Ref. Ares(2026)8764424 - 15/09/2026

Artificial Intelligence in Manufacturing leading to Sustainability and Industry5.0

Project Acronym:

AIMS5.0

Grant agreement no: 101112089

Deliverable no. and title	**D5.3 – Results on AIMS5.0 Tier 1 and OEM Use Cases**	
Work package	WP 5	Use cases for AI in Production for OEM and Tier1 stakeholders
Task	T5.1, T5.2, T5.3	Methods, Processing and Tools AI-based Logistics Systems AI techniques applied to Industrial Domains
Subtasks involved		
Lead contractor	Infineon Technologies AG Thomas Gutt, mailto:thomas.gutt@infineon.com	
Deliverable responsible	Vahid Salehi (AIDI)	
Version number	V7.0	
Date	15/09/2026	
Status	Final	
Dissemination level	Public (PUB)	

Copyright: AIMS5.0 Project Consortium, 2026

Authors

Part. no.	Part. short name	Author name	Chapter(s)
2, 4	AIDI, CETTO	Vahid Salehi, Maximilian Cetto	2.1
40	PCL	Jan Wytze van der Weit	2.2
47	HQ	Adrien Mathieu	2.3
38	EDI	Kārlis Freivalds	2.6
43	SIGNIFY	Marc de Samber	2.7
45	DAC	Krzysztof Wołk	2.7
51	SU	Oguz Icoglu	2.8
28	UPM	Jaime Sancho, Manuel Villa, Eduardo Juárez	2.4
25	IDEKO	Itziar Ricondo, Mikel Villos	2.5
3	BMW	Daniel Jess	2.4
28	UPM	Eduardo Juárez	Editor

Document History

Version	Date	Author name	Reason
v0.0	16.07.2026	Thomas Irrenhauser, Eduardo Juárez	Initial Template
v1.0	03.08.2026	Thomas Irrenhauser, Eduardo Juárez	First consolidated version
v2.0	10.08.2026	Thomas Irrenhauser, Eduardo Juárez	Second consolidated version
v3.0	18.08.2026	Thomas Irrenhauser, Eduardo Juárez	Third consolidated version
v4.0	25.08.2026	Thomas Irrenhauser, Eduardo Juárez	Fourth consolidated version
v5.0	01.09.2026	Thomas Irrenhauser, Eduardo Juárez	Fifth consolidated version
v6.0	08.09.2026	Thomas Irrenhauser, Eduardo Juárez	Sixth consolidated version
v7.0	15.09.2026	Thomas Irrenhauser, Eduardo Juárez	Seventh consolidated version

Publishable Executive Summary

Purpose of this deliverable

This document constitutes Deliverable D5.3 of the AIMS5.0 project, entitled "Results on AIMS5.0 Tier 1 and OEM use cases". It provides a comprehensive account of the publicly available results produced across all eight use cases of WP5 during the third and final reporting period of the project, together with an evaluation of the innovation outcomes achieved against the objectives and key performance indicators defined at the outset.

D5.3 serves two complementary purposes. It acts as a public record of the technical outputs made available by the consortium — including published papers, open-source software, datasets, and validated demonstrators — and as an evaluation instrument that assesses the extent to which each use case has met its defined objectives and contributed to the broader goals of WP5: the application of AI as an integrated, cross-disciplinary approach in real industrial environments.

Summary of results by use case

UC1 delivered a CO_2 fleet scorecard and energy management tool for injection moulding, enabling real-time tracking of machine energy consumption and CO_2 emissions in a production environment at Cetto. The tool provides manufacturers with a practical instrument for sustainable production planning and energy-source optimisation.

UC2 advanced the automation of robot testing for Philips IPL devices through human motion modelling, and developed process mining tools that enabled measurable improvements in production quality monitoring and fall-of-rate reduction at Philips Drachten. A peer-reviewed publication on significant discriminative sequential pattern mining was produced and is publicly available.

UC3 achieved the most operationally mature result reported in this deliverable. An edge-AI service architecture was deployed in Husqvarna's commercial Smart Lawn Service in Germany across more than 100 customer contracts, demonstrating more than 80% reduction in administrative effort and 92% reduction in AI model development time. The platform was validated against sustainability, productivity, resilience, and human-centricity dimensions and scaling-tested for fleets exceeding 10,000 connected devices. Multiple public results are available including a peer-reviewed conference paper and open-source Arrowhead 5.2 software on GitHub.

UC4 developed an autonomous drone-based scanning ecosystem for the digitalisation of production and logistics environments, combining 3D point cloud acquisition, automated AMR navigation map generation, and data quality assessment tools, validated in real BMW industrial environments with demonstrated reductions in planning effort and AMR map creation time. In parallel, a digital twin of a multi-view 3D reconstruction system was validated against a BMW wind tunnel application, where the MVS pipeline achieved sub-millimeter reconstruction accuracy for 99.26% of points at a density of 3.82 million points over a 56 m^3 bounding box. Future activities will focus on industrial deployment and gradual rollout across the global BMW production network.

UC5 delivered an industrially validated AI monitoring system for tool condition and tool breakage detection in large-scale precision machining at GOIMEK, integrating high-frequency signal processing, machine learning, edge computing, and the Arrowhead industrial framework. Tool breakage detection was demonstrated in 1.5 seconds from data acquisition to inference, with classification accuracy reaching 90% for anomaly detection in representative industrial conditions. Three peer-reviewed publications are publicly available.

UC6 reached TRL5 with a complete end-to-end AI-based robotic system for automated conveyor feeding, combining object detection, 6DoF pose estimation, differentiable-physics feedback control, and LLM-based robot task script synthesis. Key validated KPIs include 98.7% socket detection accuracy, sub-2mm positioning precision, and a 4% unplaced object rate. The system is progressing towards TRL6 and industrial introduction.

UC7 validated computer vision tools for assembly line monitoring at a Signify factory in Spain, demonstrating the capability of AI models to track takt time, identify anomalies, and support line balancing for selected product types. In parallel, a first-of-a-kind low-cost Direct Energy Deposition metal 3D printer was developed and validated with customer prototype samples, with a solid business case confirmed for relevant luminaire product categories. Synthetic defect generation and detection experiments for quality inspection, including integration with Eclipse Arrowhead and a Carbon Footprint component, were also completed and evaluated.

UC8 reported a fully integrated Industry 5.0 demonstrator combining AI-based autonomous mobile robot perception and task management with an EEG-controlled upper-body exoskeleton for workers with physical disabilities. All defined KPIs were met or exceeded: at least 25% reduction in task execution time, 30% reduction in traffic conflicts, 90% task assignment accuracy, and ergonomics improvement of more than 5%, demonstrating the potential of AI-assisted technologies to create more inclusive and sustainable manufacturing environments.

Overall evaluation

The results reported in D5.3 confirm that the AIMS5.0 project has successfully demonstrated AI integration in real industrial environments across a diverse range of manufacturing and logistics domains. All eight use cases reached operational or near-operational validation during the final reporting period, with several already in active exploitation or advanced preparation for commercialisation. The collective evidence presented in this deliverable demonstrates that AI-enabled manufacturing can simultaneously deliver improvements in productivity, sustainability, resilience, and human-centricity, and that the transition from laboratory to industrial deployment — while demanding — is achievable within the scope of a collaborative European research project.

Table of contents

1. Introduction 9
1.1 Objective and scope of the document 9
2. All Public Results Available in the Repository and Evaluations of the Innovation Project 11
2.1 UC1: Development of Green AI-Lifecycle Management Approaches (AIDI, Cetto) 11
2.1.1 Objectives of the Use Case 11
2.1.2 Development of the Demonstrator 11
2.2 UC2: AI Enhanced Value Chain (PCL, RUG, TU/e, SEMAKU) 12
2.2.1 Objectives and use-case requirements 12
2.2.2 The abstract of this paper 12
2.2.3 Evaluation of the Use Case 13
2.3 UC3: Servitization and service delivery enabled by Zero administration (HQ, SEB, SA, LTU) 14
2.3.1 Demonstrated solution 15
2.3.2 KPI achievement 16
2.3.3 Validation against the evaluation dimensions 17
2.3.4 Final improvements and lessons learned 17
2.3.5 UC3 — Public results available in the repository 18
2.4 UC4: UC4 AI based improved connection between production and logistics (BMW, IPH, UPM) 18
2.4.1 BMW Contributions 19
2.4.1.1 KPI Validation Results 20
2.4.1.2 Conclusion 20
2.4.2 UPM contribution to UC4 20
2.4.2.1 Camera hardware evaluation and technology selection 20
2.4.2.2 Multi-ToF camera rig and test sequence generation 21
2.4.2.3 Multimodal camera rig and plenoptic sequence capture 22
2.4.2.4 3D reconstruction algorithm evaluation and digital twin development 23
2.4.2.5 Reconstruction results and validation 24
2.5 UC5 AI techniques at different layers in Machine Tool Domain (IDEKO, SORALUCE, GOIMEK, SAVVY, MULTIVERSE) 25
2.5.1 Achieved results: 26
2.5.2 Evaluation of the Innovation Project 29
2.6 UC6 Reconfigurable AI-based automated conveyer feeding by a robotic system (EDI, HUA). 30
2.6.1 Object detection and segmentation (EDI) 30
2.6.2 Grasp affordance estimation (EDI) 31
2.6.3 Bottle 6DoF pose estimation (BME, EDI) 31
2.6.4 Grasp planning and execution (EDI) 32
2.6.5 Socket detection and tracking (EDI) 32
2.6.6 Differentiable physics method — robot arm feedback control (EDI, HUA) 33
2.6.7 AI-based robot control script synthesis (EDI, HUA) 34
2.6.8 Industrially relevant demonstrator 34
2.6.9 EDI's role and contribution in UC6 35
2.6.10 HUA's role and contribution in UC6 35
2.7 UC7 Digital Automated Luminaire Manufacturing platform (Signify, TU/e, possibly RUG) ... 35
2.7.1 On the use of Computer Vision and AI for monitoring assembly lines 36
2.7.2 3D metal printing using DED technology 37
2.7.2.1 Further developments on the low-cost DED printer machine 37

2.7.2.2 Results of prototype printing 38
2.7.2.3 Business case 39
2.7.2.4 On KPIs: 39
2.7.3 DAC contribution: synthetic defect generation, detection experiments and adaptable AI workflows 40
2.7.3.1 Controlled benchmark on the MVTec metal_nut dataset 40
2.7.3.2 Production-video experiment 41
2.7.3.3 API-backed image-generation verification 41
2.7.3.4 Exploratory approximate class-removal experiment 42
2.7.3.5 Open Access Platform integration 42
2.7.3.6 Consolidated evaluation against UC7 objectives 42
2.7.3.7 Public-release and repository status 43
2.7.3.8 Final DAC statement for D5.3 43
2.8 UC8: AI supporting and strengthening human and manufacturing cycles (SU, TUBITAK) 43
3. Conclusions 47

List of figures

Figure 1: The as-realised process model of the OneBlade manufacturing process in which differences with the as-designed model are highlighted in red. 13
Figure 2: Screen shot of the developed dashboard 13
Figure 3: A Philips IPL device 14
Figure 4. Measured accuracy for the cameras under test. 21
Figure 5. Multi-ToF camera rig 21
Figure 6. Multimodal camera rig 22
Figure 7. 3D (xyz) actuator 22
Figure 8. Wind tunnel simulation. The cameras are positioned 5 meters apart from the car. 23
Figure 9. On top left, the RGB image of the wheel, on top right, the depth map generated using a depth estimation algorithm, on bottom, the 3D reconstruction of the wheel using the depth map generated (in black and gray) over the ground truth 3D data (in red). 24
Figure 10. Generated point cloud of the whole car 25
Figure 11. Detailed point cloud of the car wheels 25
Figure 12. Multi-layered infrastructure for data processing and AI deployment 27
Figure 13: Arrowhead integration in SAVVY multi-layered solution 28
Figure 14: FUMO machine 37
Figure 15: More complex centro-symmetric shapes (including overhangs) 38
Figure 16: Material finishing (powder coating) and odd-shaped geometries (square-like) 38
Figure 17: Customer samples 39
Figure 18: UC8 Human detection and tracking 44
Figure 19: UC8 Gesture recognition 44
Figure 20: UC8 EEG-based intelligent exoskeleton control 45

List of tables

Table 1: UC3 public results 18
Table 2: Evaluation table 29
Table 3: UC7 controlled benchmark on the MVTec metal nut dataset 40
Table 4: UC7 production-video experiment results 41
Table 5: UC7 exploratory class-removal experiment 42
Table 6: Evaluation against UC7 objectives 42
Table 7: Public-release and repository status 43
Table 8: UC8 achieved KPIs and results 45

1. Introduction

WP5 contributes to the application of AI as an integrated approach for different disciplines and domains in the production. Despite the numerous relevant AI-technological impacts and advances, AI is not fully integrated in a real industrial context and application. The reason therefore is that currently the different domains and disciplines have the challenge to understand the full processes, methods and systems which can be applied in different areas in a production or logistic environment. Furthermore, today AI is more considered as a “Black Box” in the industry and needs more understanding about how the whole process chain could look like in the production environment. Furthermore, the different approaches and domains for AI like Machine Learning, Computer Vision, Robotics, NLP have to be attacked as an integrated approach inside the production or e.g., logistics. Therefore, there are new approaches, methods, systems and tools which must be developed to integrate AI as a fundamental part in an industrial context. For the above-mentioned reasons the WP5 is necessary to make this application more visible.

The project AIMS5.0 is generating a huge contribution for that. The development of artificial intelligence (AI) research is at an all-time high and is expected to continue to grow. This development needs to be supported by enabling technologies, such as improving access to high-quality cyberinfrastructure and data. These state-of-the-art methods of artificial intelligence can be transferred to real industrial applications by focusing on strengthening research and innovation from the laboratory to the field. It is also necessary to provide a framework for experimenting and testing these solutions in real industrial environments. To improve system control in a complex environment, systems must be able to recognize different scenarios and then derive all necessary information from them in an animation or simulation. Deep learning, for example, can enable AI-based systems to recognize defined objects and obstacles or undefined pending problems as anomalies and thus help to identify deviations from the control process. However, given the large amount of data, there is currently a need for methods to extract the most important information for a particular system. In addition, systems must be able to examine the increasingly complex process sequence for deviations and anomalies, even for long-term observation.

