Project acronym:	Lasers4MaaS
Project title:	Laser-as-a-Service Digital Platform with Dynamic Beam Shaping for Acceleration of Smart, Decentralised and Sustainable Factory of the Future
Call	HORIZON-CL4-2024-TWIN-TRANSITION-01-03 Manufacturing-as-Service: technologies for customised, flexible, and decentralised production on demand
Grant agreement No:	101178719
Work-package 7:	Algorithm development of real-time AI-based decision module using XAI/IML
Deliverable D7.1:	Algorithm for process synthesis, virtual sensors and process control, fast process development
Owner:	USTUTT
Due date:	31/05/2026

Type		
R	Document, report (excluding the periodic and final reports)	x
DEM	Demonstrator, pilot, prototype, plan designs	
DEC	Websites, patents filling, press & media actions, videos, etc	
DATA	Data sets, microdata, etc	
DMP	Data management plan	
ETHICS	Deliverables related to ethics issues	
SECURITY	Deliverable related to security issues	
OTHER	Software, technical diagram, algorithms, models, etc	

Dissemination level		
PU	Public, fully open, e.g. project website	x
SEN	Sensitive, limited under the conditions of the Grant Agreement	
Classified R-UE/EU-R	EU RESTRICTED under the Commission Decision No2015/444	

Disclaimers: this project has received funding from the European Union's HORIZON EUROPE research and innovation programme under grant agreement No. 101178719. The content of this publication is the sole responsibility of the Consortium partners listed herein and does not necessarily represent the view of the European Commission or its services.

HISTORY OF CHANGES

Version	Publication date	Change
1.1	16/04/2026	First draft
1.2	13/05/2026	Addition from TUW
1.3	15/05/2026	Addition from WMG
2.0	27/05/2026	Final version

LIST OF AUTHORS

Name	Organization	Email
Volkher Onuseit	USTUTT	volkher.onuseit@ifsw.uni-stuttgart.de
Michael Sawannia	USTUTT	michael.sawannia@ifsw.uni-stuttgart.de
Kim Glumann	USTUTT	kim.glumann@ifsw.uni-stuttgart.de
Jakob Neuhauser	TUW	jakob.neuhauser@tuwien.ac.at
Pasquale Franciosa	WMG	p.franciosa@warwick.ac.uk
Sattar Ghasemi	WMG	sattar.ghasemi@warwick.ac.uk

REVIEWED BY

Name	Organization	Email
Pasquale Franciosa	WMG	p.franciosa@warwick.ac.uk

TABLE OF CONTENT

1. Executive summary 8
2. Introduction 8
3. Objectives and Scope of WP 7 8
4. Methodological Framework 9
5. Work Performed and Main Results 12
6. Integration and Contribution to Project Objectives 25
7. Impact and Exploitation Potential 25
8. Conclusions and Future Work 27

LIST OF FIGURES

Fig. 1: Overall framework of the AI-based decision module 10
Fig. 2: Two-dimensional loss surface as created by the forward model for a target depth of 3 mm. 13
Fig. 3: Relationship between inputs and outputs for the process controller 14
Fig. 4: Logic flow of the process controller with integrated sensors 14
Fig. 5: Interaction between process optimiser and process controller 14
Fig. 6: PID controller developed around the modal ROM with constant setpoint in Simulink 15
Fig. 7: RL controller developed around the modal ROM with constant setpoint in Simulink 15
Fig. 8: PID controller constant temperature setpoint tracking 16
Fig. 9: Control action (laser power) by PID control law to follow the constant thermal setpoint 16
Fig. 10: RL agent constant temperature setpoint tracking 16
Fig. 11: Control action (laser power) by RL control policy to follow the constant thermal 16
Fig. 12: PID controller constant oscillating setpoint tracking, convergency issue after 4 s 17
Fig. 13: Control action (laser power) by PID control law to follow oscillating thermal setpoint 17
Fig. 14: RL agent oscillating setpoint tracking 17
Fig. 15: Control action (laser power) by RL control policy to follow oscillating thermal setpoint 17
Fig. 16: Optimisation workflow (Task 7.3). The visualization is adapted from a forthcoming publication by the authors. 19
Fig. 17: Visualization of model responses over a two-dimensional latent space (z[0], z[1]) 22
Fig. 18: Bayesian optimisation progress showing the evolution of cost values. 23
Fig. 19: Bayesian optimization progress showing parameter evolution, weld quality metrics, and cost values. Green shaded area: initialization phase. 24
Fig. 20: Synergies between process controller, remote monitoring and predictive maintenance 26

LIST OF TABLES

Table 1: Key challenges and methodological solutions 21

LIST OF ABBREVIATIONS

AI	Artificial Intelligence
API	Application programming interface
DBS	Dynamic Beam Shaping
FEM	Finite Element Method
FOM	Full Order Models
GA	Genetic Algorithm
KPIs	Key Performance Indicators
Log-EI	Logarithmic expected improvement
ML	Machine Learning
PID	Proportional Integral Derivative controller
RL	Reinforcement Learning controller
ROM	Reduced Order Models
VAE	Variational Autoencoder
WP	Work-package

1. Executive summary

This deliverable summarises the final outcomes of WP7, which addressed the development of AI-based algorithmic methods for optimisation-based decision support in welding-related manufacturing applications. The work focused on three functional areas within the project:

- process synthesis (T7.1)
- virtual sensing and process control (T7.2)
- fast process development (T7.3)

The main result of WP7 is a modular methodological and software framework that supports efficient parameter identification and optimisation across different welding tasks. A particular strength of the developed approach lies in its ability to operate across different levels of prior knowledge, ranging from scenarios without prior information to the use of simulation-derived models, historical data, or a combination of both. In addition, the framework is designed to support efficient parameter exploration under limited evaluation budgets while remaining compatible with integration into the project platform.

Overall, WP7 delivered an important contribution to the project's CONTROL-related functionality by providing integration-ready algorithmic components for process-oriented decision support. The results strengthen the project's capability to realise adaptive and resource-efficient manufacturing services in flexible production environments.

2. Introduction

The transition towards customised, flexible and decentralised manufacturing requires production systems that can adapt efficiently to changing product requirements, process conditions and quality targets. In such environments, conventional trial-and-error procedures and purely experience-based process development are often no longer sufficient, especially when process behaviour depends on multiple interacting parameters and when experimental evaluation is time- and cost-intensive.

Within this context, WP7 addressed the development of algorithmic methods for optimisation-based decision support in welding applications. The work package focused on approaches that combine data-driven modelling with available process knowledge to support process synthesis, fast process development and the integration of sensor-based information for process monitoring and control.

A particular challenge in this context is the need to develop methods that are not only technically effective, but also adaptable to different welding tasks, compatible with platform-based deployment and applicable across different levels of prior knowledge. WP7 therefore concentrated on the development of a modular framework that can operate in scenarios without prior information, with partial knowledge (e.g. simulation models or historical data), as well as in cases where both sources are available.

This deliverable presents the methodological basis, implementation approach and main outcomes of WP7. It describes the objectives and scope of the work package, summarises the developments achieved in its main task areas, and discusses their contribution to the overall project objectives.

3. Objectives and Scope of WP 7

WP7 focused on the development of a modular algorithmic AI framework for optimisation-based decision support in welding applications. The objective was to develop a methodological and software framework that uses machine learning to enhance welding-related decision-making and process adaptation. It can be

flexibly applied to different welding tasks, integrated into the project platform, and operated under varying levels of prior knowledge.

To achieve this, the framework was designed to operate across different levels of prior information. In scenarios without prior knowledge, optimisation starts from the admissible parameter space and relies on newly generated observations. In cases where prior information is available, simulation-derived knowledge (e.g. Reduced Order Models) and/or historical process data can be incorporated to guide and accelerate the optimisation process. This flexible strategy enables efficient parameter exploration across a wide range of use cases.

WP7 therefore contributes primarily to process synthesis and fast process development by supporting the identification and refinement of suitable process parameters for varying welding tasks and target requirements. Regarding virtual sensing and process control, the work package did establish interfaces and an algorithmic structure that enables the integration of such information into the optimisation workflow and platform architecture.

The scope of WP7 was thus centred on the development, adaptation and integration-readiness of the underlying optimisation algorithms.

Main objectives addressed in WP7

The main objectives of WP7 are closely aligned with its three core task areas:

T7.1 – Process synthesis

- Development of a modular optimisation framework for welding-related decision support
- Support of process synthesis through structured parameter identification and adaptation

T7.2 – Virtual sensing and process control

- Preparation for the integration of sensor-based and data-driven inputs into the optimisation and control workflow
- Development of standard and interoperable data format to bridge between high-fidelity numerical simulations (e.g., thermal modelling, and multi-physics simulations) and real-time sensor measurements, along with ROM techniques to solve the computational burdens during control design and optimisation
- Provision of interfaces and algorithmic structures enabling data-based and physics-based development of control algorithms and codes, ready for deployment to the actual laser welding machine

T7.3 – Fast process development

- Support of fast process development across different levels of prior knowledge
- Efficient parameter exploration under limited evaluation budgets
- Provision of a platform-compatible optimisation approach applicable across different welding tasks

4. Methodological Framework

This section describes the methodological framework developed in WP7 as the basis for AI-based optimisation-driven decision support in welding applications. The framework is designed to be modular, adaptable to different welding tasks, and compatible with the overall project platform.

A key characteristic is its ability to operate under different levels of prior knowledge, ranging from knowledge-free optimisation to the integration of simulation-derived information, historical data, or their combination. In addition, the framework is designed to allow the future incorporation of further process information, such as sensor-based inputs, without requiring changes to the optimisation core.

