidence of improved predictive maintenance capabilities.

Papers on two use cases accompany the research reports MS10. These two use cases are:
* "Benchmarking Existing Solutions for Anomaly Detection in Time Series" which evaluates deep learning approaches for anomaly detection in industrial time series, focusing on predictive maintenance of turbomachinery. INFN proposed a novel approach, combining Generative Adversarial Networks (GANs), a reconstruction-based framework, and Graph Neural Networks (GNNs) for effective and interpretable anomaly detection. A GNN-based model and an RNN-based model, both using adversarial autoencoders, were benchmarked

on real sensor data from ENI's COVA plant. Results show the RNN is less computationally demanding and achieves better detection performance than the GNN and a PCA-based baseline, and demonstrates the effectiveness of our approach in terms of robustness and interpretability.
* "Optimization of a Natural Gas Liquefaction Unit (LNG) with Deep Reinforcement Learning" where Deep Reinforcement Learning (RL) is used to optimize natural gas liquefaction plants by dynamically adjusting process parameters using IoT sensor data. LNG units are energy-intensive systems requiring precise control of refrigeration cycles, compressor operations, and mixed refrigerant compositions to maintain efficiency under dynamic lean gas feed rates. Traditional optimization methods, such as nonlinear programming (NLP) or heuristic algorithms, often struggle to adapt to real-time operational uncertainties. RL agents improve efficiency under supervision, bridging the gap between static design and real-time operation. While fully autonomous control is not yet safe, this approach shows strong potential for adaptive, supervised plant optimization.

AGRI@INTESA

The previous issues on providing data from agriculture companies, due to privacy concerns, have been addressed. During this period, the team formalized a process for data exchange with a target company, addressing confidentiality issues through an NDA validated by ISP's Legal office and signed by one of the target companies.

A literature review was conducted to investigate correlations between agronomic and financial indices—particularly in relation to predicting default risks in sectors such as viticulture—highlighting existing methodologies, findings, and research gaps. A set of papers were prepared, focusing on several types of cultivation and on different kinds of satellite data. Additionally, the team analyzed financial statements from 2022–2024 to derive key variables and ratios (e.g., revenues, EBITDA, profit, leverage, working capital, and liquidity indicators) to assess the financial standing of the company.

The MS10 report is not explicit on the fulfillment of the additional tasks planned for WP1 (TAR1.4, the delivery of the final dataset ensuring that it meets the project quality standards and objective, TAR1.5 leveraging current and future tools) and for WP2 (TAR2.7 refining databases with updated financial statements, and TAR2.8 calculation of financial variables for selected companies).

A few TAR activities (1.5, 3.5, 3.6, 3.7), are still under refinement, evaluation, or development, and presumably will be shifted to milestone 11 (at the moment of writing the MS11 research plan is not yet available).
TAR3.5 involves developing a pipeline for estate-level analysis, including parcel extraction from AGEA, manual segmentation, and downloading satellite images; according to the report, this activity is ongoing, with efforts directed at refining and automating steps. TAR3.6 aims to build algorithms using optical satellite data (e.g., LGB, Sentinel) to accurately identify areas of interest, though challenges remain due to mismatches between maps and reality and incomplete correlations with KPIs; the report notes this work is still in progress and under evaluation. TAR3.7 concerns transforming AGEA data into structured formats to enable parcel-level scaling and automation; its status is also reported as ongoing, particularly regarding tasks like text data analysis.

It is not stated if they will be completed in MS11, whether, among the objectives for M11, is to focus on refining the definition of the sample of target companies.

SAFE

During the first reporting period, WP3 (Generative AI) faced delays due to the lack of a dedicated resource at the University of Bologna. IFAB initiated preliminary studies, conducting a literature review and evaluating synthetic data generation approaches (GANs, augmentation, interpolation). Original targets were not fulfilled and the team shifted efforts to improving real-data modelling. These efforts laid the foundation for a structured strategy, though no formal design or model testing was completed.

WP4 (Development and Training of Edge AI Model) achieved substantial progress in developing an anomaly detection model for edge deployment. A robust preprocessing and exploratory analysis pipeline was established to aggregate, clean, and structure sensor data into reliable time series, and statistical visualisations confirming data quality. Building on this foundation, extensive benchmarking was conducted across CNN, LSTM, CNN–LSTM hybrids, TCN, Time-BERT, DeepAnT, DeepAR, and gradient boosting, with CNN–LSTM hybrids providing the best trade-off between accuracy and drift sensitivity, while CNNs alone excelled on stable vibration features. Optimisation techniques—including quantisation, pruning, and kernel simplification—preserved accuracy while reducing computational load, ensuring feasibility on STM32-class microcontrollers. The optimised models, trained on real datasets, achieved high accuracy ($R^2 > 0.97$ for stable features), and the embedded workflow was successfully validated in STM32CubeIDE.

Building on these advances, WP5 (Approximate computing) initiated embedded workflow preparation, exporting quantised models to C arrays and confirming deployment feasibility on STM32U545RE and SECO Easy Edge hardware.

Finally, WP6 (System testing) reached the TAR6.1, defining criteria and datasets for upcoming simulation tests (anomaly injection, drift simulation). However, execution of the tests was explicitly postponed to the next milestone, once embedded prototypes are consolidated.

All targets and activities in the WPs are extensively described with their corresponding KPI in the accompanying report document “SAFE – Methods, Experiments, and Edge Deployment Findings, September 2025”.