1.1 Objective and scope of the document

This document constitutes Deliverable D5.3 of the AIMS5.0 project (HORIZON-KDT-JU-2022-01-IA), entitled "Results on AIMS5.0 Tier 1 and OEM use cases". It covers the results produced across all eight use cases of WP5 during the third and final reporting period, providing a comprehensive account of the technical achievements, demonstrators, and key performance indicators reached by the project consortium.

The objective of this document is twofold. First, it presents an overview of all publicly available results produced within each use case, including published papers, open-source software, datasets, and demonstrators deposited in project and public repositories. Second, it provides an evaluation of the innovation project by assessing the extent to which each use case has met its defined objectives and KPIs, and by reflecting on the contribution of the results to the broader goals of AI integration in European manufacturing and logistics.

The document covers the following use cases: UC1 (Green AI-Lifecycle Management), UC2 (AI Enhanced Value Chain), UC3 (Servitization and Zero Administration), UC4 (AI-based connection

between production and logistics), UC5 (AI techniques in Machine Tool domain), UC6 (Reconfigurable AI-based automated conveyor feeding), UC7 (Digital Automated Luminaire Manufacturing), and UC8 (AI supporting human and manufacturing cycles). Each use case is described in terms of its objectives, technical developments, demonstrator results, KPI achievement, and public availability of outputs.

2. All Public Results Available in the Repository and Evaluations of the Innovation Project

2.1 UC1: Development of Green AI-Lifecycle Management Approaches (AIDI, Cetto)

Running machine in Production based on green and sustainable energy is one of the most important points in the future. Based on vison of Zero CO2-Production environment is one of the base points. In this use case we will define a green Munich Agile concept to realize the Idea of green Energy an Environment. Provision of AI-enhanced algorithms for faster time to market, better quality and improved performance, by: aggregating requirements from the use cases and developing and refining algorithms to fit the needs of European Industry or supporting the integration of algorithms into specific use-case scenarios, as well as with presenting use-case specific results as gains reached by the applied algorithms. The algorithm should define the best path for Energy sourcing approach. Research & evaluate a concept for manufacturing execution systems (MES) based on AI improving production cycle for a Green and more sustainable background.

2.1.1 Objectives of the Use Case

The following chapter explains the results of Deliverable D5.3. The focus was on demonstrating the technical objectives actually achieved using a prototype. The use case pursued four key objectives. These included, in particular, identifying sustainability aspects related to the use of plastic injection molding machines and analysing productivity. In addition, a robot was deployed at the Cetto plant to handle material transport.

2.1.2 Development of the Demonstrator

The developed demonstrator was used to record and display the energy and CO_2 status of the production environment under consideration. Like the analysis of CO_2 footprints in vehicle manufacturing, a so-called CO_2 fleet scorecard was developed for the injection Molding machines in use. The goal was to enable developers, with the help of this CO_2 fleet scorecard and a software tool, to determine the actual energy consumption of the machines used to manufacture the injection-molded parts. Depending on the available infrastructure, the machine data was collected using an energy meter. The data was then transferred to a database for later processing using a software tool. In a second step, the measured values of machine consumption were synchronized over time and processed within the production environment. This made it possible to determine and utilize the energy consumption of the individual machines.

At Cetto, a corresponding measurement system was set up as an example. This system was used to systematically record the energy consumption of the individual injection Molding machines and to make the data available for further processing and the calculation of average consumption. In a third step, the measurement data was transferred into an energy-based, CO_2-based labelling system. The determined values were processed using a key performance indicator (KPI) system so that the corresponding CO_2 emissions could be calculated based on the resulting energy consumption. The key performance indicators focused on the energy consumption of a specific production order or on the energy consumption of the machines currently in operation. The energy consumption of an order was determined by entering the recorded consumption values into a software-based tool. This allowed the various machines to be recorded in the system. The machines had different characteristics. Depending on the machine type, manufacturer, and machine runtime, the corresponding consumption values could be assigned to them.

This made it possible to accurately allocate the respective machine consumption figures to the individual injection molding machines.

Determination of CO_2 Fleet Values: Based on this, the average CO_2 value for the entire injection Molding machine fleet was determined, and a corresponding CO_2 fleet value was calculated. This made it possible to illustrate the machines' current energy consumption in real time and to determine and establish corresponding energy and CO_2 fleet values for the entire injection Molding machine fleet. Based on the recorded energy consumption and CO_2 values, an average CO_2 emissions value was then determined. In addition, various energy pathways were examined. Of particular interest was the question of what effects would result if the machines were not operated exclusively with conventional electricity based on the German electricity mix, but if part of the energy demand were replaced by sustainable energy sources. The focus was particularly on solar energy. At Cetto, it is possible to generate a portion of the energy required to power the machines from solar installations. These installations result in a lower CO_2 footprint than powering the injection molding machines exclusively with electricity based on the conventional German electricity mix.

2.2 UC2: AI Enhanced Value Chain (PCL, RUG, TU/e, SEMAKU)

2.2.1 Objectives and use-case requirements

The goal of Philips Drachten for AIMS5.0 is to enable a step towards the overall vision of mass customization by creating a value chain leading from the consumer back to the consumer. To achieve this goal, consumer data must be fed back through several connected data and modelling platforms to be able to offer the consumer a customized and improved product upon renewal. The key result of this AIMS5.0 use-case of AIMS5.0 is a proof of concept in a real-life value chain (TRL7) of an architecture for an AI enhanced Value chain that enables mass customization for added consumer value, a graphical representation is provided.

Use case 2 consists of several parts. One part is the investigation into the OneBlade production process. A collaboration of RUG and PCL. From the RUG Michel Medema and Mare Smits worked together with Erik Bloemsma and Jan Wytze van der Weit from PCL. A publication was written and published on December 2, 2025. The title is 'Mining Significant Discriminative Sequential Patterns from Event Logs.

2.2.2 The abstract of this paper

Business processes are at the heart of any organisation. Advances in digital technology are changing the way organisations manage and execute their processes. These developments offer opportunities to improve performance and productivity, but also increase technical complexity and raise concerns about social, environmental, and ethical issues. Process mining can support organisations in ensuring that their processes run efficiently and comply with regulations by leveraging the large amounts of event data recorded by the information systems that typically support process execution. Comparative process mining identifies patterns that distinguish cases in one event log from those in another, revealing potential root causes of inefficiencies and undesired outcomes. Unfortunately, existing tools and techniques primarily focus on the control flow of a process, disregarding other contextual information, or require manual construction of features, which forces the complex sequence of events to be reduced to a small selection of values. This paper presents an algorithm that combines sequential pattern mining and statistical testing to discover significant discriminative sequential patterns in event logs. These patterns capture sequential combinations of event attributes related to one or more process perspectives that are significantly more frequent for the cases in one event log compared to those in another log. A case study on identifying potential causes of defective products

in a manufacturing process demonstrates the use of the proposed algorithm within a larger process mining project. This case study shows that the algorithm generates highly selective sets of patterns, making them more interpretable and easier to use for process improvement.

Keywords: Business Processes, Process Mining, Sequential Pattern Mining, Significant Discriminative Sequential Patterns, Manufacturing.

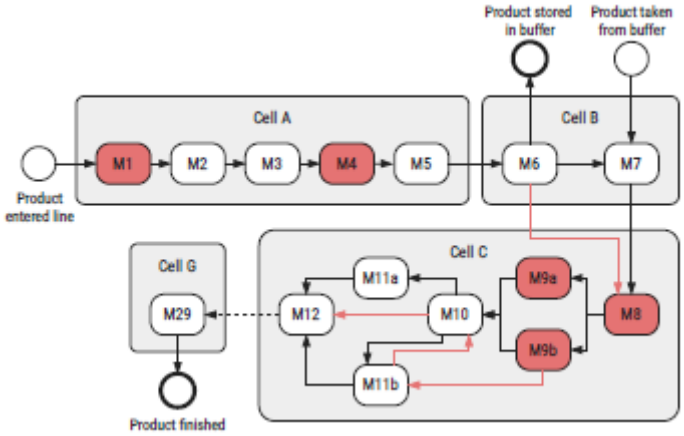

Figure 1: The as-realised process model of the OneBlade manufacturing process in which differences with the as-designed model are highlighted in red.

2.2.3 Evaluation of the Use Case

Part of the One Blade part of the Use Case, was the development and further improvement of a dashboard. See the next figure. The dashboard has made it possible to find the main root causes of Fall of Rate (FOR) earlier and lower the overall FOR. Within the processes, the efficiency and FOR are continuously monitored, and it has shown that both have shown an improvement (to within targets) compared to earlier years.



Figure 2: Screen shot of the developed dashboard



Figure 3: A Philips IPL device

Another part of the Use Case is the work done to improve the robot testing of IPL (Intense Pulsed Light) devices. The work was done in collaboration with TUE (Oscar Robben, Berk Türkçapar, Mike Holenderski and Nirvana Meratnia). A way of working is determined to generate and validate human motions by introducing the human skeleton in the models. This skeleton determines the possible movements. The way to translate data from one skeleton model to another is studied, because there is already extensive data available, but in different models. A presentation about the status and results was held during the meeting in Gdansk. In January 2026 two formats CMU (Carnagie Mellon University) and SMPL (Skinned Multi Person Linear) were used. The first has more restrictions (but more human like) and the second less restrictions (less human like) but is often used in literature of comparable research. From the 3 possible model types (Diffusion, Variation Auto Encoders and General Adversarial Network) a choice has still to be made. Towards the requirement KPI Time to Market, the goal is to show that the robot movements are made more human like. One difficult part is how to judge human likeness. Converting models to be able to use public data sets is done. Developing and testing motion trails is still on going.

The work is planned to continue after the closure of AIMS5.0, to reach later the goal of improved robot testing.

A third part of the Use Case was the improvement of a (ECG) process for production of male shavers. The model used in this process suffered from a drift, making it necessary to frequently readjust the model. The work was started in collaboration with TUE, but later it was decided to phase out the specific process (because of phasing out the parts made by the process).

2.3 UC3: Servitization and service delivery enabled by Zero administration (HQ, SEB, SA, LTU)

Connecting market actors involved in service delivery through a legally binding smart contract. This smart contract will serve as a hub for data sharing and transparency, it will initiate financial transactions and activities as well as support audits.

A key feature of the platform is a rapid interactive work process where a user can go back and forth between the above steps freely. This is enabled by the conceptual approach of working with streaming data throughout the full work process. The platform supports a broad range of models to cater for different use cases. Models include mathematical-statistical function, Ma-

chine Learning models including Neural Networks. Different model categories can also be combined for more advanced use cases. On device training of models and federated learning is possible. The platform is based on a real time data stream architecture designed for edge devices and microcontrollers enabling data streams on very resource constrained devices. The smallest version of the streaming engine only requires 100 kilobytes of RAM which makes it possible to enable data streams and edge AI models on midrange microcontrollers. The engine is independent of any other software and is portable to different configurations of processor architectures and OSes. The platform is designed to be modular, extensible, and interoperable with other IT systems, data sources and AI model frameworks which enables integration with other architectures, systems and components developed within the project. A standard frontend solution is integrated with the platform offering a good UX for the users. The platform is designed to be easy to use for data scientists, engineers, domain experts and other groups interested in developing edge AI solutions, without the need to have deeper coding skills. HQ is porting Stream Analyze edge AI streaming engine to chosen target device (MCU) in lawn mowers.

HQ, SEB and SA are responsible for the implementation of AI model framework(s) suitable for the use case. This includes implementation of techniques for visualization of real time data streams and model results suiting the use case.

SEB will take care for integration with Husqvarna and SEB IT systems via integration with the project architecture for utilization of model results in their IT infrastructures.

Use Case 3 has been led industrially by Husqvarna Group and demonstrates, using a real-world example, how for instance autonomous lawn mowing can evolve from a product-centric offering into an AI-enabled, service-oriented business model based on **Zero Administration**. The use case combines Edge-to-Cloud (E2C) continuum AI, service orchestration, Financial Digital Twin concept and secure distributed communication into a unified architecture that supports large-scale service delivery with minimal manual intervention.

During the final project period, the solution was validated in a real commercial environment through Husqvarna Group's **Cut Grass as-a-Service** offering Smart Lawn Service in Germany, operating with more than 100 customer contracts. This provided a realistic validation setting for the project outcomes, including the analytical models, orchestration concepts, and automation capabilities developed in UC3. This solution also provided valuable insights from current end-customers. The operational environment also supported scaling validation of the service platform for deployments exceeding 10 000 connected devices, demonstrating that the architecture is relevant not only technically but also at industrial fleet scale.

The final UC3 implementation reached operational maturity. The edge-AI layer ran on Husqvarna robotic mowers, with model lifecycle management handled centrally through Stream Analyze Federation Services (SAFS). Two live data paths were used in the use case: a ROS 2 path for robot telemetry and an acoustic path for blade condition. These enabled the AI models to be deployed on the mowers and used in an operational service context.

2.3.1 **Demonstrated solution**

The final UC3 demonstrator integrates several building blocks into a complete service delivery architecture.

A **Lawn Quality Assessment** model derives garden-condition scores from mower motion data, including accelerometer, IMU and fused local position information. These outputs can be used for service planning, customer transparency, and dynamic service-cost adjustment.

A complementary **Blade Sharpness Assessment** model uses acoustic analysis of cutting sound to estimate blade condition and predict when blades need replacement. This added a new sensing modality without changing the underlying analytics infrastructure and confirmed the flexibility of the edge-AI architecture.