4.1. Overall Framework Concept

The framework follows a modular architecture that separates optimisation logic, information sources and platform interaction. This structure enables adaptation to different use cases and facilitates integration into the project software environment.

At its core, the framework consists of:

- an optimisation engine based on Bayesian Optimisation, responsible for parameter exploration and decision-making,
- a set of interchangeable information sources providing process-related knowledge, and
- defined interfaces for communication with the platform and external components.

This modular design ensures that the same optimisation approach can be applied across different welding tasks while allowing the integration of additional data sources or models when available.

The framework is illustrated in Fig. 1. The figure highlights the interaction between platform components, information sources and the optimisation engine, as well as the integration of the different WP7 task contributions.

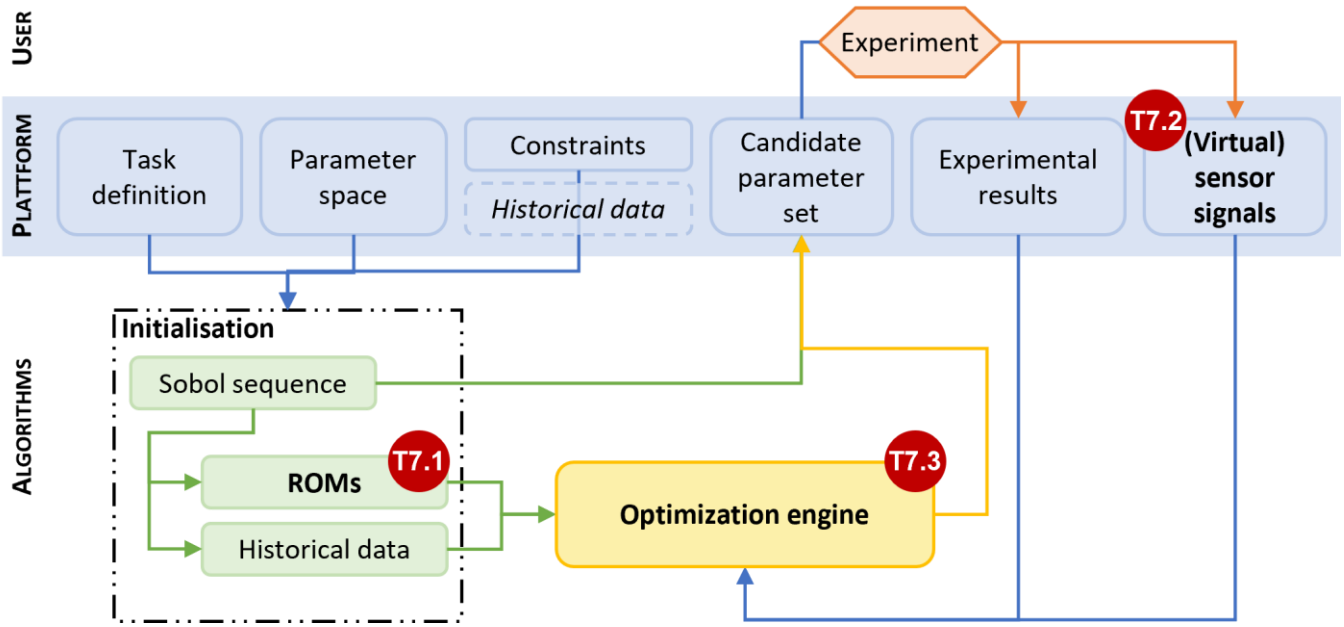

Fig. 1: Overall framework of the AI-based decision module

4.2. Handling of Different Levels of Prior Knowledge

A central methodological feature of the framework is its ability to operate across different levels of prior knowledge, rather than relying on a fixed optimisation mode.

Depending on the available information, the optimisation process can incorporate:

- no prior knowledge, relying solely on exploration of the admissible parameter space,
- partial prior knowledge, such as either historical data or simulation-derived information (e.g. Reduced Order Models), or
- combined prior knowledge, where multiple information sources are used jointly.

These information sources are not treated as separate modes, but as interchangeable and complementary inputs to the optimisation process. As illustrated in Fig. 1, they can be integrated flexibly, for example to guide initial sampling, inform the surrogate model, or support decision-making during the optimisation loop.

This approach enables the framework to operate robustly across different use cases, independent of the availability, quality, or transferability of prior knowledge.

4.3. Integration of Reduced Order Models (ROMs)

Reduced Order Models play a central role in enabling knowledge-informed optimisation within the WP7 framework. The ROMs are derived from high-fidelity simulations and provide computationally efficient approximations of relevant physics-based process relationships.

Within the framework, ROMs are not treated as fixed predictors, but as an additional source of information that can be queried during the optimisation process. This allows the optimisation algorithm to:

- guide the initial sampling of the parameter space,
- provide prior trends for the surrogate model, and
- support the interpretation of optimisation results.

This flexible integration allows simulation knowledge to be leveraged without making the optimisation process fully dependent on it, while enabling seamless combination with other information sources such as historical data. This ensures robustness in the presence of model inaccuracies or domain gaps between simulation and real-world behaviour.

4.4. Optimisation Strategy

The optimisation component is designed to iteratively explore the process parameter space and identify parameter settings that meet user-defined objectives, such as quality criteria, geometric targets, or other process-specific KPIs.

In the standard implementation, the optimisation engine is based on Bayesian optimisation using a Gaussian Process (GP) as a probabilistic surrogate model. The GP is iteratively updated with newly obtained process evaluations and provides both predictions of the objective function and associated uncertainty estimates.

Candidate parameter settings are selected using an acquisition function that balances exploration of uncertain regions and exploitation of promising regions of the parameter space (logarithmic Expected Improvement (log-EI)).

The optimisation process follows a sequential loop:

1. Initial parameter configurations are generated (randomly or guided by prior information such as ROMs),
2. Process evaluations are performed (experimentally, via simulation, or using existing data),
3. The Gaussian Process model is updated with the new observations, and
4. New candidate parameters are selected based on the acquisition strategy.

The loop continues until predefined stopping criteria are met, such as convergence of the objective function or a maximum number of iterations.

4.5. Interface Design and Platform Integration

A key requirement of WP7 was compatibility with the project platform. The framework was therefore designed with clearly defined interfaces for the exchange of task definitions, candidate parameters, evaluation results and optimisation states.
This interface structure supports coupling to external components, including experimental workflows, simulation-based evaluation and, in future, sensor-based information sources.

4.6. Applicability Across Welding Tasks

The framework was developed to support different welding scenarios and task definitions within a common methodological structure. This is achieved through generic task specification, modular handling of parameter bounds and constraints, and task-independent optimisation logic.
As a result, the framework can be adapted to different welding processes, geometries and target criteria, supporting its use in a broader manufacturing-as-a-service context.
Overall, the methodological framework developed in WP7 provides a flexible and integration-ready basis for optimisation-driven decision support. Its modular design, flexible handling of different levels of prior knowledge, and optional use of simulation-derived prior knowledge make it suitable for a wide range of industrial scenarios with varying levels of available information.

5. Work Performed and Main Results

5.1. Task 7.1 – Process Synthesis

Objective
High-fidelity physics-based simulation models, such as the compressible multiphase Mass-of-Fluid model used in the project, have demonstrated strong predictive capabilities across a broad range of laser-welding conditions without requiring process-specific calibration [1,2]. The high-fidelity model used is a two-fluid approach, meaning it explicitly resolves both the vapor and liquid phases, which is the key source of its advanced predictive capabilities. This predictive power, however, comes at a computational cost approximately one order of magnitude greater than that of simpler one-fluid alternatives [2].

Approach
Neither the one-fluid nor the two-fluid method is computationally cheap, making direct use of such models for real-time decision support or large-scale parameter exploration impractical. To overcome this limitation, data-driven reduced-order models as developed in WP9 and WP10 offer a compelling path forward. The core idea is to train machine learning algorithms on synthetic datasets generated from validated high-fidelity simulations. Once trained, these reduced-order models can infer results within seconds, making them well-suited for real-time decision support and large-scale parameter exploration. Furthermore, it alleviates the problem of the two-fluid approach being computationally more expensive, since this additional cost arises only during training data creation.

Results
In an initial demonstration for 316L stainless steel, using a limited variety of beam shapes, it was shown that weld pool geometry can be predicted from process input parameters, with median relative accuracies exceeding 97%. While this result is exceptionally good, it only solves the forward problem, in which we seek process outcomes given the input parameters. Solving the inverse problem, in which we aim to predict suitable input parameters for a given welding task, is more difficult because the relationship between input and output data is no longer a mathematical function. This means that a given welding result is achievable not by a single set of input parameters, but by a multitude of parameter combinations, thereby confusing

the reduced-order models.
For this problem, a solution could be found by sampling many points within a certain range and using forward prediction to assess whether the input parameters yield a satisfactory result. This brute-force approach is feasible because reduced-order models can make predictions almost instantly, enabling thousands of predictions within seconds. Fig. 2 shows a loss surface for a given target depth, where, for each point, the predicted penetration depth was compared to the target depth to compute an error. Laser power and the shape radius were varied, the welding speed was fixed at 0.25 m/s. Using this surface, it is possible not only to find an optimum but also to find alternative points.