An analytics dashboard visualizes the lawn-quality score per device, the blade-sharpness output, and the difference between raw sensor volume and the much smaller inferred results sent to the backend. This clearly demonstrates the data reduction enabled by on-device inference.

The service orchestration layer was further developed to support large-scale as-a-service operations. The platform integrates telemetry from more than 250 robotic lawn mowers and supports automated work order creation, intelligent routing of mower error codes to service partners, configurable customer contracts, and telemetry-based SLA monitoring. Personalized interfaces were refined to improve visibility for customers, service providers, and solution providers.

In parallel, UC3 further developed the SEB **Financial Digital Twin** concept, preparing the integration of operational device data with financial KPIs and future asset valuation. This supports the understanding of how connected products can contribute to circular business models and financial reporting.

UC3 also validated a secure distributed service architecture through LTU's Arrowhead 5.2 reference implementation. The proof-of-concept demonstrated secure communication between geo-distributed stakeholders and consistent authorization across both the control plane and the application data plane, reinforcing the scalability and trustworthiness of the overall architecture.

2.3.2 KPI achievement

UC3 successfully achieved the key performance indicators defined for the use case.

The KPI for reducing the cost of administration in service delivery was validated in the Smart Lawn Service implementation, where automated work order creation, telemetry-based SLA monitoring and service orchestration produced **more than 80% reduction in administrative effort**. This was confirmed in a real operational setting and not only in a laboratory environment.

The KPI for reducing AI model development time was also achieved. The project target of 80–90% reduction was exceeded, with end-to-end model development for service delivery reduced from an estimated 40 working days to approximately 3.25 days, corresponding to around **92% reduction in effort** as explained in our previous reporting.

The middleware and platform KPIs were met through the operational integration of data streams, edge inference, centralized model lifecycle management, and downstream service systems. The platform supports all stages of AI model development and deployment, and it was already accepted at TRL 7 during the project and carried into the final period.

The platform scalability KPI was also demonstrated through the service orchestration environment, which supports telemetry and insights from hundreds of active devices and has been scaling-tested for fleets exceeding **10 000 connected devices per single edge orchestrator**.

The academic impact KPI was addressed through publications, theses and teaching-related dissemination, including the published work on terrain issue detection and additional research outputs in progress.

2.3.3 Validation against the evaluation dimensions

UC3 delivered clear results in relation to the evaluation dimensions used in WP5.

From a **sustainability** perspective, on-demand blade replacement reduces unnecessary material use and service trips, while on-device inference lowers data transfer and cloud load. The Financial Digital Twin further supports circular business models and more responsible asset management.

From a **productivity** perspective, the substantial reduction in AI model development time and the automated edge-to-service-order loop reduce engineering and administrative effort per deployed model and per serviced machine.

From a **resilience** perspective, the edge runtime supports operation across a connectivity range from fully autonomous/offline to interactive, so analytics can continue without constant cloud connectivity. The secure distributed architecture also improves robustness across organizational boundaries.

From a **human-centricity** perspective, garden-quality insights are surfaced to gardening personnel and service teams, while the low-code/no-code development environment lowers the barrier for domain experts to contribute to model development and refinement.

2.3.4 Final improvements and lessons learned

During RP3, model porting and parameterization onto the SA Engine were completed and moved into field use during the mowing season. The acoustic blade-sharpness modality and the on-device dashboard were added in this period, strengthening the operational demonstrator.

A key lesson learned is that a generic, transport-neutral ingestion path made it possible to add a new sensing modality without changing the infrastructure, which was a major enabler of the rapid model-development result. The main limitation is data breadth: high in-distribution accuracy did not yet transfer across sites, confirming that fleet-scale, multi-site data is required for robust generalization.

LTU's work strengthened the interoperability, security and scalability of the UC3 architecture through the development of a fully compliant Arrowhead 5.2 Go reference implementation with Arrowhead Data-plane Authorization Policy Interfaces functionality. The proof-of-concept demonstrated policy-based authorization across both the Arrowhead control plane and the application data plane, using PEPs, an external PDP and certificate management to support secure distributed collaboration. LTU also contributed to predictive maintenance research, including terrain issue detection and audio-based blade-sharpness analysis, which supported the validation of AI-based maintenance methods for outcome-based service delivery.

2.3.5 UC3 — Public results available in the repository

Table 1: UC3 public results

Result	Type	Status	Public availability / repository	Owner(s)
Terrain Issue Detection for Servitization of Autonomous Lawn Mowing	Conference paper	Published (open access)	IEEE Xplore — DOI 10.1109/ETFA65518.2025.11205594	LTU, SA (Araujo et al.)
Arrowhead 5.2 Go reference implementation (ADAPI-enabled; PEP/PDP + CA-as-PIP)	Software (open source)	Released	GitHub — https://github.com/ulfbod/Arrowhead-experimental-core-systems	LTU
Multi-stakeholder result-oriented product-service systems incl. AI-based edge-to-cloud data analytics: a technology survey	Journal Paper	In pipeline	<tbd>	LTU (Araujo et al.)
Audio-based blade sharpness detection system for predictive maintenance of autonomous lawn mowers	Conference Paper	In pipeline	<tbd>	LTU, SA (Araujo et al.)
Multiplayer Online Game Technologies for Cyber-Physical Systems: A Survey	Journal Paper	Submitted	<tbd>	LTU (Borngrund et al.)

2.4 UC4: UC4 AI based improved connection between production and logistics (BMW, IPH, UPM)

A combination of the three intralogistics domains Handling, Transport and Commissioning via AI holds large innovation potential. By developing a three-layer AI system, implemented on an intelligent, autonomous transport vehicle, a versatile solution for application in intralogistics in different European Industries will be established.

Furthermore, the implementation of IoT devices with new functionalities (e.g., M2M) and a blockchain based eternalization of important data points will enhance the data collection and quality in order to establish a reliable foundation for AI algorithms.

IPH will support BMW to find adequate autonomous scan device technologies to cover the complete production site. Since ground bound devices like AGV do not cover all areas, innovative unmanned aerial vehicles (UAV) could be used as complement to fulfill the need of getting gapless production data.

▪ Construction of autonomous 3D-Scans [M1-M12] (BMW, IPH)

Autonomous scanning device with associated automated process for low-effort and sustainable re-scanning of the BMW production sites. Using an existing AGV for simple production integration and low maintenance costs and a reduced carbon footprint. The scanning device will be transported on a smart transport robot which is already in use at the BMW plants. Thus, the scanning can be accomplished whenever a transport robot is free. In the end of the day, we automated the digitalization for a more sustainable footprint, during the currently running digitalization. For recognizing scanned data, we use Artificial Intelligence to identify objects.

▪ Transfer to Omniverse & provide collected properties [M13-M24] (BMW, IPH)

The goal of this use case is to enable a step towards the virtual reality available production plant. By collecting and transferring the data into Omniverse, our data is used for training Artificial Intelligence, which is later used for running autonomous transport systems. Thus, with the use of Artificial Intelligence identified objects and data we create sustainable and reusable data.

▪ Training Artificial Intelligence [M25-M36] (BMW, UPM)

Digital and up to date image of our production sites, which is used for assembly line planning and training of industry capable Artificial Intelligence. Digital world with physical properties for training Artificial Intelligence based on BMW requirements. ▪ Concept creation for delivery to the assembly line [M1-M12] (BMW, IPH) Supply robot benefitting of the information, gained of the 3D scans, for the automated execution of the exchange of full and empty containers at the supply location

2.4.1 **BMW Contributions**

Within UC4, a customized autonomous drone solution was developed to support the digitalization of production and logistics environments. The platform is capable of autonomously navigating complex industrial areas while capturing both high-resolution 3D point clouds and panoramic imagery. The collected data provides a comprehensive digital representation of factory environments and serves as the basis for planning, analysis, and optimization activities across production and logistics processes.

The developed solution was successfully validated in industrial pilot environments. Autonomous drone-based data acquisition significantly reduces the effort required for capturing and updating factory layouts compared to conventional manual surveying approaches. The generated digital representations enable planning activities based on up-to-date and remotely accessible information, reducing the dependency on physical site visits and manual measurements.

To ensure efficient utilization of the captured data, several supporting software solutions were developed as part of UC4. A mission planning tool simplifies the preparation and execution of scanning missions by allowing users to efficiently define and manage autonomous flight paths. Furthermore, a dedicated processing workflow was developed to automatically generate navigation maps for Autonomous Mobile Robots (AMRs) directly from captured 3D point cloud data. This significantly reduces the effort required for AMR deployment and map updates.

In addition, a data quality assessment tool was developed to evaluate the completeness and quality of the acquired point cloud data. The tool automatically identifies insufficiently scanned areas and highlights regions that require additional data acquisition. This functionality supports the targeted planning of re-scans and helps ensure the creation of high-quality digital factory models suitable for industrial applications.

The developed software and hardware components were validated through pilot deployments in real production and logistics environments. The evaluation demonstrated significant improvements regarding planning efficiency, data availability, automation of scanning processes, and the generation of digital assets required for downstream applications.

2.4.1.1 KPI Validation Results

The validation activities performed within UC4 demonstrated positive results regarding the defined KPIs.

Time Reduction in Inventory and Planning Activities: The use of digital 3D scans enables planning activities to be performed remotely, reducing the need for repeated on-site inspections, travel, coordination efforts, and manual documentation. Pilot use cases demonstrated a qualitative improvement in planning efficiency and confirmed the potential to significantly reduce the number of required site visits.

Reduction of Planning Effort per Action: The availability of up-to-date digital factory information, combined with dedicated software tools, considerably reduces the effort required for planning activities. Validation in pilot scenarios confirmed a substantial reduction in planning effort compared to traditional approaches.

Reduction of Person-Hours per Scan: Autonomous drone-based data acquisition significantly decreases the manual effort required for factory scanning. Compared to conventional methods involving multiple personnel and extensive on-site work, the autonomous solution demonstrated a clear reduction in person-hours required for each scanning operation.

Reduction of AMR Map Creation Time: A major achievement of UC4 is the automated generation of AMR navigation maps from 3D scan data. Validation activities confirmed that map creation and update processes, which previously required several days of manual work, can now be completed within a significantly shorter timeframe.

2.4.1.2 Conclusion

The results achieved within UC4 demonstrate the successful development and validation of an autonomous drone-based scanning ecosystem consisting of hardware, software, and data processing tools. The combination of autonomous data acquisition, automated AMR map generation, and intelligent data quality assessment creates a scalable approach for maintaining up-to-date digital representations of production and logistics environments.

The developed solution contributes to increased productivity by reducing manual effort and planning time, supports resilience through continuously updated factory information, and provides the foundation for future digitalization and automation initiatives in industrial environments. Future activities will focus on the industrial deployment of the solution, including the training of designated operators across multiple BMW sites and the gradual rollout of the technology within the global BMW production network.

2.4.2 UPM contribution to UC4

2.4.2.1 Camera hardware evaluation and technology selection

UPM carried out a systematic analysis and evaluation of a wide set of depth cameras as candidates for UC4. The evaluation focused on key performance parameters including accuracy, range, frame rate, and cost, providing a comprehensive comparison to support technology selection. Measured results were obtained using a 3D model object as ground truth reference,

ensuring a controlled and repeatable assessment. Among all candidates evaluated, the PhotoneoS achieved the best accuracy results. However, its high price and limited operational range represent significant drawbacks for deployment within the project context, making it a less practical choice despite its superior precision. The Intel D405 achieved the second-best accuracy, but its range also raised concerns for the specific working distances required by the use case. Taking all parameters into account, the Intel D435f — used with its active infrared projector — emerged as the most balanced option, combining low reconstruction error, competitive cost, reasonable frame rate, and an operational range compatible with the UC4 requirements. The full comparison of measured results across all evaluated cameras is illustrated in Figure 4.

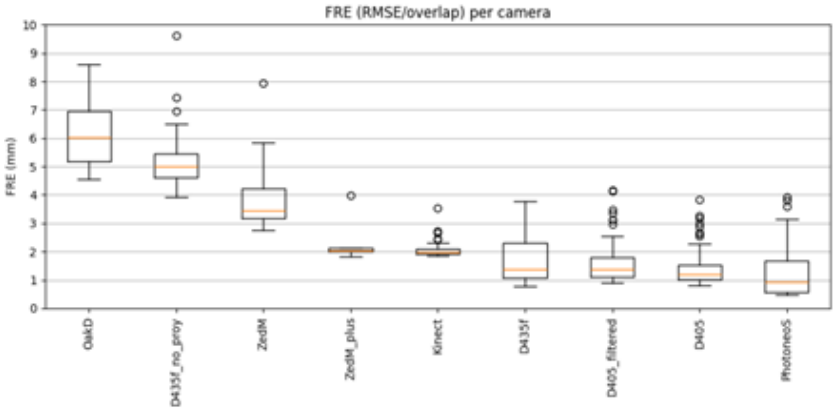

Figure 4: Measured accuracy for the cameras under test.

2.4.2.2 Multi-ToF camera rig and test sequence generation

UPM has generated two natural dynamic test sequences using a multi-ToF (Time of Flight) camera rig for test purposes within UC4. The two sequences differ in their capture configuration. The first sequence uses an 8-camera multi-ToF RGBD static array, providing a wide baseline and dense angular coverage of the scene. The second sequence uses a 7-camera multi-ToF RGBD static array arranged in a line at different heights, complemented by an additional movable camera. This movable camera is mounted on an x-y and pan-tilt actuator, which allows it to change its position in a plane parallel to the ground and adjust its orientation along two axes. This combination of a static array and a controllable movable camera provides greater flexibility in capturing dynamic scenes from multiple perspectives and at varying viewpoints. All RGBD cameras used in both sequences provide aligned full HD RGB images and corresponding depth maps, ensuring spatial consistency between colour and depth information. The capture setup is shown in Figure 5.



Figure 5: Multi-ToF camera rig

2.4.2.3 Multimodal camera rig and plenoptic sequence capture

A multimodal camera rig has been designed and built to facilitate the generation of controlled laboratory tests within the UC4 context. The rig integrates two ToF cameras and a light-field plenoptic 2.0 camera into a single compact setup. This combination allows simultaneous capture of depth information from the ToF sensors and light-field data from the plenoptic camera, enabling the study of sensor fusion and cross-modal reconstruction techniques. The plenoptic 2.0 camera, in particular, provides rich angular information within a single capture, which is valuable for depth estimation and view synthesis tasks. The camera rig is shown in Figure 6.



Figure 6: Multimodal camera rig

Additionally, two dedicated sequences have been captured using a plenoptic 2.0 camera mounted on an xyz actuator, which allows precise positioning of the camera in three-dimensional space. The xyz actuator is shown in Figure 7. The first sequence consists of 25 lenslet frames arranged in a regular 5×5 grid, with a spatial separation of 1 cm between positions along both the x and y axes. This grid configuration provides a structured light-field dataset suitable for evaluating multi-view reconstruction and depth estimation algorithms under controlled spatial sampling. The second sequence is composed of 192 lenslet frames of a dynamic scene, offering a significantly larger and more challenging dataset for testing reconstruction methods under real motion conditions. Together, these sequences provide a diverse and complementary set of test data covering both static structured captures and dynamic real-world content, supporting the evaluation of the full reconstruction pipeline developed within UC4.



Figure 7: 3D (xyz) actuator

2.4.2.4 3D reconstruction algorithm evaluation and digital twin development

To test the possibilities of the above mentioned technologies, UPM has been working in a particularly demanding application provided by BMW that introduces a set of challenging constraints. This application involves 3D reconstruction for aerodynamic modelling of a car during wind tunnel testing. First, the distance from the camera system to the car is set to be approximately 5 meters, which places strong demands on the depth estimation range and signal-to-noise ratio of the sensing setup. Second, the accuracy requirements are very high, with reconstruction errors required to remain below the millimeter level, which is significantly more demanding than typical industrial 3D scanning applications. Third, the system must be capable of capturing the car while its wheels are rotating at speeds up to 140 km/h, introducing motion blur and dynamic scene complexity that standard structured-light or passive stereo systems struggle to handle reliably.

To address these constraints, UPM has worked closely with BMW on two parallel lines of activity. The first is the generation of a 3D simulation environment that accurately mimics the conditions found in the wind tunnel, including camera placement, lighting, object geometry, and motion dynamics. This simulation allows rapid prototyping and evaluation of different sensing configurations without requiring physical access to the wind tunnel facility at every iteration. The second line of activity focused on the selection of the most suitable sensing technology. Following the evaluation and simulation work, a stereo camera setup combined with an active light projector was selected as the most appropriate configuration – see Figure 8. This setup provides the combination of active depth sensing and stereo geometry needed to achieve high accuracy at the required working distance.

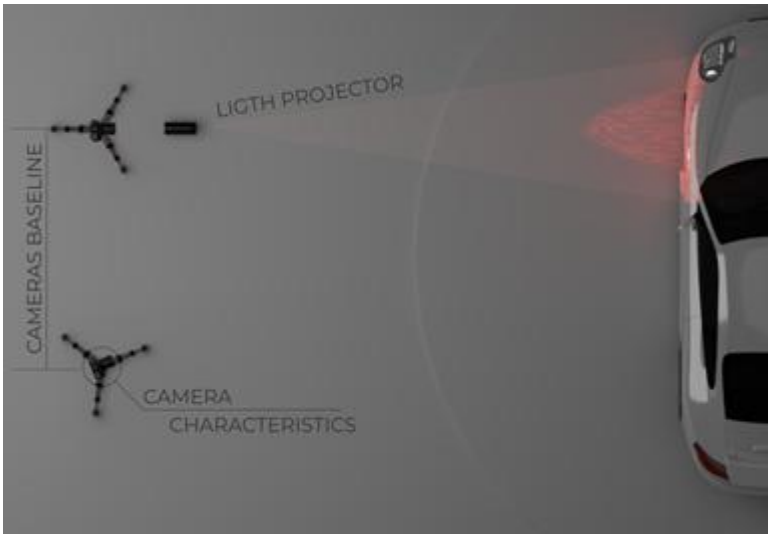

Figure 8: Wind tunnel simulation. The cameras are positioned 5 meters apart from the car.

UPM has developed a digital twin of a multi-view camera system for 3D reconstruction. The goal of this work has been to develop a simulation environment that accurately replicates the behaviour of a real multi-view camera setup, enabling the evaluation and optimization of 3D reconstruction pipelines without requiring physical hardware at every iteration. This digital twin approach reduces development time and allows systematic benchmarking of reconstruction algorithms under controlled conditions.

Several algorithms were evaluated as candidates for the reconstruction pipeline, including GPU-accelerated multiview depth estimation (GoRG), 3D Gaussian Splatting (3D-GS), and Multi-View Stereo (MVS). GoRG was originally developed for medical imaging applications and offers strong performance in structured environments. 3D-GS is a recent technique that represents scenes

as a collection of 3D Gaussian primitives, offering fast rendering and compact scene representation, but with limitations in geometric precision for high-accuracy industrial applications. MVS, which reconstructs dense 3D point clouds or surface models from a set of images taken at multiple viewpoints, was selected as the best candidate to meet the target accuracy requirements. Its robustness, maturity, and ability to produce geometrically accurate reconstructions from calibrated multi-view setups make it well suited for the precision demands of this use case.

2.4.2.5 Reconstruction results and validation

The evaluation was carried out by comparing the reconstructed point clouds against the ground truth — the actual 3D models of the car — using reconstruction error as the key performance indicator (KPi). Results obtained with both the stereo setup for the wind tunnel case and MVS for the digital twin case are promising. For the stereo setup, Figure 9 illustrates the results showing the reconstructed wheel as an RGB image (top left), the generated depth map (top right), and the 3D reconstruction overlaid with the ground truth model (bottom, in black and grey and red respectively). The close agreement between reconstruction and ground truth confirms the viability of the selected approach for the static case. On the other hand, for the multi-view camera setup, the reconstruction error is below 1 mm for 99.26% of the points in the point cloud, meeting the target accuracy defined for this application. This level of precision is relevant for industrial inspection and quality control scenarios, where sub-millimeter accuracy is often a hard requirement. The reconstruction density is also high, with approximately 3.82 million points covering a bounding box of 56 m^3. This density ensures that fine geometric details of the reconstructed object are well captured, including surface curvature, edges, and small features, without resorting to post-processing interpolation that could introduce artefacts.

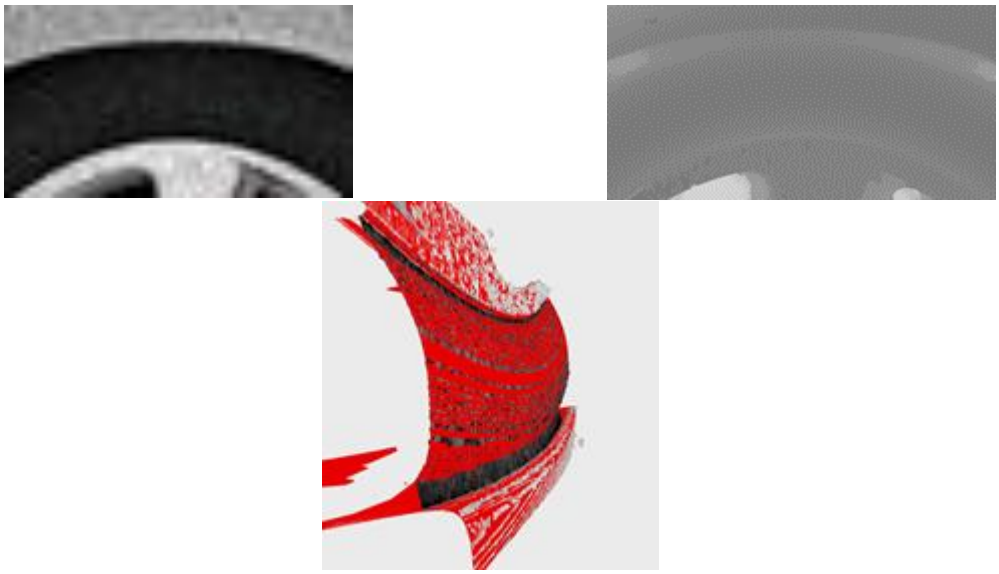

Figure 9: On top left, the RGB image of the wheel, on top right, the depth map generated using a depth estimation algorithm, on bottom, the 3D reconstruction of the wheel using the depth map generated (in black and grey) over the ground truth 3D data (in red).

Subjective results are shown in the figures below for the multi-view camera setup, which provide a visual assessment of the reconstruction quality alongside the quantitative metrics. Green points correspond to errors below 1 mm — within the target accuracy — and red points correspond to errors above 1 mm. Figure 10 shows the point cloud of the whole car, giving an overview of the global reconstruction quality, while Figure 11 shows a detail of the wheels, where the reconstruction is more challenging due to the complex geometry, reflective surfaces, and repeated patterns typical of wheel structures. As can be seen, errors are concentrated in some

parts of the wheels only, which is consistent with the known limitations of passive multi-view reconstruction in regions with low texture or high specularity. These localized errors do not affect the overall quality of the reconstruction.

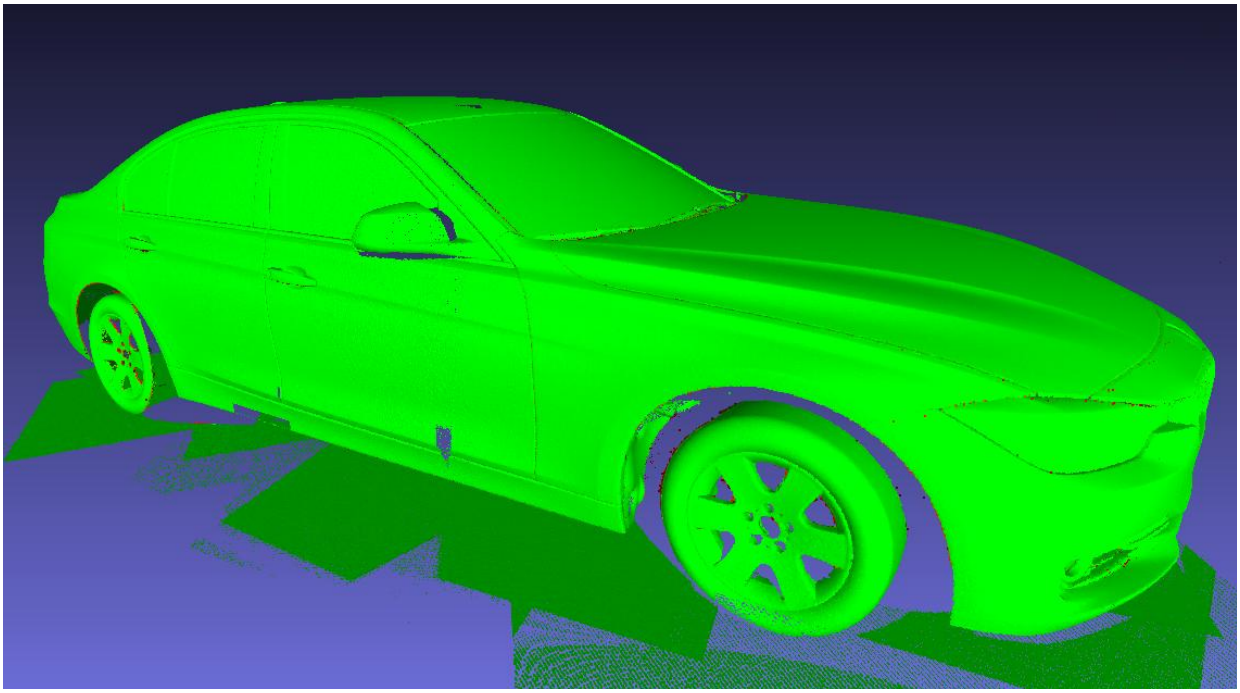

Figure 10: Generated point cloud of the whole car

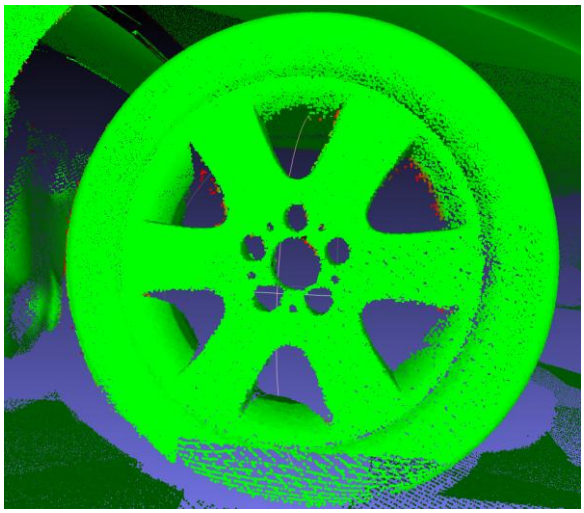

Figure 11: Detailed point cloud of the car wheels

2.5 UC5 AI techniques at different layers in Machine Tool Domain (IDEKO, SORALUCE, GOIMEK, SAVVY, MULTIVERSE)

The precision machining of large-size parts is used to operate with large part heavy-duty precision machining. The variability in the execution time for the same CNC machining program is very high (e.g., some parts can take from 3 up to 5 hours). The CNC program has some automated parts but it is mainly executed manually; thus, the expertise and skills of the operator become crucial. The process takes long time and the operator is not usually aware of the status

of the current operation nor has full control of the process. Following results/impacts are expected from this use case:

- Improvement of AI Methods for Machine tool domain processes => Optimization and automatization of execution of CNC programs (gain independency from operator's knowledge)
- Improved robustness of Machine Tool processes with monitoring and advanced control solutions.
- Increased product quality by avoiding faulty conditions using optimized and automatized controls.
- Increase efficiency in the company management processes by linking and making interoperable information in different layers of the company.
- Increase productivity of Machine Tool domain for a full industrial adoption

2.5.1 **Achieved results:**

The UC5 objective has been achieved: an AI-enabled monitoring and decision-support system has been developed for real-time tool condition monitoring and tool breakage detection in high-precision machining of large industrial components. The solution combines advanced signal processing, machine learning, edge computing, and secure industrial data infrastructures to detect tool failures, support predictive maintenance, and improve process reliability. We are now working on transferring and validating the solution in different machining scenarios and industrial environments. UC5 supports more productive, efficient, and sustainable manufacturing by reducing scrap, minimizing unplanned downtime, enabling earlier fault detection, and improving machine utilization and process stability.