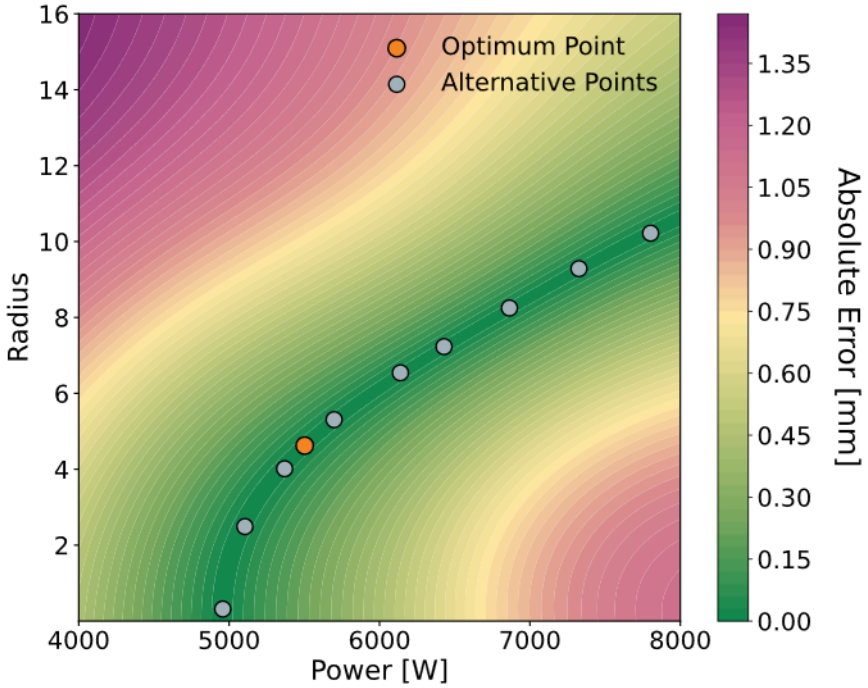

Fig. 2: Two-dimensional loss surface as created by the forward model for a target depth of 3 mm.

Key Takeaway
The results demonstrate that combining validated high-fidelity simulation with data-driven reduced-order modeling is a viable and efficient strategy for optimizing laser welding processes. High-fidelity models provide the physical accuracy needed to generate reliable synthetic training data, while reduced-order models translate this accuracy into practically usable tools capable of both forward prediction and inverse parameter identification at near-instantaneous speeds. Extending this framework to the full dynamic beam-shaping parameter space, as targeted in WP9 and WP10, is therefore a well-motivated and technically grounded objective that builds directly on the demonstrated methodology.

5.2 Task 7.2 – Virtual Sensors and Process Control

Objective
To develop algorithms for virtual sensors and process control as the key element for the service process controller as shown in Fig. 3. Accurate, computationally efficient, interpretable and explainable models are a key component of controller design and development for laser welding processes; yet systematic frameworks for control-oriented modelling of laser welding remain limited. While High-fidelity physics-based simulation models (FEM-based) represent the transient heat transfer dynamics with high fidelity, their large dimensionality limits their application in real-time control. The main objective is the exploit ROM techniques for fast development of control architectures capable to embed the fundamental physics of the process. Additionally, the approach does promote the standardisation of the data representation, through seamless integration across different workflows, allowing process simulations, sensor measurements, and control algorithms to exchange information consistently and efficiently.

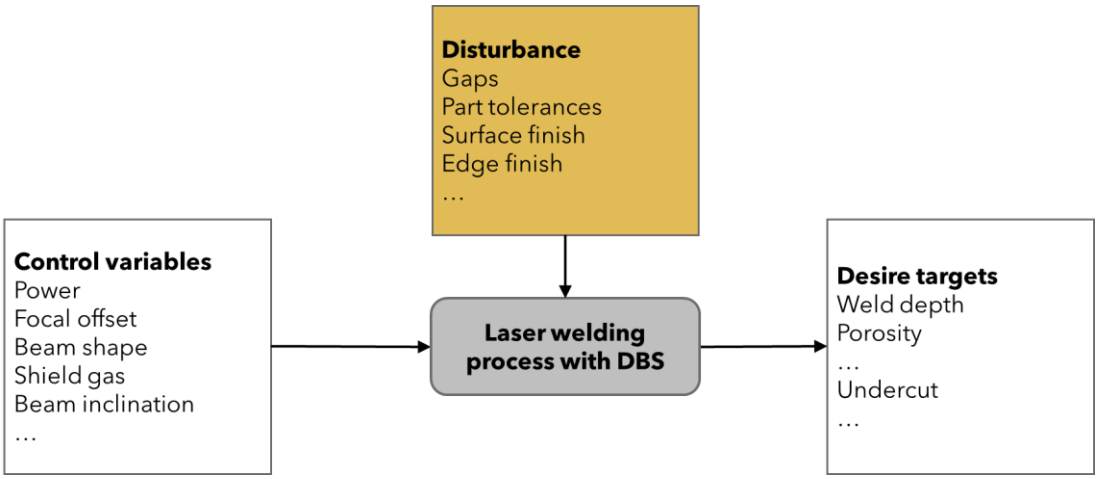

Fig. 3: Relationship between inputs and outputs for the process controller

Approach

This work provides one of the first systematic evaluations of ROM techniques for control-oriented modelling of laser welding processes as illustrated in Fig. 4 and Fig. 5. The approach is based on the development of set of ROM libraries, including balanced truncation, modal truncation and moment matching via Krylov subspace method. In parallel, data-driven models are derived from experimental data using linear and nonlinear system identification techniques including ARX, nonlinear ARX and Hammerstein model.

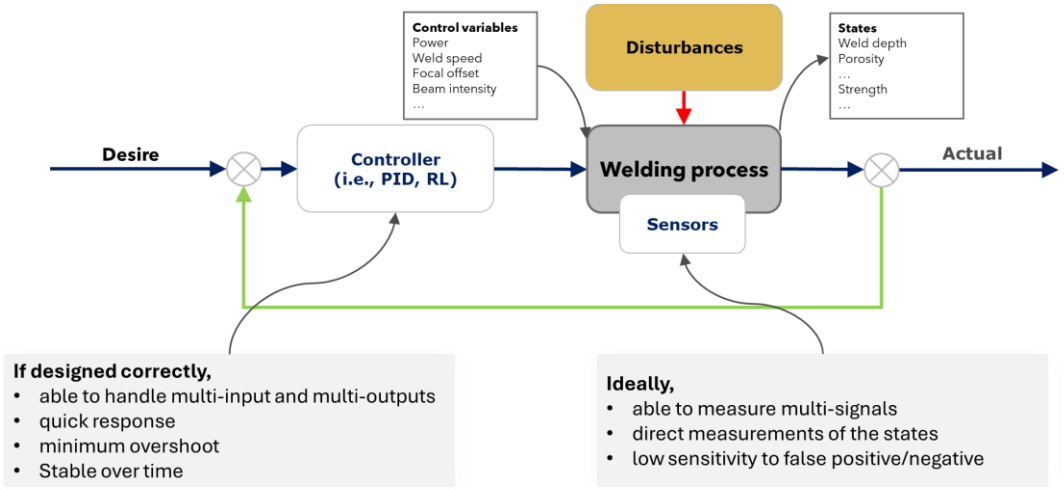

Fig. 4: Logic flow of the process controller with integrated sensors

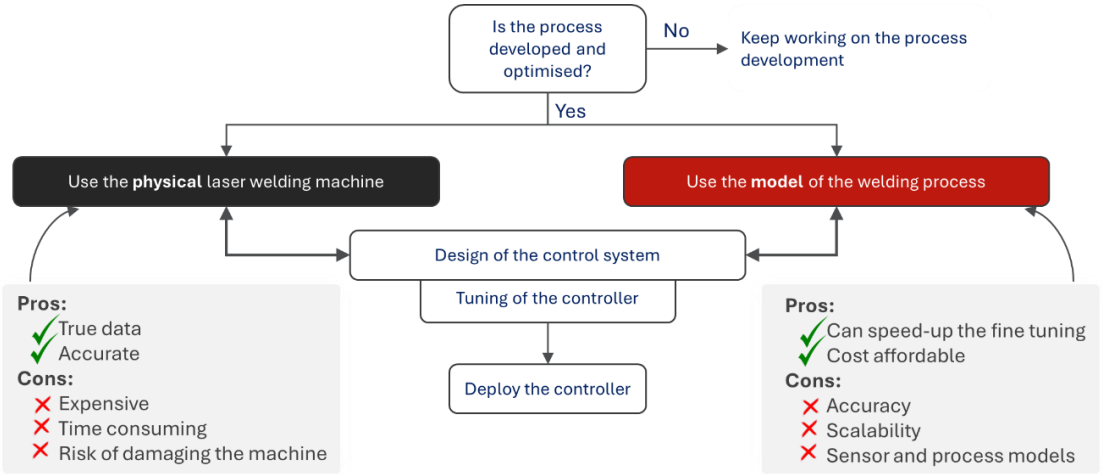

Fig. 5: Interaction between process optimiser and process controller

Results

The initial validation has been based on AA5754 (ref. battery cooler use case) laser welding. The results show that the main computational bottleneck stemming from FOM can be solved via ROM. Thanks to this approach, reinforcement learning (RL) policy training is expedited by orders of magnitude. A comparative assessment was performed between RL controller described in Fig. 6 and a conventional Proportional–Integral–Derivative (PID) controller described in Fig. 7 to evaluate their suitability for adaptive process control for laser welding tasks. The comparison focuses on tracking performance, stability, interpretability, and controller tuning complexity.

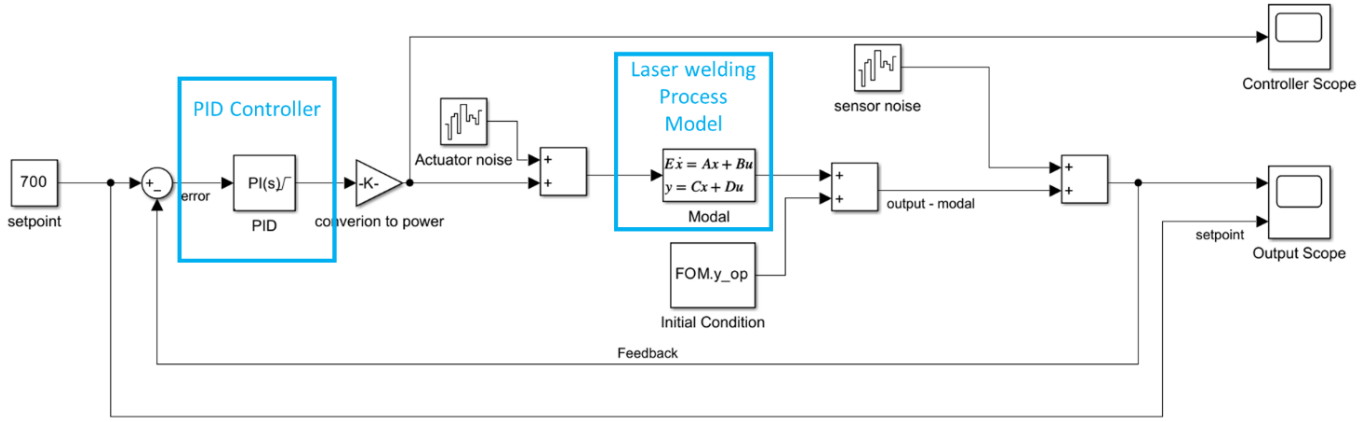

Fig. 6: PID controller developed around the modal ROM with constant setpoint in Simulink

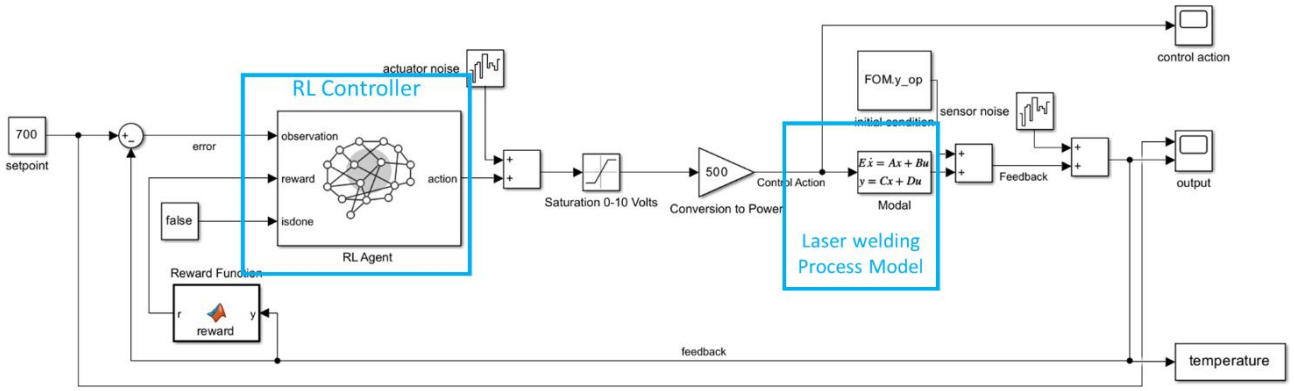

Fig. 7: RL controller developed around the modal ROM with constant setpoint in Simulink

The PID controller demonstrated smoother trajectory tracking behaviour with limited oscillatory response during operation. Due to its deterministic structure and well-established control formulation, PID control provides predictable and stable responses under relatively fixed operating conditions. However, the controller performance is strongly dependent on appropriate tuning of the gains. In highly nonlinear or dynamically changing environments (set point changing from constant to sinusoidal as shown in the Fig. 8, Fig. 9, Fig. 12 and Fig. 13), the PID controller may exhibit reduced robustness and lower stability margins, particularly when process dynamics vary significantly over time.

In contrast, the RL agent demonstrated the capability to learn adaptive control policies directly from interactions with the environment. When properly trained, the RL controller achieved highly stable behaviour across varying operating conditions and showed strong adaptability to nonlinear process dynamics. Nevertheless, the learned control actions occasionally exhibited jittering or oscillatory behaviour (as shown in Fig. 10, Fig. 11, Fig. 14 and Fig. 15) in the tracking response due to exploration effects, sensitivity to reward formulation, or high-frequency policy updates. Additional smoothing constraints or action regularisation methods may therefore be required for deployment in industrial systems.

From an interpretability perspective, the PID controller remains advantageous because its behaviour can be

directly understood and analysed through the physical meaning of the controller gains. Engineers can intuitively relate changes in gains to system response characteristics such as rise time, overshoot, and steady-state error. Conversely, the RL agent relies on deep neural network representations, which inherently operate as black-box models with limited transparency regarding internal decision-making processes. This reduces explainability and can complicate validation and certification.

The tuning effort associated with each controller also differs significantly. PID controllers typically require moderate tuning effort, which may be performed manually through empirical methods or automatically using optimisation and evolutionary algorithms (for example, GA). In comparison, RL-based controllers require substantially higher configuration effort, including design of the reward function, selection of observation spaces, neural network architecture definition, hyperparameter optimisation (e.g., learning rate, discount factor, replay buffer size), and activation function selection. Training additionally requires large quantities of representative interaction data and computational resources to ensure convergence and robustness.

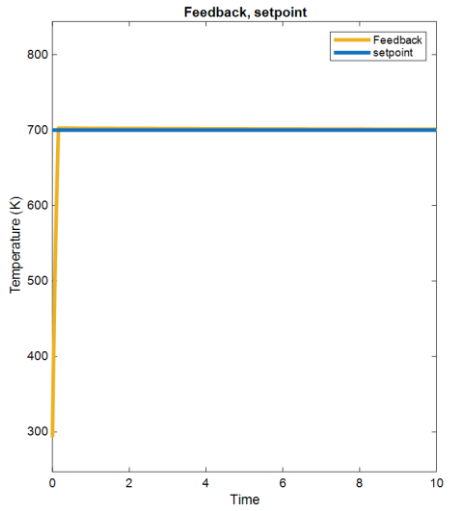

Fig. 8: PID controller constant temperature setpoint tracking

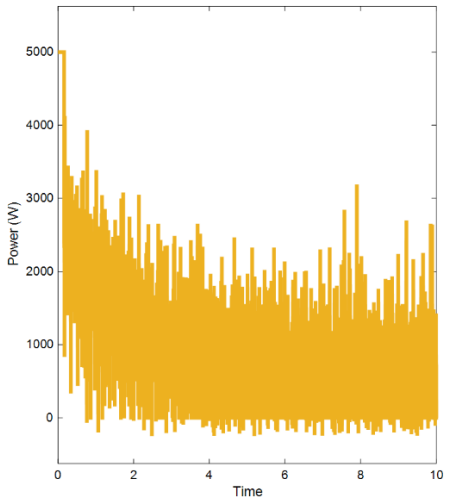

Fig. 9: Control action (laser power) by PID control law to follow the constant thermal setpoint

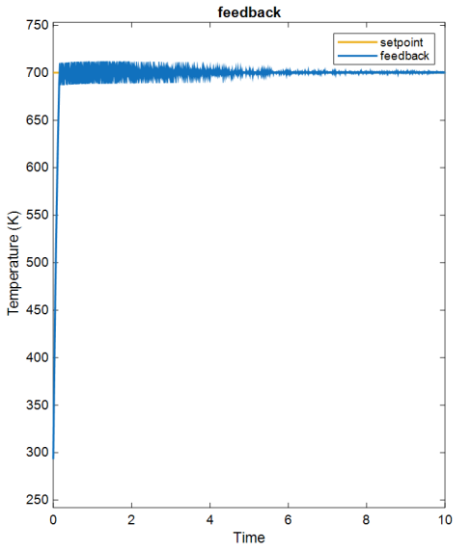

Fig. 10: RL agent constant temperature setpoint tracking

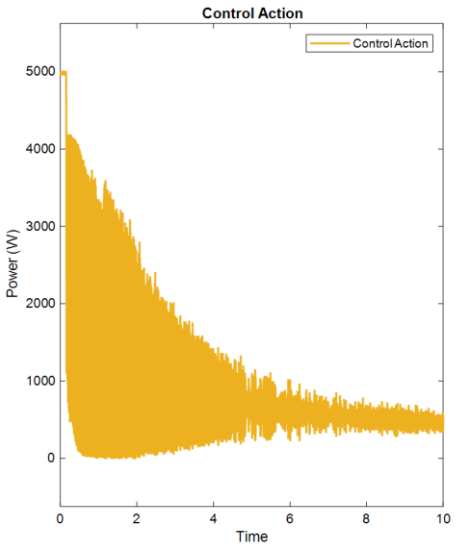

Fig. 11: Control action (laser power) by RL control policy to follow the constant thermal