1. Multi-layered infrastructure for data processing and AI deployment (SAVVY, IDEKO)

A secure edge-fog-cloud architecture has been developed to acquire, process, store, and analyse machine data in industrial environments. The infrastructure supports real-time monitoring, AI model deployment, production traceability, and integration with the Arrowhead framework. The architecture enables scalable deployment of AI applications through containerized microservices and efficient communication between production equipment and analytics services.

At the edge layer, industrial Jetway IPCs equipped with custom firmware provide sufficient computational resources to execute AI inference tasks in real time and include embedded microservice orchestration capabilities. They are directly connected to machine sensors and controllers, enabling local processing of vibration, temperature, and other machine signals.
At the fog layer, rack-mounted servers' aggregate data from multiple edge devices, enabling site-level analytics and coordinated AI services.

At the cloud layer, centralized platforms support large-scale AI model training, long-term data storage, fleet management, and global monitoring dashboards.

This architecture combines and integrates the component solutions described in the following.

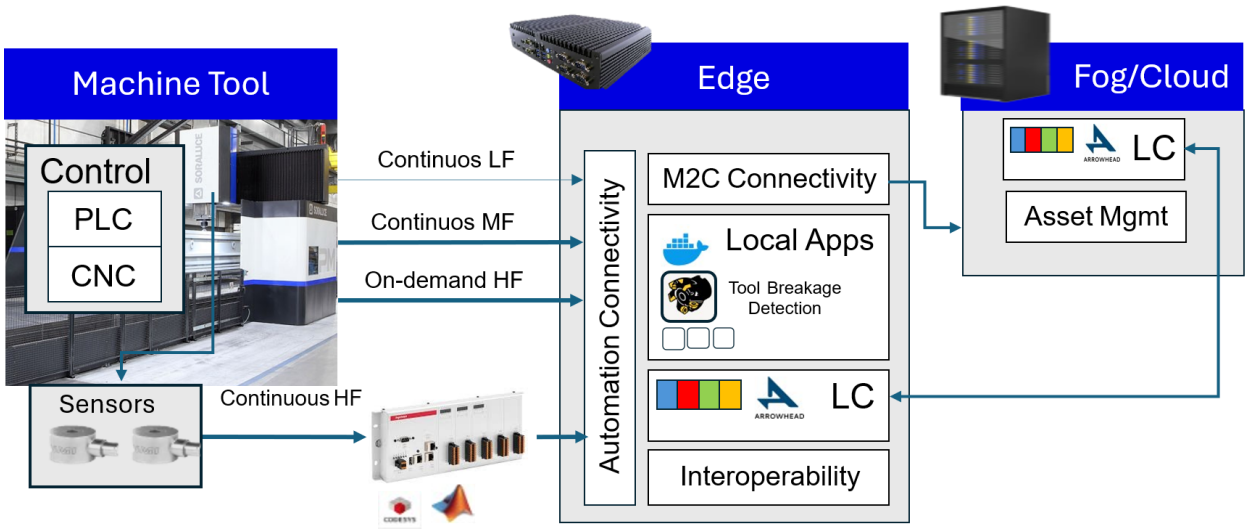

Figure 12: Multi-layered infrastructure for data processing and AI deployment

2. Intelligent Tool Monitoring & Breakage Detection Solution (GOIMEK, SORALUCE, IDEKO)

Integrated AI-based solution for real-time monitoring of machining processes and automatic tool breakage detection in milling operations. The solution combines machine-tool data, sensor information (vibration data), data preprocessing and AI algorithms to identify abnormal tool conditions and prevent defective machining operations.

3. Enhanced acquisition capabilities of sensor and machine data (SORALUCE, IDEKO, SAVVY)

For continuous high frequency (HF) acquisition of vibration sensors, a solution composed of a Programmable Automation Controller (PAC) device and the Edge device has been developed, suitable for the requirements of AIMS5.0 Use Case 05. Raw accelerometer data is ingested, and frequency domain-based features are extracted at 1s frequency and transmitted through Modbus to the Edge.

For continuous medium frequency (HF) acquisition, a communication library has been developed for the Edge platform, capable of extracting and storing internal data from Siemens Sinumerik CNC and Heidenhain TNC7/TNC640 controllers, without requiring modifications to machine programs or introducing additional processing load on the industrial PC. The data acquisition system enables medium-frequency capture in parallel with the machine's operational cycles.

4. Arrowhead integration in SAVVY multi-layered solution for AI solutions (SAVVY)

SAVVY has integrated the Arrowhead Local Cloud framework to orchestrate the UC5 machine-learning pipeline. The Local Cloud provides a secure service-oriented environment where data acquisition services, AI models, databases, monitoring dashboards, and deployment tools can communicate and exchange information through standardized interfaces. This enables AI models developed for tool breakage detection to be trained, deployed, executed, and monitored within the factory network, while maintaining interoperability, scalability, and low-latency communication between edge, fog, and cloud layers.

UC05 deploys the Savvy SmartBox edge device as the Arrowhead Local Cloud gateway at machine-tool level, with the SmartBox interfacing to CNC equipment via industry standard communication protocols, such as OPC UA, and feeding collected data to a fog-level AI-gym for predictive maintenance and process optimisation model development and deployment. Distinct

Arrowhead Local Clouds have been implemented for Edge, Fog, and Cloud layers, each hosting its own mandatory core systems (Service Registry, Authorization, Orchestrator), ensuring autonomy and localized service management within each tier.

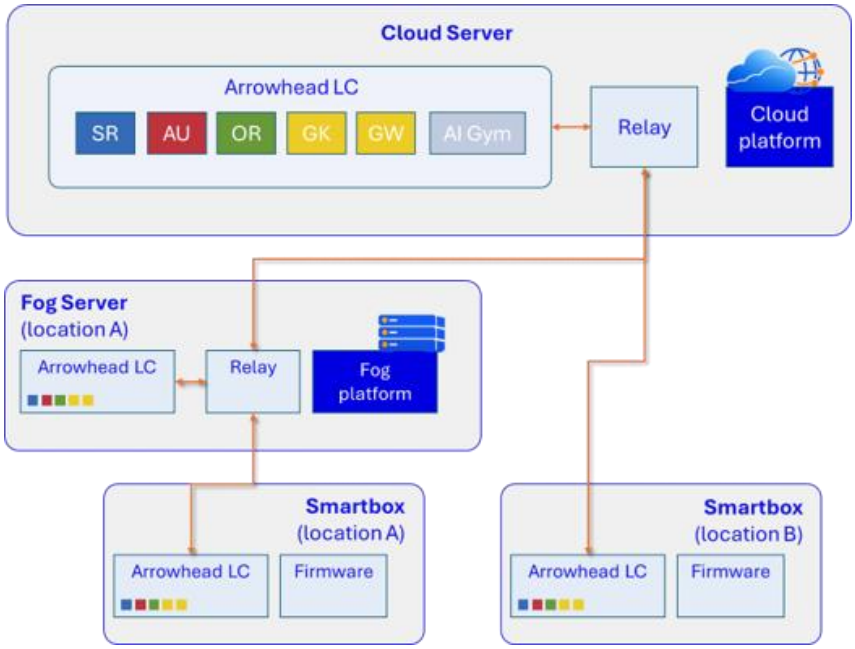

Figure 13: Arrowhead integration in SAVVY multi-layered solution

5. Motif Discovery and time-series segmentation (MULTIVERSE COMPUTING)

This technology automatically identifies recurring operational patterns in machining data and reduces the need for manual data segmentation. By combining motif discovery, clustering, and change-point detection techniques, the system helps identify normal process behaviour and detect deviations related to tool degradation and failure events.

6. Industrially relevant demonstrator

The developed technologies have been integrated and validated on a Soraluce PM and Soraluce FS production milling machine at GOIMEK, demonstrating end-to-end monitoring of machine, process, and tool condition. The demonstrator shows how AI methods, edge computing, advanced analytics, and industrial digital infrastructures can be combined to improve machining reliability, reduce defects, and support sustainable manufacturing. The UC has progressed from laboratory validation towards industrial deployment and exploitation.

7. Publications

Conferences and journal articles:

- A. Egibar and I. Llanos, 'Automatic signal segmentation using unsupervised learning techniques for process fingerprint identification in smart manufacturing systems', *Procedia CIRP*, vol. 138, pp. 186–191, Jan. 2026, doi: 10.1016/j.procir.2026.01.033.
- M. Uranga, G. Gutiérrez, I. Ricondo, *Federated Learning for Tool Breakage Detection*, 3rd European Symposium on Artificial Intelligence in Manufacturing (ESAIM 2025), Lecture Notes in Mechanical Engineering (proceedings pending publication).
- I. Llanos, A. Sanz, I. Ermina, A. Robles, and J. Munoa, 'Data-Driven Tool Condition Monitoring Platform for Grooving Operations with Chip Control Issues', Machines, vol. 14, no. 7, p. 811, Jul. 2026, doi: 10.3390/machines14070811.

2.5.2 Evaluation of the Innovation Project

Overall, UC5 has demonstrated how AI methods, machine-monitoring technologies, Arrowhead-enabled cyber architectures, and edge computing can be combined to create an industrially relevant solution for predictive monitoring and automatic tool-breakage detection in large-scale machining environments. The use case has progressed from laboratory validation towards industrial deployment and exploitation.

The main achievements are:

- **Implementation of a multi-layer edge-fog-cloud infrastructure** capable of acquiring, processing, and managing machining data in industrial environments. The architecture supports scalable deployment of AI services and production monitoring applications.
- **Development of FFT-based signal processing and KPI extraction methods** for monitoring tool condition through high-frequency vibration data. Tool harmonics and other relevant indicators are automatically calculated and analyzed in real time.
- **Development and validation of machine-learning algorithms for tool breakage detection and tool wear estimation**, using vibration-based features and supervised classification techniques.
- **Development of a real-time tool-breakage detection application** based on Dockerized microservices, integrating data acquisition, AI inference, notifications, MQTT messaging, and Grafana-based visualization.
- **Development of the Motif Discovery Toolkit**, enabling automatic segmentation and analysis of machining time series, reducing manual data-processing effort and supporting anomaly detection in production data.
- **Integration of Arrowhead Local Cloud technologies and AI-Gym**, providing a secure and interoperable AI/ML pipeline for model training, deployment, monitoring, and lifecycle management across the manufacturing infrastructure.

Table 2: Evaluation table

Prototype / Demonstration	Validation Outcome	Achieved KPI / Result
FFT-based signal processing and KPI extraction methods for monitoring tool condition	Validated in representative industrial environment	High-frequency data processing at 1-second interval
Regression and classification algorithms for tool breakage and tool wear	Validated in representative industrial environment	70% accuracy for specific detection of cutting-edge breakage. This increases to 90% when detecting other anomalies (e.g., runout caused by various issues), >95% in lab conditions.
Tool breakage detection application: tool failure detection in real environment	Validated in representative industrial environment	Detection in 1.5 s(from data acquisition to inference)
Motif Discovery Singularity Toolkit:	Testing tool by selected users.	Reduction in data segmentation time in 50% (compared to actual procedure of part program manual segmentation);

2.6 UC6 Reconfigurable AI-based automated conveyer feeding by a robotic system (EDI, HUA)

In this Use Case, an AI-based robotic system for conveyer feeding is developed, integrated and validated for organic cosmetics products. The robotic system will consist of several highly configurable AI software modules supported by respective hardware components: a) deep learning-based perception system, which provides the robots with the perceiving, comprehending, and reasoning abilities, that includes sensory data processing, representation, and interpretation; b) adaptive Deep Reinforcement Learning based planning module, c) high-level AI task instruction and context-aware task execution which improves this synergy. The following results/impacts are expected from this use case:

- A demonstrator: Reconfigurable AI-based automated conveyer feeding robotic system demonstrated an industrially relevant environment.
- A set of reconfigurable modules: Deep Learning-based perception modules (possible subcomponents such as sensing, information extraction and interpretation, knowledge acquisition and generalization), adaptive planning and acting modules.

The UC6 objective has been achieved: a modular and reconfigurable AI-based robotic system has been developed for automated bottle feeding onto a conveyor, capable of working with randomly arranged bottles of different sizes, weights, and types. We are now working to make the system transferable to other production environments and to adapt it to them. UC6 supports more productive, efficient, and high-quality manufacturing by reducing repetitive manual work, enabling reuse of generic robotic parts, lowering environmental impact, improving use of production floor space, and making production lines more flexible and resilient.

Within UC6 we have developed the following technologies, together forming a learning-based perception, planning, and control system that recognizes bottles and sockets, determines where they are located, separates each bottle or socket from the background and from other objects, understands object position and orientation, and plans and adapts the robot's actions accordingly. Two software results were co-developed by EDI and HUA: (i) the differentiable-physics robot-arm feedback controller for dynamic bottle placement, and (ii) the AI-based robot-control and simulation-synthesis pipeline that generates and validates task logic from natural-language descriptions. These modules complement the perception, pose-estimation, grasping and socket-tracking components in the end-to-end UC6 demonstrator.

2.6.1 Object detection and segmentation (EDI)

This technology detects and segments individual bottles in RGB camera images or video frames (see Figure: Object detection and segmentation). It is designed for scenarios where multiple bottles may appear in the same image, including unstructured or cluttered arrangements such as bottles lying in a pile, partially overlapping, or appearing at different orientations. The method processes each RGB frame independently and produces two main outputs for every detected bottle: an object detection output — a bounding box indicating the approximate location of each bottle in the image — and an instance segmentation output — a pixel-level mask that precisely separates each individual bottle from the background and from other nearby objects.

The processing pipeline consists of two stages. First, an RT-DETR object detection model is applied to the input image; it detects all visible bottles and returns bounding boxes with confidence scores, providing an initial estimate of where each bottle is located in the frame. In the second stage, each detected bounding box is used as an input prompt for the MobileSAM segmentation model, which refines the detection by generating a detailed pixel mask for the

corresponding bottle. This yields a more accurate object shape than a bounding box alone can provide.

Key features and KPIs

- Detection of bottles and sockets in RGB-D images
- Segmentation of regions of interest
- Robust perception in cluttered scenes
- Extraction of object masks for downstream processing
- Calculation time: < 30 ms
- Reconfigurability time: ~16 h up to a couple of days, using hand-labeled data

2.6.2 Grasp affordance estimation (EDI)

This technology calculates a suitable grasp pose for a detected bottle. The grasp pose defines the position and orientation that the robot gripper should reach in order to grasp the bottle safely and reliably (see Figure: Grasp affordance estimation). Using the technologies described above, the 6DoF pose of the object to be grasped can be used directly as a target gripper pose for motion planning in MoveIt2: once the bottle position and orientation are known, the system calculates one or more possible grasp poses based on the object geometry, the preferred grasping region, and the required gripper orientation.

Other approaches do not require the bottle pose to be estimated explicitly before grasp generation. For example, methods such as GraspGen can use RGB and depth streams from depth cameras directly as input; the grasp pose is then generated from the visual and depth information without a separate 6DoF pose estimation step in the pipeline. Regardless of the selected method, the final output is always a 6DoF grasp pose consisting of a position (the x, y, z coordinates of the desired gripper position) and an orientation (typically represented as a quaternion). The generated grasp pose is forwarded to the robot motion planning system — in the current architecture, to MoveIt2, which plans a feasible trajectory from the current robot configuration to the desired gripper pose. The grasp affordance estimation module thus acts as an important link between object perception and robotic manipulation.

KPIs — calculation time by method

- > 5 s using GPD (https://github.com/atenpas/gpd)
- < 6.1 s using GraspGen + upscaled depth map (https://github.com/NVlabs/GraspGen)
- 0.55–1 s using GraspGen + simplified segmentation
- < 100 ms using the ICP result + simple vector calculations for the grasp pose

2.6.3 Bottle 6DoF pose estimation (BME, EDI)

This technology estimates the 6DoF pose of a detected bottle — its three-dimensional position and orientation in space — using the known 3D model of the target object (see Figure: 6DoF pose estimation). It is intended for RGB-D camera-based perception, where visual information from the RGB image is combined with depth information to recover the spatial location of the bottle. The output is the bottle pose in a camera or robot reference frame, depending on the system configuration and calibration.

Two methods have been developed and tested. The first is based on the Iterative Closest Point (ICP) algorithm and uses the output of the object detection and segmentation module as input: the pixel mask of the detected bottle is used to extract the corresponding subset of the depth-based point cloud, representing the visible part of the bottle in the real scene. The ICP algorithm then aligns this observed point cloud with a reference point cloud generated from the full 3D model of the bottle; the resulting transformation provides the estimated 6DoF pose. The

second method is based on FoundationPose, a learning-based 6DoF pose estimation framework. It uses synchronized RGB and depth frames from the camera stream as input and requires a 3D object model of the target bottle in .obj format, specified in the corresponding configuration files. FoundationPose estimates the pose of the known object directly from the RGB-D data and can be used as an alternative to the ICP pipeline, especially where learning-based estimation is more robust under partial visibility, clutter, or challenging object orientations.

Both methods return the 6DoF pose of the bottle: position (x, y, z coordinates) and orientation (a quaternion). In the current implementation, the estimated pose is published in ROS2 using the geometry_msgs/msg/PoseStamped message type, so the module integrates directly into the robotic perception and control pipeline, where the pose is consumed by grasp planning, trajectory planning, and real-time robot control components. Overall, the technology provides a bridge between visual perception and robotic manipulation: the system determines where the bottle is located and how it is oriented, enabling the robot to reason about suitable grasping and manipulation actions.

Key features and KPIs

- RGB-D input and 3D point cloud processing
- 6DoF orientation and position estimation
- Reliable grasp planning for insertion
- Calculation time: < 30 ms; overall object/grasp pose estimation time: 200–500 ms
- Accuracy at distance 20 cm: angle < 3°, position < 17 mm
- Accuracy at distance 1.5 m: angle < 60°, position < 325 mm
- Reconfigurability: it has been validated that the technology works without reconfiguration when a new object 3D model is given

2.6.4 Grasp planning and execution (EDI)

This technology calculates and executes the robot trajectory required to grasp and pick up a detected bottle. The input is a target 6DoF grasp pose defining the desired position and orientation of the robot gripper. Standard MoveIt2 planning algorithms generate a feasible trajectory for the UR5e robot, taking into account the robot kinematics, joint limits, current robot state, and possible collision constraints in the workspace. A custom interface with the UR5e is used to send the planned trajectory to the robot controller and execute the grasping motion, which typically includes moving towards the grasp pose, closing or activating the gripper, and lifting the bottle after a successful grasp.

KPIs

- Planning time: 15–180 ms, depending on the distance between the initial and target position
- Unpicked object ratio: 3/100

2.6.5 Socket detection and tracking (EDI)

This technology detects the center position of sockets on a moving conveyor belt (see Figure: Socket detection) and classifies each socket as either filled or empty. The detected socket position is used by the robotic system to determine where the bottle should be inserted and to support synchronization between perception, conveyor motion, and robot actions. Detection and classification are performed by a YOLOv8m neural network trained on manually labelled image data: the model identifies visible sockets in RGB camera frames, returns their location including the center point of each detected socket, and simultaneously classifies the socket state — whether a bottle is already present or the socket is available for insertion.

The dataset for training and testing was collected using Intel RealSense D435 cameras operating at 30 FPS with an image resolution of 1280 × 720 pixels, capturing the conveyor belt and sockets during motion. The output — the detected socket center position together with its filled/empty classification — is forwarded to the robot control and task planning modules to support bottle placement, conveyor tracking, and decision-making during automated feeding.

Key features and KPIs

- Detection of socket centers and target geometry
- Tracking of targets on the production line; maintaining precise alignment during operation
- Calculation time: < 15 ms
- Accuracy at distance 50 cm: mAP@0.5 > 70 %, position accuracy 2 mm
- Socket detection accuracy: 98.7 %
- Filled/empty socket classification accuracy: 95 %
- Reconfigurability time (new socket, camera angle): ~3 h

The following two technologies form an adaptive planning and acting system that covers the complete robot trajectory task: how to grasp the bottle, along which trajectory the robot should move, where to place the bottle, and how to adapt the motion if the socket position, the object, or the environment changes.

2.6.6 Differentiable physics method — robot arm feedback control (EDI, HUA)

This technology implements a robot trajectory feedback control policy for placing a grasped bottle into a moving socket. A neural network model adjusts the robot trajectory during execution, allowing the robot to compensate for the motion of the socket on the conveyor belt. The model outputs target joint position and velocity values for the next robot-control time step; it receives as input the current state of the robot arm — joint angles and joint velocities — as well as the estimated positions of the target socket and the grasped bottle. The control model is trained using a gradient-based approach together with a differentiable physics simulator, which allows the system to learn how changes in robot joint commands affect the final bottle position and to improve the insertion trajectory during training. We developed this differentiable physics method as an alternative to — and more capable than — reinforcement learning. For safety, the robot state is continuously monitored to detect possible collisions with itself or with objects in the environment; this monitoring supports safer execution and can provide feedback to higher-level control modules if replanning or error recovery is required. The output is a sequence of robot joint commands that guide the robot arm during the final bottle insertion task.

Key features and KPIs

- Enabling easy adaptation and real-time feedback; complex tasks with low computational demands
- Unplaced object rate: 4 %
- Static socket error: 1.42 ± 0.77 mm; moving socket error: 1.55 ± 0.91 mm
- Orientation error: 0.21°–0.23°
- Policy step evaluation time: 0.3 ms per action
- Faster training and smoother results than reinforcement learning; differentiable hard-constraint penalties
- Architecture: feed-forward neural network, 1 hidden layer × 256 units, nonlinear activations
- Training time: ~22 h (small network, fast convergence); this is also the setup/reconfiguration time for a new object type, new task, or new pickup/drop location

2.6.7 AI-based robot control script synthesis (EDI, HUA)

In the simulation environment we have developed LLM-based code generation: this technology generates high-level robot control logic from a brief natural-language task description, and also generates configuration code for simulating the task in IsaacSim. The generated simulations can automatically randomize chosen aspects of the task environment and goal criteria — for example the number and initial positions of objects in the environment and goal positions for pick-and-place operations (see Figure: Robot control script synthesis). Execution of the task controller code in the task simulation is used during the code generation process, by an agentic code generation assistant, to test and correct the code being generated, and also allows the user — visually, and by inspecting the final state of simulation runs — to validate that the initial task description has been correctly interpreted by the code generation agent and that the generated implementation correctly achieves the task objectives. Validated task implementations accumulate in a code repository that is available to the code generation agent when generating new task implementations, serving as a reference for available APIs and coding patterns.

High-level AI task instruction means that the robot's task is given at a high, human-understandable level rather than as a detailed program with every motion command. Context-aware task execution means that the robot performs the task taking the specific situation — the context — into account, rather than acting identically every time. Together these enable a reconfigurable robotic line and efficient adaptation to environmental changes.

Key features and KPIs

- Input: natural-language text, free-form, 1–3 sentences; usable by a production engineer with no robotics coding skills
- Generated code runs fully offline
- Synthesis time: typically < 2 h of interaction end-to-end; tasks that do not require significantly new or more complex functionality than tasks already in the repository are typically generated in 30–45 min, while tasks requiring back-and-forth dialog with the code generation agent (to debug or refine new functionality) take 2–3 hours, depending on the complexity of the new features and the number of iterations
- Initial validation success: 97 %; full simulation success rate: 80 %; for tasks substantially similar to previously generated tasks: > 90 %
- Sim-to-sim transfer accuracy: 94.3 % (from the differentiable physics simulator to IsaacSim)

2.6.8 Industrially relevant demonstrator

We are now adapting our demonstrator to an industrial environment close to a real production setting, in order to demonstrate its practical applicability for automated bottle handling and placement tasks. We demonstrate how differentiable-physics methods, combined with vision models, RGB-D cameras, and grasp and trajectory planning, can make robotic conveyor feeding more flexible, reliable, and easier to configure. Our robot executes an end-to-end scenario, picking bottles from a bin and placing them into a socket with dynamically determined bottle and socket positions.

- TRL5 has been achieved; in the exploitation phase, work is progressing towards TRL6 (introduction into industry)

- We are working to advance the UC6 technology beyond the concept and laboratory-validation stage towards a sufficiently mature solution that can be demonstrated in industrially relevant conditions and further prepared for practical use and potential commercialization
- Setup time at a new location is being measured as part of this validation

2.6.9 EDI's role and contribution in UC6

- EDI developed advanced computer vision algorithms for object detection in cluttered environments, ensuring high-reliability robotic grasping.
- EDI developed a differentiable-physics planning module (in place of deep reinforcement learning), which showed better results and was trainable in roughly 20× fewer time steps than classical reinforcement learning; a scientific article has been written about this work.
- The differentiable-physics-trained model supports high-level task instructions and context-aware task execution — an adaptive planning and acting system that anticipates events and prepares for them in advance, thereby coping with unforeseen situations.
- EDI implemented a deep-learning-based perception system that provides robots with perception, understanding, and reasoning capabilities, including sensor data processing, representation, and interpretation, among other algorithms.
- EDI developed the workflow and system architecture to support the integration of AI algorithms into the automated conveyor-feeding robotic system.
- System reconfigurability is achieved both through a flexible middleware framework and through custom deep-learning-based algorithms — LLM-based code generation.
- EDI developed a modular and reconfigurable AI-based robotic system to automate the conveyor feeding of various types of randomly arranged bottles with a high degree of efficiency and effectiveness, replacing manual human work in the production process as it is currently performed in industry.

2.6.10 HUA's role and contribution in UC6

- HUA co-developed with EDI the differentiable-physics robot-arm feedback-control policy for placing a grasped bottle into a moving socket, including its simulator-based training and quantitative evaluation.
- HUA co-developed with EDI the AI-based robot-control and simulation-synthesis workflow, which translates short natural-language task descriptions into control logic and Isaac Sim configurations and validates generated implementations through simulation.
- These HUA/EDI modules are reusable software assets within the integrated demonstrator and provide transferable capabilities for real-time dynamic manipulation and rapid robot-task reconfiguration.

2.7 UC7 Digital Automated Luminaire Manufacturing platform (Signify, TU/e, possibly RUG)

Signify aims at further developing and bringing to realization an E2E digitally driven manufacturing method for the Luminaires (Professional market, B2B), with as characteristics:

- Agile and adaptive product architecture and processes
- Easy and effective manufacturing (assembly and disassembly) on- premises digitally fabricated parts, high level of modularity and late-stage configuration

- Maximally using local (raw) materials as to reduce BOM and remote sourcing: impact on transport, geopolitical disturbances and response time
- Securing maximally options for re-use of parts and recycling of materials (sustainability)
- Aiming at reshoring to Europe

Data will be used as to optimize productivity, predict and anticipate on anomalies using ML/AI, and achieve zero defect and zero waste, just in time, etc.

In year 3 of AIMS5.0 the focus was on the sub use cases (1) CV and AI for assembly monitoring and (2) Metal Printing. The other sub-use cases of Signify were finalized earlier and reported accordingly in the Y1 report, Y2 report and in the various deliverables.

We would like to mention that detailed explanation of this work has been reported in deliverables of WP1 (for HW and SW for 3D metal printing), WP2 (for AI algorithms) and obviously also in WP5.