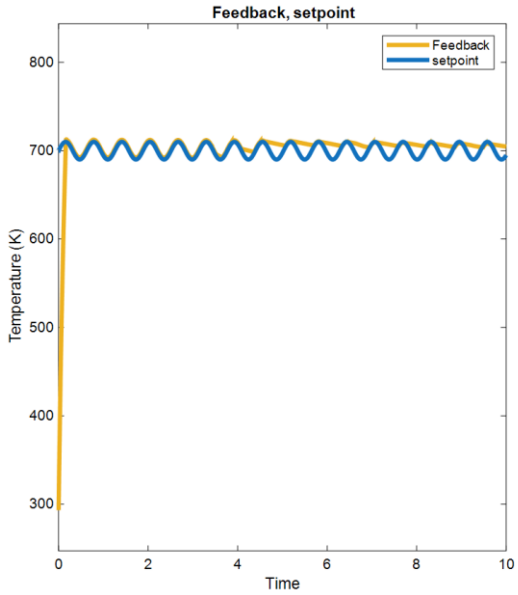

Fig. 12: PID controller constant oscillating setpoint tracking, convergency issue after 4 s

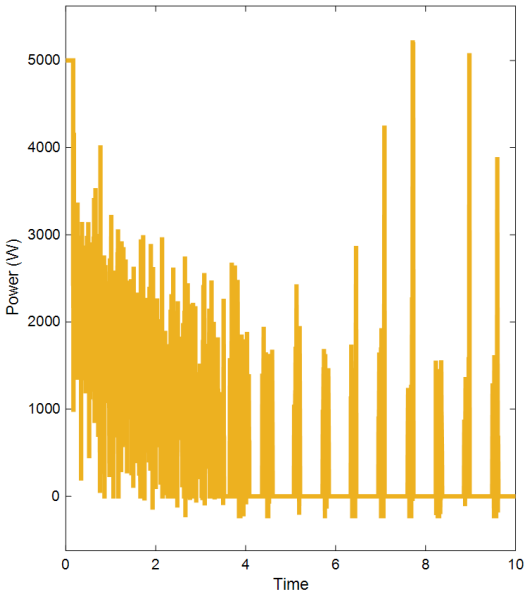

Fig. 13: Control action (laser power) by PID control law to follow oscillating thermal setpoint

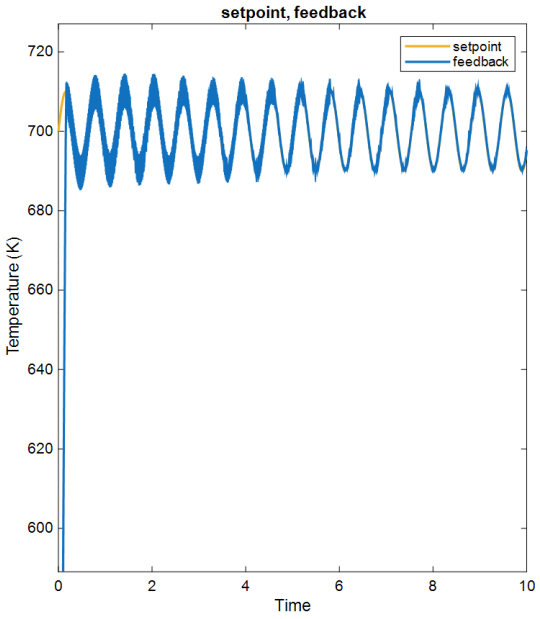

Fig. 14: RL agent oscillating setpoint tracking

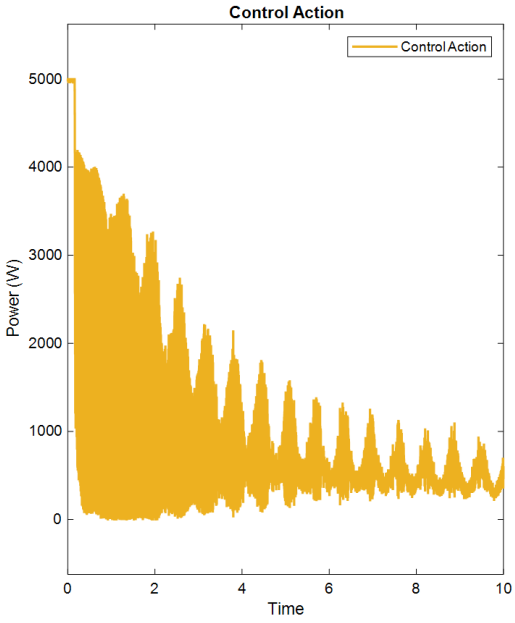

Fig. 15: Control action (laser power) by RL control policy to follow oscillating thermal setpoint

Key Takeaway

Overall, the comparison highlights that PID controllers remain effective for stable and well-understood processes requiring transparency and smooth control action, whereas RL-based controllers provide greater adaptability and robustness for complex nonlinear systems at the expense of increased implementation complexity and reduced interpretability. The choice of controller therefore depends on the required balance between adaptability, explainability, computational effort, and process variability within the target application. The work in WP11-12 will provide additional testing data for the DBS system in order to further validate the control methodology, and train the developed ML-algorithm in WP8 along with

integration of sensor data.

5.3 Task 7.3 - Fast Process Development

Objective
The objective of Task 7.3 was to develop a modular optimisation framework for fast process development in welding applications, enabling efficient identification and refinement of suitable process parameters across different tasks and target objectives. A key requirement was that the approach should operate across different levels of prior information, ranging from scenarios without prior knowledge to the use of simulation-derived models, historical data, or their combination. The framework was designed to be flexible with respect to different parameter spaces, objectives, and constraints, while remaining compatible with platform-based deployment. The focus of Task 7.3 was on the development of underlying optimisation algorithms and integration-ready components, which will be verified and validated in WP8.

Key Challenges and Design Requirements
The development of an optimisation framework for welding process design is associated with several key challenges, which guided the methodological decisions in WP7.

C1. High-dimensional parameter spaces
Modern welding processes, particularly those involving high-dynamic beam shaping, lead to high-dimensional parameter spaces with a large number of degrees of freedom. Direct optimisation in such spaces is computationally demanding and limits the applicability of standard optimisation techniques. Efficient representations are therefore required to make the optimisation problem tractable.
C2. Limited evaluation budget
Process evaluations in welding are typically costly and time-consuming, as they involve either physical experiments or high-fidelity simulations. As a result, only a limited number of evaluations is available, requiring optimisation methods that can efficiently explore the parameter space and converge with minimal data.
C3. Multi-criteria process objectives
Welding processes are characterised by multiple, often competing quality criteria, such as geometric properties, stability, and defect minimisation. This makes the formulation of suitable objective functions non-trivial and requires structured approaches to combine multiple targets into a meaningful optimisation objective.
C4. Varying availability of prior knowledge
The availability of prior process knowledge varies significantly across use cases. In some cases, neither historical data nor suitable simulation-derived models are available, while in others either or both can be used. Since ROMs are not universally applicable across all materials and process conditions, the optimisation framework must operate robustly across scenarios with no, partial, or combined prior knowledge.
C5. Low requirement for user expertise
A key requirement of the framework is its usability within a platform context by users who may not have deep expertise in process modelling or optimisation methods. The system must therefore minimise the need for manual parameter tuning and expert knowledge and instead provide structured and largely automated workflows for task definition, objective construction, and optimisation.
C6. Integration into platform-based workflows
The optimisation approach is not intended as a standalone tool, but as part of a larger digital platform. This requires clearly defined interfaces, separation between optimisation logic and evaluation, and compatibility with heterogeneous execution environments, including experimental setups and simulation workflows.

Approach
To address the challenges and design requirements outlined above, a modular Bayesian optimisation framework based on Gaussian Process models was developed. The framework is designed to handle

different welding tasks within a unified optimisation logic while allowing task-specific adaptation through the definition of parameter bounds, constraints, objectives, and prior knowledge sources.

To enable integration into the project platform, a dedicated application programming interface (API) was defined in collaboration with the project partner Futurice. This API specifies the data structures and communication protocols required for the interaction between the platform and the optimisation module. In particular, it supports the exchange of:

- task definitions, including parameter spaces, constraints, and objective specifications,
- candidate parameter configurations generated by the optimisation algorithm,
- evaluation results obtained from experiments or simulation workflows, and
- the serialised optimisation state, including model parameters and observed data points.

This explicit interface design ensures a clear separation between optimisation logic and execution environment and directly addresses the requirement for integration into platform-based workflows. It enables deployment across different platform contexts and supports asynchronous and distributed evaluation processes, which is particularly relevant for computationally expensive or experimentally constrained welding applications. The overall optimisation workflow is illustrated Fig. 16.

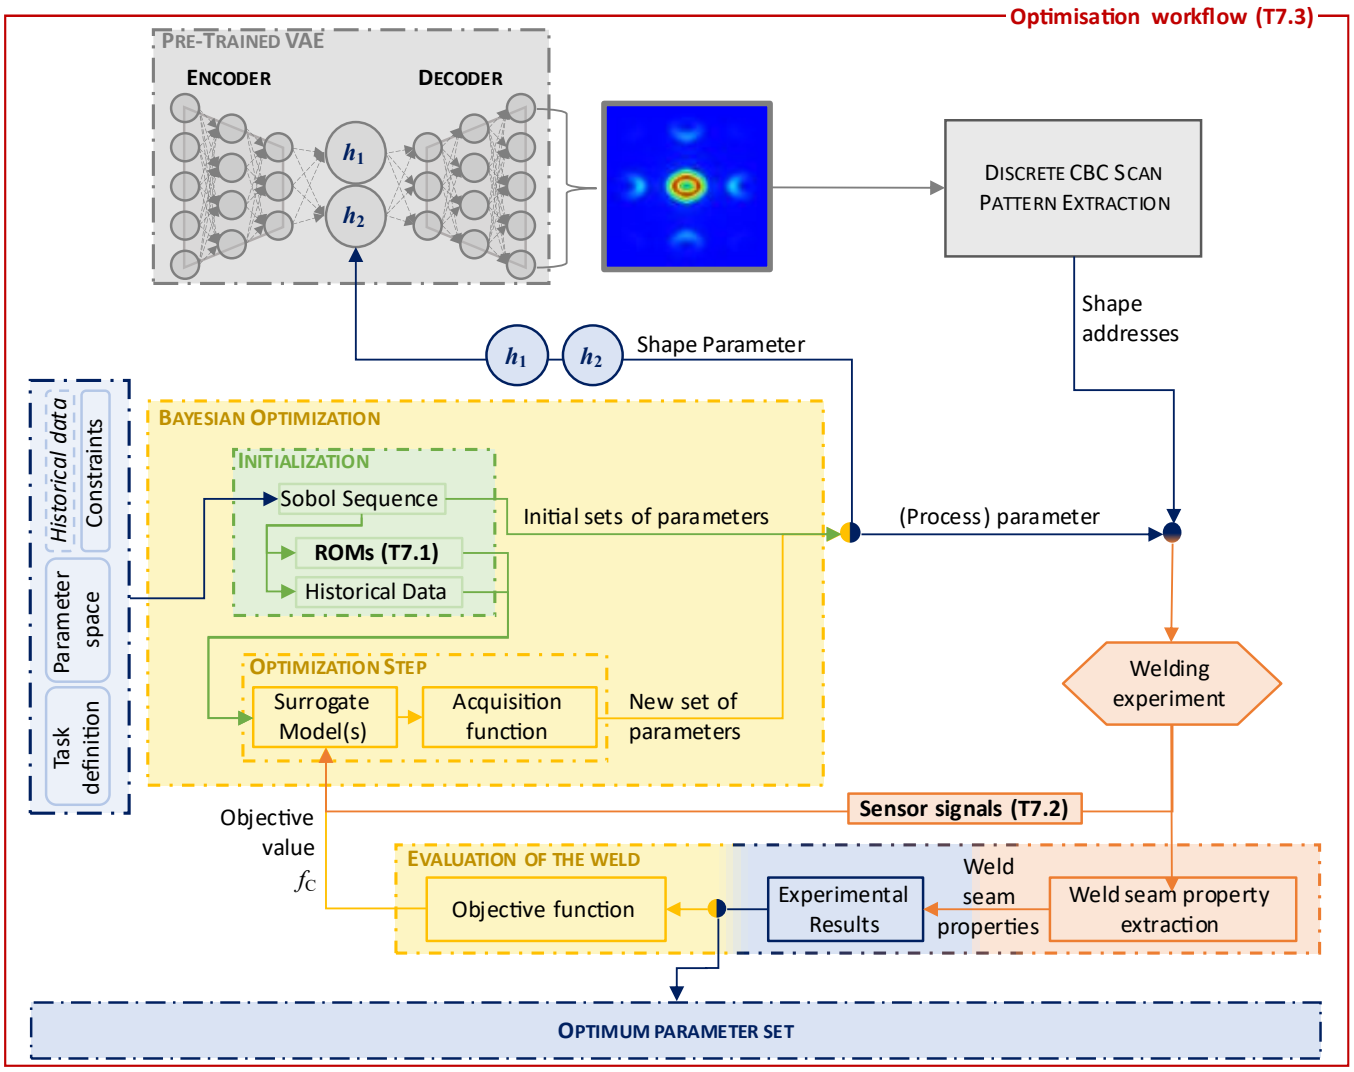

Fig. 16: Optimisation workflow (Task 7.3). The visualization is adapted from a forthcoming publication by the authors.

At the beginning of the optimisation, the task is defined, including the parameter space, admissible parameter ranges, constraints, and target criteria. To reduce the need for expert intervention, the framework provides structured workflows for task definition and objective construction, thereby supporting users with limited prior knowledge in optimisation or process modelling. These structured workflows for task definition and objective construction specifically address the requirement for low user expertise by reducing the need for manual parameter tuning and domain-specific knowledge.

The formulation of the objective function represents a critical component of the overall approach, as it defines the mathematical representation of the optimisation goal. Inadequate objective definitions may lead to technically convergent optimisation processes that fail to achieve the desired process outcome. In welding applications, the construction of suitable objective functions typically requires in-depth domain knowledge, as multiple target criteria must be combined and appropriately weighted.

To address this challenge, a structured workflow for objective function design was developed and implemented as a dedicated service within the platform architecture. This service generates a consistent and reproducible objective function definition based on the specified target criteria and provides it to the optimisation module as part of the task configuration. While the objective function is typically defined at the beginning of the optimisation process, it can be adapted during the experimental campaign if required. In this case, previously observed process results remain available, and the corresponding objective values are recomputed according to the updated definition. This approach directly addresses the challenge of multi-criteria process objectives by providing a systematic and reproducible way to combine multiple target criteria into a single optimisation objective.

To address the challenge of high-dimensional parameter spaces arising from advanced beam shaping, a dimensionality reduction step based on a Variational Autoencoder (VAE) is applied. The VAE maps high-dimensional beam representations to a compact latent space, which is then used as the input for the optimisation. This enables efficient optimisation while preserving the relevant structure of the process parameter space.

Based on the defined search space, an initial set of candidate parameter configurations is generated. In knowledge-free operation, this is achieved using a Sobol sequence to obtain a space-filling initial design. If prior knowledge is available, historical data and information derived from Reduced Order Models (ROMs) can be used to guide the initialisation. This flexible initialisation strategy directly addresses the varying availability of prior knowledge across different use cases, enabling the framework to operate both in data-scarce and data-rich scenarios.

The optimisation then proceeds in an iterative loop. Candidate evaluation is performed externally, for example through physical experiments or simulation-based workflows. The resulting weld seam properties are processed using the current objective function to obtain scalar objective values.

These observations are returned to the optimisation module and used to construct or update a Gaussian Process surrogate model. Since the optimisation module operates on exchanged data rather than maintaining a persistent internal model state, the surrogate model can be reconstructed from the available dataset at each iteration. This stateless design ensures consistency, particularly in cases where the objective function is modified during the optimisation process.

Based on the surrogate model, new candidate parameter configurations are selected via an acquisition function, balancing exploration of uncertain regions and exploitation of promising regions of the parameter space. This directly addresses the challenge of limited evaluation budgets, as Bayesian optimisation enables efficient parameter exploration with a minimal number of costly evaluations.

This iterative loop of candidate generation, evaluation, model construction, and selection continue until predefined stopping criteria are reached, such as convergence of the objective function or exhaustion of the evaluation budget.

The framework was implemented using the BoTorch library, employing Gaussian Process models with Automatic Relevance Determination kernels and logarithmic Expected Improvement as acquisition function. Where appropriate, logarithmic transformations are applied to variables spanning multiple orders of magnitude. To further support platform integration, the optimisation state is not stored internally but serialised and exchanged explicitly via the platform interface, enabling stateless interaction between optimisation module and execution environment.

The relationship between the identified challenges and design requirements (C1–C6) and the corresponding methodological solutions implemented within the WP7 framework is summarised in Table 1.

Table 1: Key challenges and methodological solutions

ID	Challenge / Requirement	Methodological Solution in WP7 Framework
C1	High-dimensional parameter spaces	Variational Autoencoder (VAE) for dimensionality reduction and latent space optimisation
C2	Limited evaluation budget	Bayesian optimisation using Gaussian Process surrogate models and acquisition functions
C3	Multi-criteria process objectives	Structured objective function design workflow (objective service) for scalarisation
C4	Varying availability of prior knowledge	Dual initialisation strategy (Sobol sampling and ROM / historical data integration)
C5	Low requirement for user expertise	Structured workflows for task definition and objective construction
C6	Integration into platform-based workflows	API-based architecture and stateless optimisation design

Results

As a result of Task 7.3, a first modular optimisation framework for fast process development in welding applications was successfully implemented and tested. The framework supports optimisation across different task definitions by separating task-specific elements from the optimisation core and by enabling both knowledge-free initialisation and the integration of prior knowledge (C4).

Baseline evaluation: Bayesian optimisation workflow

A first experimental campaign was conducted at University of Stuttgart using AlSi316L material to validate the core Bayesian optimisation workflow without dimensionality reduction. The optimisation was performed over a parameter space including laser power, feed rate, shape frequency and predefined beam shape parameters. Quality criteria were welding depth, pore area, seam elevation and undercuts.

Initialisation was carried out using a Sobol sequence with 10 parameter configurations. Based on this initial dataset, the Bayesian optimisation loop was applied. Within fewer than 40 additional iterations, the optimisation successfully identified a parameter configuration fulfilling the defined process requirements. This demonstrates that the proposed optimisation framework is capable of efficiently guiding parameter exploration under limited evaluation budgets (C2).

The results of this study form the basis of a conference proceeding to be presented at LANE 2026 in Fürth, Germany.