2.7.1 On the use of Computer Vision and AI for monitoring assembly lines

The pilot study in our Valladolid (Spain) outdoor luminaire factory was executed successfully. The AI method (as described in WP2) allowed us, with the Operations team of the factory, to make conclusions. The final tool that was tested is a commercial one (not disclosed here) and was selected based on the experimentation with own Research algorithms. Our initial internal Research work allowed us to build up our competences and be able to 'ask the right question/set the correct specifications' towards the commercial tool provider. This will also facilitate further developments (see D5.4, on the exploitation potential).

The purpose of the CV tools was to recognize individual assembly activities from videos of the assembly process and to process/translate these data into dashboards that give information on:

- Average and variation of takt time, individual assembly activities and waiting times
- Anomalies in takt time and individual assembly activities
- Line-balancing for each different product
- Actual human minutes for each different product

And next to use this information to:

- Further improve the assembly efficiency
- Further improve quality
- Avoid quality issues
- Feed a (potential) digital twin

For one selected product type in the family (Coreline Tempo Medium) videos were recorded, annotated, the model was trained, and the first (standard) dashboards were filled. Next, 3 more product types (all Coreline Tempo) were added (recordings, annotation, training, dashboarding). The details of the outcome of the models are discussed in D9.7.

The conclusions are that for the assembly process monitoring and analysis for pre-selected products the tool is well performing (as the AI model was trained on this), but that for adjacent products (from the same family but not pretrained/no earlier data sets) often poor quality of results showed up. This is related to mainly the fact that:

(1) Small differences make big impact on recognition

(2) Assembly does not always take place in the same area in the LoS of the camera (the 'freedom' of the human worker to decide) and
(3) The assembly activities are too diverse and differ too much between the operators for proper recognition.

2.7.2 3D metal printing using DED technology

3D printing, in which a product is generated only based on a digital file (design + process) and the raw material is a very well understood and validated sustainable fabrication method, as it is an additive process that allows achieving zero waste production. Signify has successfully developed and launched polymer 3D printing technology for luminaire fabrication.

Signify has now developed a consumer-targeted 3D metal printing technology, focusing on outdoor luminaires manufacturing. This is a sustainable alternative method for energy-consuming die casting technology. A first-of-a-kind FUMO (functional module) proof-of-concept machine for metal printing, based on Direct Energy Deposition (DED) technology was developed. In contrast to powder bed printing, this technology allows for efficient manufacturing of large parts (products) but suffered (before the start of the AIMS5.0 project) from too expensive machinery, too low productivity (for our Signify luminaire business) and limited process quality and geometry options. Hence, these 3 limitations were set as the target of Signify in AIMS5.0. In year 3 following results were achieved:

2.7.2.1 Further developments on the low-cost DED printer machine

The most important additions to the baseline machine are:

- A laser curtain shield for personal protection
- Metal dust collection measures, including re-use of non-molten powder
- Implementation of various sensors for process and product control (for real-time monitoring of the process, and for use in data analysis – cf. the SW platform). Detailed analysis of, e.g., laser beam, powder stream and melt pool allowed for further process development and finding the optimal settings of the process and machine.

See below an overview of the machine realization:

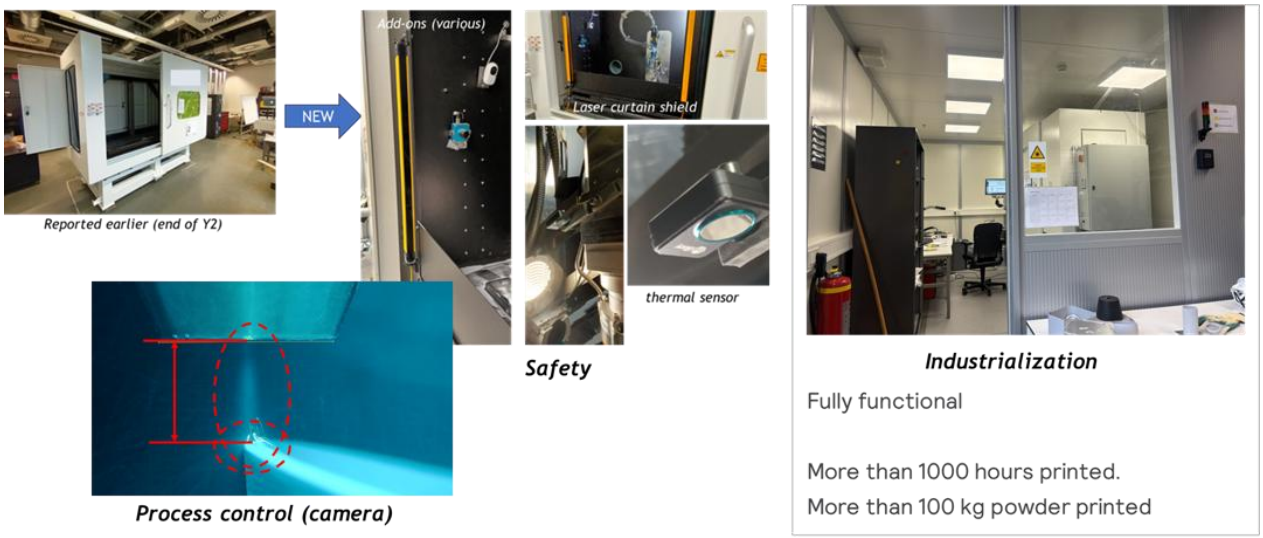

Figure 14: FUMO machine

2.7.2.2 Results of prototype printing

Below we show some examples of prototype shapes/products. We gradually approach product geometries and features that till now could only be achieved by metal powder bed printing (comment: powder bed printing which does not fulfil our Operational requirements on cost price and processing speed). Our results have brought the realization of customer products with DED 3D printing technology in reach, as shown next from our customer sample examples.

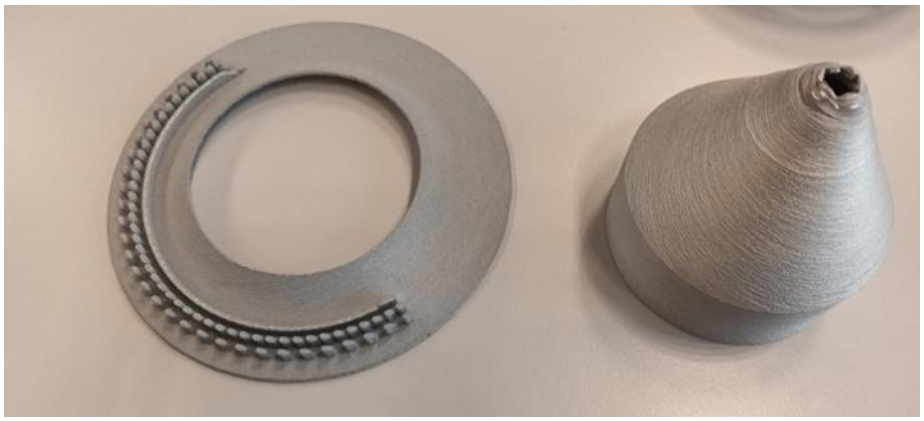

Figure 15: More complex centro-symmetric shapes (including overhangs)



Figure 16: Material finishing (powder coating) and odd-shaped geometries (square-like)

2.7.2.3 Business case

For a selection of key customers, in alignment with our business colleagues, products were fabricated for customer evaluation. The feedback is very promising and allows us to proceed towards further industrialization.

In the below figure some examples of such customer samples are shown:

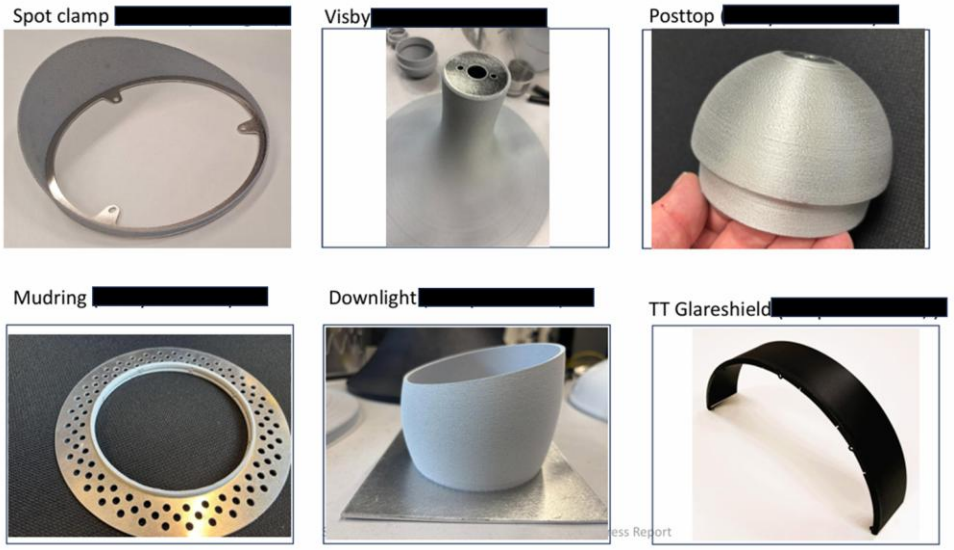

Figure 17: Customer samples

2.7.2.4 On KPIs:

Quantification of the KPIs is difficult and at the same time also confidential-critical.

- The machine we have developed is a first-of-a-kind, so it is almost impossible to estimate its cost price. Only pricing from an equipment manufacturer (something that should happen after the finalization of AIMS, in the exploitation phase) would make sense. We can however state that we expect that the price will be substantially lower than that of the high-end metal 3D printers on the markets. We dare to estimate a factor of 5x lower (comment: this is not proven, just an estimation). The goal of AIMS5.0, on this topic, was exactly to develop a lower cost/more productive machine that could meet our business case requirements.
- As we have not done comparable process development with an existing industrial (high end, high cost) metal 3D printer, comparison is not possible. Doing such process development with a high-end machine would not make sense anyway for Signify, as we know from the start (from industry-known cost price and productivity performance) that such exploration can't meet our business requirement anyway. So, it would have been a waste of time (and capex that would be wasted). Our estimations of machine investment combined with the results on operational parameters (print speed, machine usage) and metal powder price as we know by now are such that the business case for Signify, in relevant product categories, is judged to be solid.
- As we explained earlier, during the Madrid GA meeting (with the PO and reviewers), further detailing is also too confidential to share.

We can confirm at this time that we are on track on meeting the business case requirements on the use of DED metal printing for luminaire manufacturing, allowing to replace the die casting process of a substantial part our Outdoor luminaire product portfolio by the sustainable 3D metal printing method.

2.7.3 DAC contribution: synthetic defect generation, detection experiments and adaptable AI workflows

Scope. DAC investigated methods for creating labelled defect data when real defective samples are scarce, training and evaluating object detectors, applying the workflow to production video, and exploring approximate class removal from a trained model. The contribution also includes integration of the Open Access Platform with Eclipse Arrowhead and a Carbon Footprint component in cooperation with Infineon. The experiments support UC7 at proof-of-concept level; they do not constitute a validated production inspection system.

2.7.3.1 Controlled benchmark on the MVTec metal_nut dataset

A controlled benchmark was performed on one MVTec category (metal_nut) using the classes OK, SCRATCH and STAIN. The real-data split contained 162 training, 56 validation and 59 test images. One hundred profile-driven synthetic defect images were added to the training set. YOLO and RT-DETR models were trained with identical settings within each baseline/augmented comparison; the training-data composition was the intended difference.

Table 3: UC7 controlled benchmark on the MVTec metal nut dataset

Model	mAP@50	mAP@50-95	Precision	Recall	FP on OK
YOLO baseline	0.0311	0.0136	0.0361	0.0909	0.0%
YOLO + synthetic data	0.5665	0.3065	0.6083	0.5569	0.0%
RT-DETR baseline	0.6135	0.3261	0.7910	0.5961	0.0%
RT-DETR + synthetic data	0.6906	0.3700	0.8057	0.4848	4.08% (2 OK images)
Class / model		Precision	Recall	F1	mAP@50
SCRATCH - YOLO baseline		0.0723	0.1818	0.1035	0.0570
SCRATCH - YOLO + synthetic		0.7182	0.6971	0.7075	0.6581
STAIN - YOLO baseline		0.0000	0.0000	0.0000	0.0053
STAIN - YOLO + synthetic		0.4985	0.4167	0.4539	0.4748
SCRATCH - RT-DETR baseline		0.5820	0.7273	0.6466	0.7196
SCRATCH - RT-DETR + synthetic		0.8132	0.6364	0.7140	0.7317
STAIN - RT-DETR baseline		1.0000	0.4648	0.6347	0.5074
STAIN - RT-DETR + synthetic		0.7983	0.3333	0.4703	0.6496

Interpretation. For YOLO, synthetic augmentation increased mAP@50 from 0.0311 to 0.5665, reported in the notebook as a 1719.42% relative increase, without increasing false positives on the verified OK images. For RT-DETR, mAP@50 increased from 0.6135 to 0.6906 (12.58%), but recall fell from 0.5961 to 0.4848 and the FP-on-OK rate rose from 0% to 4.08%. The result therefore shows that synthetic augmentation can help, but its intensity and ratio must be tuned to the detector architecture.

2.7.3.2 Production-video experiment

Two production videos were available and the longer sequence was selected. Six hundred clean frames were extracted at approximately one frame per second. The frames were split temporally into 420 training, 90 validation and 90 test frames, so the three sets were disjoint. A fixed region of interest constrained synthetic defects to the product area. Four defect classes were generated: SCRATCH, STAIN, DISCOLORATION and MISSING_COMPONENT.

The generated dataset contained 1,560 synthetic defect images. Including retained OK frames, the final split contained 1,344 training images, 108 validation images and 108 test images. A YOLOv8n detector was trained for 50 epochs at 320-pixel image size. The training log reported 0.235 hours, approximately 14.1 minutes, on a Tesla T4 GPU.