VAE-based parameterisation: Elliptical beam shapes

In a second study, a Variational Autoencoder (VAE) was introduced to investigate the use of latent representations for beam shape parameterisation. The VAE was trained exclusively on elliptical beam shapes. Due to this restricted training domain, no dimensionality reduction was achieved in the strict sense. However, the study demonstrated that latent parameters can be used to represent meaningful variations in beam geometry. A latent space traversal (see Fig. 17 - each tile represents the corresponding output or feature map for a specific latent variable combination, highlighting how spatial patterns evolve across the latent domain) shows that variations in beam shape are reflected by continuous changes in the latent space. The latent dimensions exhibit structured behaviour, with z[0] and z[1] controlling complementary aspects of the beam geometry. The traversal is performed by varying each latent dimension within the range of approximately [−3, 3], corresponding to the high-probability region of the assumed standard normal prior.

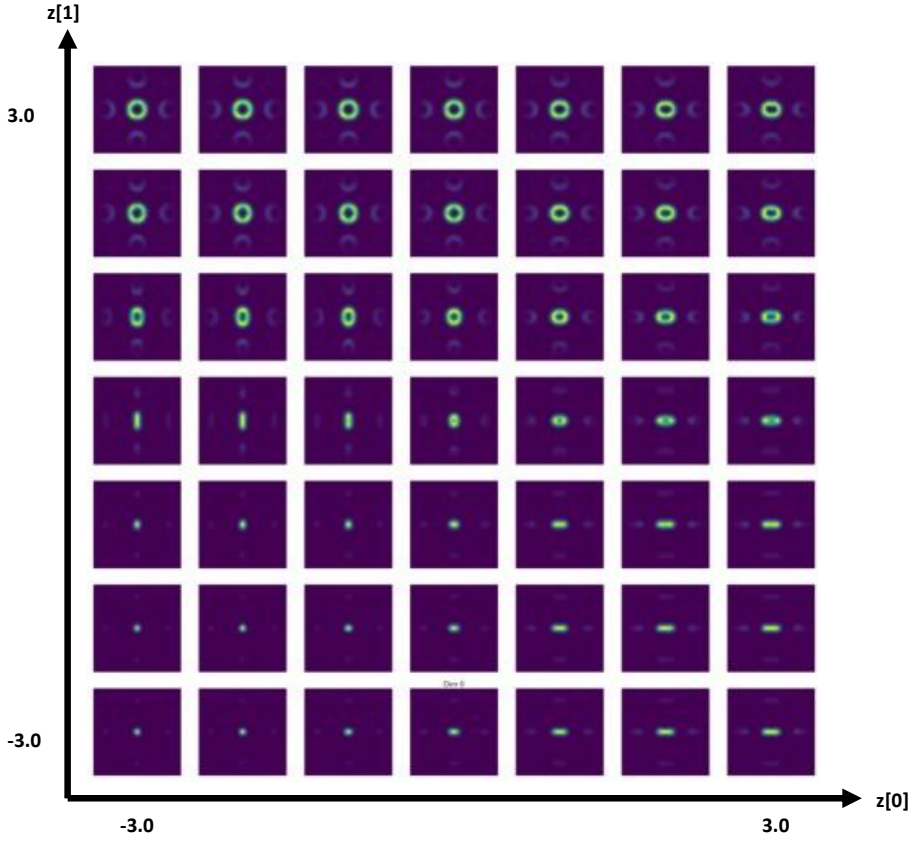

Fig. 17: Visualization of model responses over a two-dimensional latent space (z[0], z[1])

Similarly to the baseline evaluation, laser power, feed rate, shape frequency, and two latent variables were optimised to achieve a defined penetration depth while minimising pore area, undercuts, and seam elevation. Fig. 18 shows the evolution of the objective (cost) values over the course of the Bayesian optimisation. After an initial exploration phase (green shaded area), low cost values—corresponding to high weld quality—are reached repeatedly. The optimisation is thus able to identify parameter configurations that satisfy the target criteria.

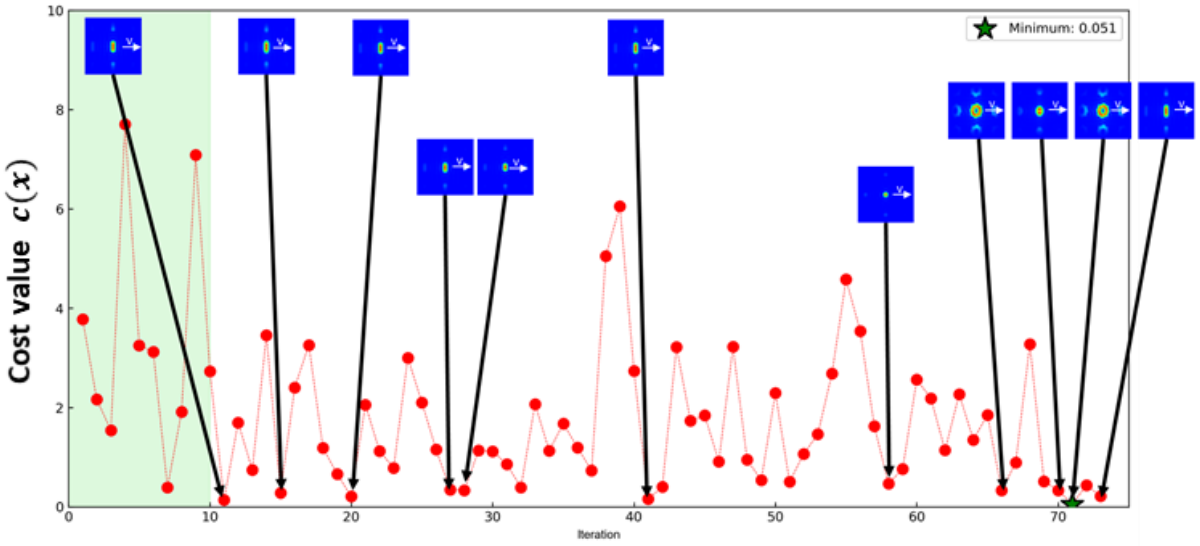

Fig. 18: Bayesian optimisation progress showing the evolution of cost values.

The corresponding parameter evolution and weld quality metrics are shown in Fig. 19. Initially, a broad range of parameter combinations is explored before the optimisation increasingly focuses on regions associated with lower cost values. Notably, these low-cost solutions are not restricted to a single beam configuration. Instead, multiple distinct combinations of process parameters and latent variables yield similarly favourable objective values. This indicates that the latent representation enables structured exploration of the shape space and supports the identification of multiple viable process solutions.

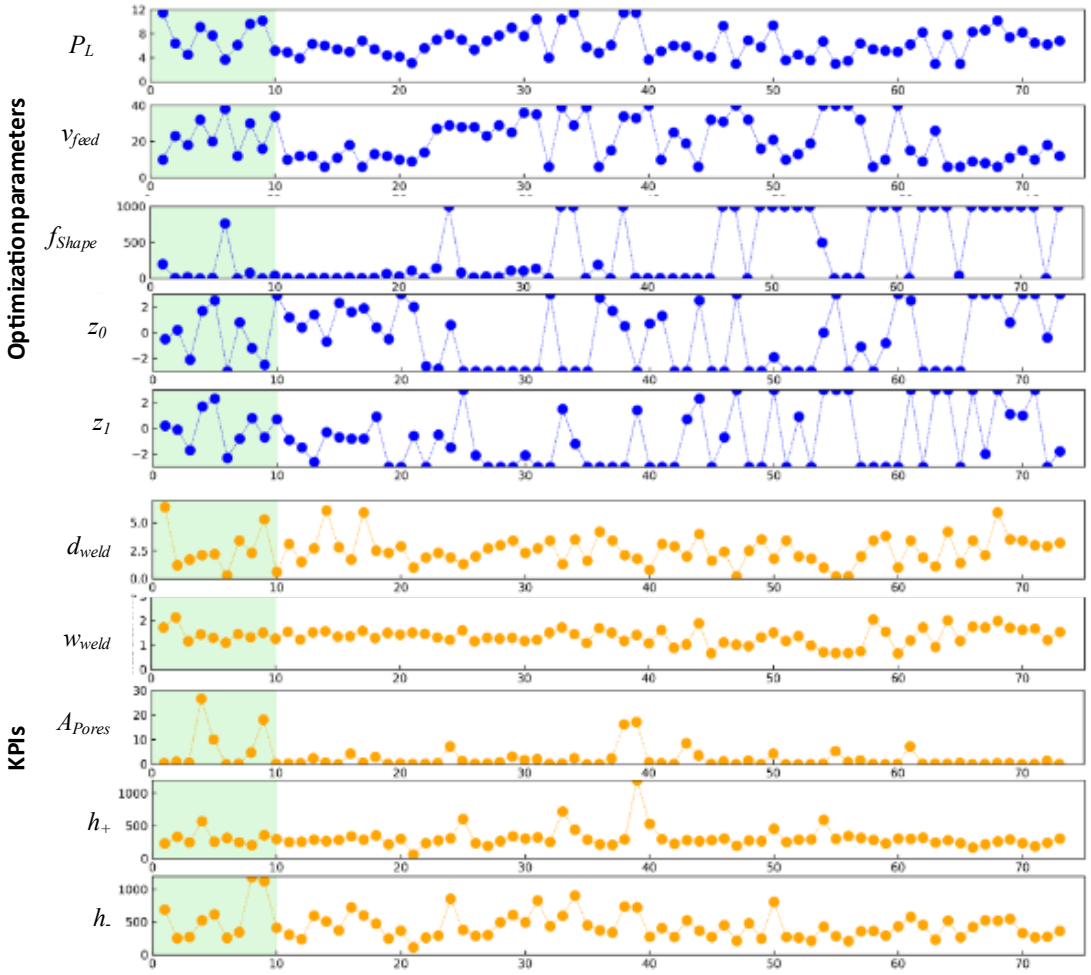

Fig. 19: Bayesian optimization progress showing parameter evolution, weld quality metrics, and cost values. Green shaded area: initialization phase.