Table 4: UC7 production-video experiment results

Metric	Final value	Interpretation	
mAP@50	0.8706	Measured on synthetic defects placed on held-out production-video frames	
mAP@50-95	0.6389	Synthetic test set only	
Precision	0.9612	Defect classes only	
Recall	0.8096	Defect classes only	
FP-on-OK rate	33.3% (6/18)	High false-alarm rate on clean test images	
Validation inference time	3.8 ms/image	Tesla T4; excludes transfer and application integration	
Defect class	Precision	Recall	mAP@50
SCRATCH	0.932	0.762	0.831
STAIN	0.924	0.955	0.976
DISCOLORATION	1.000	0.522	0.681
MISSING_COMPONENT	0.989	1.000	0.995

Validation boundary. The production-video experiment proves that the complete synthetic-data and training pipeline can operate on real product geometry. It does not prove detection of real manufacturing defects: all defect examples in the test set were synthetic, no actual defective production samples were available, and the 33.3% FP-on-OK rate is too high for an operational inspection system without further tuning and real-data validation.

2.7.3.3 API-backed image-generation verification

The API-backed NanoBanana variant was executed only in verification mode. On the MVTec data it generated one SCRATCH and one STAIN example. On the production-video data it generated one example for each of the four defect types plus one OK image. These runs verified technical image-generation feasibility, but no separate full training run or quantitative detector benchmark attributable to the API-backed variant was completed. No performance claim is therefore made for that variant.

2.7.3.4 Exploratory approximate class-removal experiment

An exploratory procedure started from the production-video YOLO model and removed the STAIN class by fine-tuning on a restricted three-class dataset for 10 epochs. A from-scratch three-class model was trained for 50 epochs as a comparison. The notebook reported 3.1 minutes for restricted fine-tuning and 15.4 minutes for retraining, corresponding to a 4.9x speed-up.

Table 5: UC7 exploratory class-removal experiment

Model	mAP@50	Precision	Recall	FP-on-OK	Training time
Original four-class model	0.7123	0.9185	0.6746	74.63%	Reused from earlier task
Three-class retraining	0.8196	0.8685	0.7415	5.97%	15.4 min / 50 epochs
Three-class restricted fine-tuning	0.6885	0.9484	0.5897	2.99%	3.1 min / 10 epochs

Critical limitation. This experiment is an approximate class-removal proof of concept, not certified machine unlearning. The STAIN metric was set to zero by evaluating three-class models whose output space no longer contained STAIN; the dedicated behavioural forget-set cell had no executed output. In addition, the notebook reported inconsistent dataset-directory counts between dataset construction and later evaluation scans. The timing and retained-class metrics are therefore treated as indicative, and no claim of proven or regulatory-compliant forgetting is made.

2.7.3.5 Open Access Platform integration

DAC completed integration of the Open Access Platform with Eclipse Arrowhead and the Carbon Footprint component in cooperation with Infineon. This is a cross-work-package integration result supporting secure service interaction and carbon-footprint-related data processing. The supplied notebooks contain no quantitative KPI for this integration, so no numerical performance or sustainability improvement is claimed here.

2.7.3.6 Consolidated evaluation against UC7 objectives

Table 6: Evaluation against UC7 objectives

Dimension	Supported conclusion
Productivity	The experiments automate synthetic dataset generation, training and evaluation. A direct reduction in manual annotation hours was not measured.
Quality	Strong synthetic-test metrics were obtained, but the production-video pipeline has not been validated on real defective products and has a 33.3% FP-on-OK rate.
Resilience	The workflow can generate controlled variants and adapt the class set; transfer to new products remains unmeasured.
Human centricity	No user study, operator acceptance KPI or human-review time measurement was performed.

Dimension	Supported conclusion
Sustainability	The OAP was integrated with the Carbon Footprint component; no quantified CO2 or energy reduction was measured in the supplied experiments.

2.7.3.7 Public-release and repository status

DAC does not release the source code, notebooks, model weights, datasets, production videos, annotations, logs, deployment configurations or experiment visualisations in an external public repository. No public repository URL is provided by DAC. The public contribution is limited to the high-level textual description and aggregated numerical results contained in the consortium deliverable. The implementation artefacts remain internal and proprietary because of commercial interests, partner confidentiality, production-data restrictions and third-party licensing conditions.

Table 7: Public-release and repository status

Artefact	Public status
Source code and deployment configuration	Not public
Google Colaboratory notebooks	Not public
Model weights and generated datasets	Not public
Production videos, extracted frames and annotations	Not public
Experiment logs and visualisations	Not public
External repository entry or URL	None
Text in the consortium D5.3 deliverable	Public at high level only

2.7.3.8 Final DAC statement for D5.3

DAC demonstrated a complete proof-of-concept workflow for synthetic defect generation, object-detector training and evaluation on a public benchmark and on real production-video geometry. The controlled benchmark showed substantial improvement for a weak YOLO baseline and a smaller improvement for RT-DETR, while also showing architecture-dependent trade-offs. The production-video experiment achieved mAP@50 of 0.8706 on synthetic held-out samples, but its 33.3% false-positive rate and absence of real-defect validation prevent a production-readiness claim. Approximate class removal was faster than retraining but was not behaviourally or formally validated as machine unlearning. The Open Access Platform was integrated with Eclipse Arrowhead and the Carbon Footprint component with Infineon. No DAC technical artefacts are released publicly.

2.8 UC8: AI supporting and strengthening human and manufacturing cycles (SU, TUBITAK)

AI, supporting and strengthening human and manufacturing cycles will be realized by 2 main developments that are reinforcing each other. (1) As a governor system, AI supported AMR Task & Traffic Management System that enables the safe and energy efficient intra-logistics inside any factory will be developed and demonstrated. Additionally, AMR sensor fusion, perception, obstacle detection and classification, as well as AMR human-like reaction with visual/sound

communication between machine and human obstacle will be developed and demonstrated. The intra-logistics part will be end-to-end material handling from warehouse until the work-station of human workers. (2) AI supported Human Assistance Exoskeleton using brain signals for human factory workers with disabilities will be assisting the human workers who wouldn't be able to / would have hard time without such an active exoskeleton. Material that has arrived to the workstation by AMR will be received and pro-cessed by the human worker, thanks to the Up-per-body exoskeleton to be developed and demonstrated in AIMS5.0 as first time.

The UC8 demonstrator successfully delivered an integrated Industry 5.0 solution combining AI-based perception, intelligent human-robot interaction, and assistive technologies for smart manufacturing. The project has produced several exploitable technology building blocks that can be further industrialized and incorporated into future commercial products. These include:

1. AI-based perception system for AMRs using camera and LiDAR sensor fusion;
2. Dynamic cost-map generation for safe robot navigation;
3. Human detection and tracking;

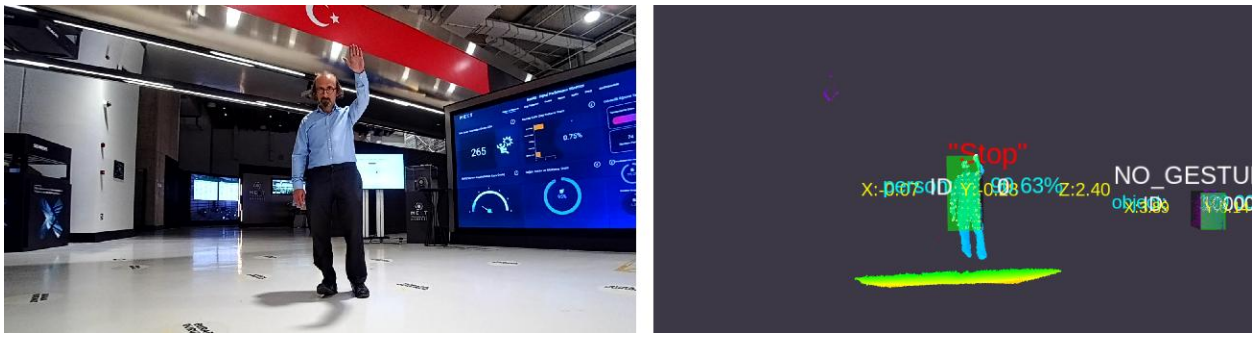

Figure 18: UC8 Human detection and tracking

4. Gesture recognition enabling intuitive human-robot interaction;

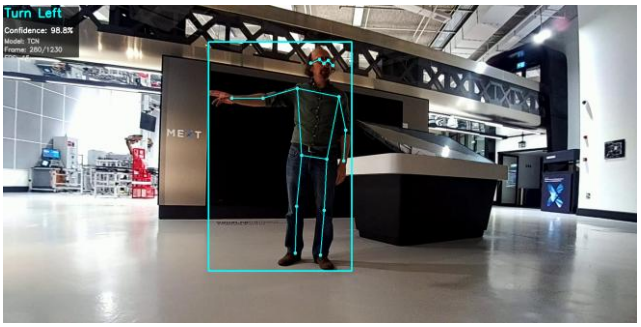

Figure 19: UC8 Gesture recognition

5. Audio-based communication module providing transparent robot intentions;
6. EEG-based intelligent exoskeleton control supporting manual material handling tasks.



Figure 20: UC8 EEG-based intelligent exoskeleton control

The developed technologies are modular and can be integrated independently or combined into complete human-aware robotic solutions depending on industrial requirements.

The developed Autonomous Mobile Robot (AMR) perception system integrates LiDAR and camera sensor fusion to generate dynamic cost maps, enabling safe navigation in environments shared with human workers. AI-based human detection, object detection, and gesture recognition allow intuitive interaction between workers and robots, while an audio feedback module enhances transparency by informing nearby operators of the robot's intentions (e.g., turning, obstacle detection, or task completion). In addition, an EEG-controlled exoskeleton was developed to support workers during physically demanding material handling tasks, contributing to improved ergonomics and worker well-being.
The developed prototypes were validated in representative industrial scenarios, demonstrating their capability to improve safety, productivity, resilience, and human-centric operation. All defined key performance indicators were successfully achieved or exceeded.

- The perception system reached a human-presence detection accuracy of 80% within a 5 meters range, gesture recognition accuracy of 85%, and 75% 3D bounding-box quality (IoU).
- Operational improvements included at least 25% reduction in task execution time, 30% reduction in traffic conflicts, and 90% task assignment accuracy.
- Furthermore, the AI-assisted exoskeleton supported the creation of new employment opportunities for workers with disabilities while improving workplace ergonomics by more than 5%, demonstrating the potential of Industry 5.0 technologies to create more inclusive and sustainable manufacturing environments.

Table 8: UC8 achieved KPIs and results

Prototype / Demonstration	Validation Outcome	Achieved KPI / Result
AI-based human detection	Validated in representative industrial environment	≥80% detection accuracy (≤5 meters)
Gesture recognition	Workers successfully interacted with AMRs using gestures	≥85% recognition accuracy
Camera-LiDAR sensor fusion & 3D perception	Reliable environmental perception demonstrated	≥75% IoU for 3D bounding boxes
Dynamic cost-map generation	Safe navigation around workers and obstacles	Supports collision-free path planning

AI-assisted task management	Validated in factory logistics scenarios	≥25% reduction in task execution time
Traffic management	Reduced robot conflicts during operation	≥30% reduction in traffic conflicts
AI task assignment	Automated task allocation demonstrated	≥90% assignment accuracy
EEG-controlled exoskeleton	Demonstrated worker assistance	Ergonomics improvement

Overall, UC8 successfully demonstrated that AI-enabled perception, human-aware navigation, natural human-robot interaction, and intelligent assistive technologies can be effectively integrated into industrial environments. The achieved KPIs confirm the maturity of the developed solutions and their potential to enhance safety, productivity, resilience, and human-centric manufacturing, providing a strong foundation for future industrial deployment and commercialization.

3. Conclusions

The AIMS5.0 project has demonstrated across eight complementary use cases that artificial intelligence can be effectively integrated into real industrial environments spanning manufacturing, logistics, and service delivery. The results reported in this deliverable confirm that the project has progressed beyond laboratory validation in several use cases, achieving operational deployment and real-world validation at industrial scale.

UC1 delivered a CO_2 fleet scorecard and energy management tool for injection moulding, providing manufacturers with a practical instrument for sustainable production planning. UC2 advanced the automation of testing through human motion data and process mining, with demonstrated improvements in production quality monitoring and fall-of-rate reduction. UC3 achieved one of the most complete industrial validation in the project, deploying an edge-AI service architecture across more than 100 customer contracts and demonstrating 92% reduction in AI model development time and more than 80% reduction in administrative effort in a live commercial service. UC4 delivered an autonomous drone-based scanning ecosystem — demonstrating significant reductions in planning effort, person-hours per scan, and AMR map creation time — alongside a multi-view 3D reconstruction pipeline achieving sub-millimeter accuracy for 99.26% of points at a density of 3.82 million points over a 56 m^3 bounding box, both validated in real BMW industrial environments with a global production network rollout planned as the next step. UC5 delivered an industrially validated AI monitoring system for tool condition and tool breakage detection in large-scale precision machining, reaching TRL deployment level with detection times of 1.5 seconds from data acquisition to inference. UC6 reached TRL5 with a complete end-to-end robotic conveyor feeding system achieving 98.7% socket detection accuracy and sub-2mm positioning precision. UC7 validated computer vision for assembly line monitoring at an industrial Signify facility and developed a first-of-a-kind low-cost DED metal 3D printer with a confirmed solid business case for relevant product categories. UC8 successfully demonstrated an integrated Industry 5.0 solution combining AI-based AMR perception and task management with an EEG-controlled exoskeleton, achieving at least 25% reduction in task execution time, 30% reduction in traffic conflicts, and 90% task assignment accuracy.

Across all use cases, the project has confirmed that the integration of AI as a cross-disciplinary, end-to-end approach — rather than as a collection of isolated tools — is both technically feasible and industrially relevant. The results provide a strong foundation for exploitation, further industrialisation, and future research building on the methods, datasets, and demonstrators developed within AIMS5.0.