These results indicate that VAE-based representations are suitable as a compact parameterisation approach, although their effectiveness strongly depends on the diversity and structure of the training data. The findings of this study are currently under review for publication and will be presented at International Congress on Applications of Lasers & Electro-Optics (ICALEO) 2026.

Extended representation: Towards dimensionality reduction

In a further step, the parameter space was extended to include both elliptical and rectangular beam shapes. This increased the variability of the shape representation and provided a more suitable basis for dimensionality reduction.

The combined representation demonstrated that latent parameterisations have the potential to significantly reduce the dimensionality of complex beam shaping problems, while preserving relevant geometric characteristics. These results highlight the importance of appropriate training strategies for the VAE and confirm its potential as a key component for addressing high-dimensional optimisation problems (C1).

Overall findings

Across all studies, the interaction between objective function definition, parameter representation, and Bayesian optimisation was successfully demonstrated. In particular:

- the optimisation framework enables efficient parameter exploration under limited evaluation budgets (C2),
- the structured objective function workflow provides consistent and reproducible evaluation of experimental results (C3, C5), and
- the modular and stateless architecture supports flexible integration into platform-based workflows (C6).

At the same time, the investigations highlighted several areas for further development, including improved VAE training strategies, more expressive beam-shape representations, and full integration into the platform feedback loop.

Key takeaway

Task 7.3 delivered and demonstrated a modular optimisation framework that enables efficient process parameter identification under limited evaluation budgets. Its key strengths lie in its adaptability across welding tasks, its ability to operate across different levels of prior knowledge, and its integration-ready architecture. Further training, validation and extension will be carried out within WP8.

6. Integration and Contribution to Project Objectives

WP7 contributes directly to project CONTROL objective, which targets real-time AI-based decision support through the integration of sensor data and digital twins.

The contributions of WP7 are structured according to its three main task areas:

- **TUW (T7.1 – Process synthesis):** The development of data-driven reduced-order models that provide rapid and accurate forward predictions, as well as forward-model-assisted inverse parameter identification. All of this can be combined into a tool that supports welding engineers and even enables autonomous, rapid decision-making.
- **WMG (T7.2 – Virtual sensing and process control):** The methodology provides a control-oriented benchmarking of modelling approaches for laser welding, identifying models that achieve an effective trade-off between predictive accuracy and computational efficiency. This establishes a modelling framework suitable for state estimation and predictive control of laser welding processes. Additional synergies are presented in D11.1, where integration with real-time data from IIoT sensors are presented, with a framework to seamlessly combine sensor data to FOM/ROM models for control actions, using standardised file protocols.
- **USTUTT (T7.3 – Fast process development):** Development of a modular optimisation framework for efficient parameter identification under varying levels of prior knowledge, integrating simulation data, experimental results, and optimisation strategies within a unified workflow.

Cumulatively, these contributions establish the methodological foundation of the AI-based Decision Module, which will be tested and validated in WP8, leading to integration into the Lasers4MaaS in WP15 and WP16.

7. Impact and Exploitation Potential

The work carried out in WP7 has generated significant scientific impact and strong exploitation potential for the Lasers4MaaS digital platform, particularly in the areas of reduced-order modelling, control-oriented modelling, AI-assisted process optimisation, process monitoring, and decision-support services for advanced laser welding.

In T7.1, TUW developed data-driven reduced-order models with significant impact and exploitation potential beyond the immediate project context. By drastically reducing the time required for process

development from lengthy and iterative experimental campaigns to near-instantaneous model-based predictions, the developed tools lower the barrier to adopting laser beam welding technologies. The ability to predict process outcomes and identify suitable input parameters on the fly makes the methodology directly exploitable as a core component of the Process Selector and Process Optimiser services within the Lasers4MaaS digital platform. These new decision-support capabilities also contribute to reducing the specialist knowledge barrier for laser welding, supporting the broader Lasers4MaaS vision of accessible, flexible, and decentralised manufacturing.

In T7.2, WMG developed the fundamental methodological framework for control-oriented modelling of laser welding. This work provides one of the first systematic evaluations of reduced-order modelling techniques for control-oriented laser welding applications. The framework combines reduced-order models with state-of-the-art PID controllers as well as next-generation reinforcement learning controllers, thereby creating a strong scientific foundation for adaptive and intelligent control strategies in laser welding. The results are planned to be presented at the ICALEO conference in October 2026.

From an exploitation perspective, the T7.2 work will support the Process Controller service of the Lasers4MaaS platform. End users will benefit from the ability to automatically generate control code ready for deployment on actual laser welding machines. This will reduce manual setup time and minimise trial-and-error approaches for controller calibration and tuning. Once deployed, the Process Controller will act as a "quality-keeping" tool under process uncertainty. The main application areas include adaptive laser welding systems in sectors such as aerospace, automotive, energy, medical devices, and food machinery, with dynamic beam shaping technology playing a central role. Beyond laser welding, the developed control-oriented modelling framework also has strong potential for transferability to other laser-based manufacturing processes, including Directed Energy Deposition.

In addition, the continuous collection and analysis of process data support predictive maintenance and remote monitoring strategies by enabling the identification of early indicators of component degradation, optical contamination, or machine instability before critical failures occur. The synergies between the Process Controller, Remote Monitoring, and Predictive Maintenance services are illustrated in Fig. 20 and further addressed in D11.1.

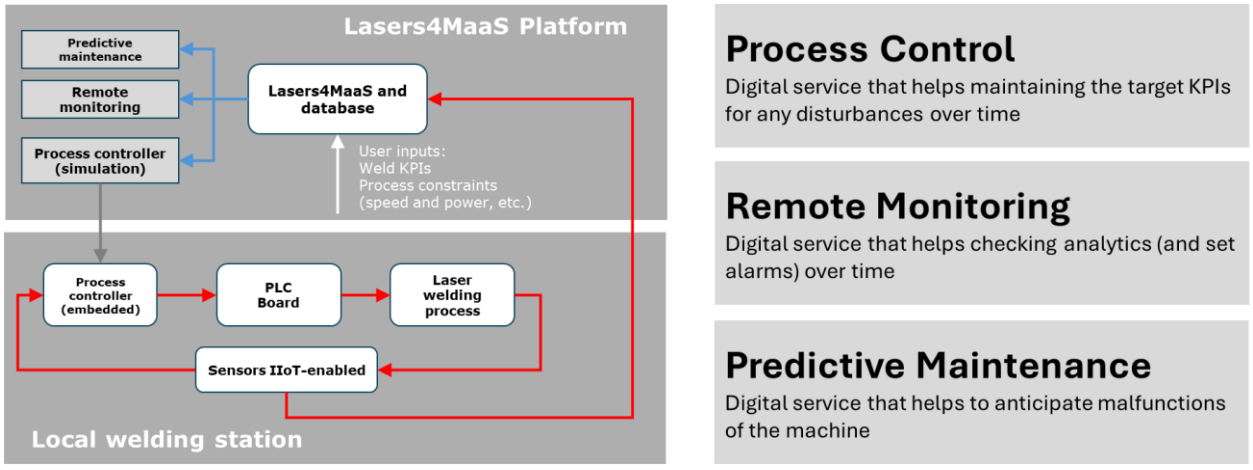

Fig. 20: Synergies between process controller, remote monitoring and predictive maintenance

In T7.3, USTUTT developed a novel methodology for the parametrisation of dynamic laser beam shaping in laser welding using Variational Autoencoder-based methods. To the best of the consortium's knowledge, no comparable approach currently exists that enables the compact and physically meaningful representation of highly dynamic beam shapes within a reduced latent space. The methodology therefore represents a significant advancement beyond the current state of the art in AI-assisted laser process optimisation.

The developed VAE parametrisation enables a systematic investigation of the relationship between beam shape evolution and relevant physical welding phenomena, including melt pool dynamics, keyhole stability, defect formation, and microstructure development. This provides an important scientific foundation for linking dynamic beam shaping strategies with process quality and weld performance. The methodology will be made accessible to the scientific community through peer-reviewed publication, with first results planned for presentation at the LANE conference in September 2026.

From an exploitation perspective, the VAE methodology forms a key enabling technology for the Lasers4MaaS digital platform and its AI-based decision-support architecture. Integrated with Bayesian optimisation and reduced-order models, the approach enables rapid multi-objective optimisation of welding processes, considering quality, productivity, energy efficiency, and defect minimisation simultaneously. Within the platform, the methodology will support the AI-based Decision Module and Process Optimiser services, enabling industrial users to automatically identify optimal beam shapes and process parameters with significantly reduced experimental effort and process development time. This exploitation pathway directly contributes to the Lasers4MaaS objectives related to controllability, prediction, and rapid process optimisation within WP7–WP10, while supporting the broader project vision of smart, decentralised, and sustainable Manufacturing-as-a-Service solutions for advanced laser welding applications.

Together, the developments in T7.1, T7.2, and T7.3 provide complementary methodological building blocks for the Lasers4MaaS digital platform. T7.1 enables fast data-driven prediction and process selection, T7.2 provides the basis for adaptive process control and quality assurance, and T7.3 introduces a compact AI-based parametrisation of dynamic beam shapes for advanced process optimisation. Combined, these contributions strengthen the platform's capabilities for prediction, optimisation, control, monitoring, and decision support. They reduce the need for specialist knowledge and extensive experimental campaigns, while enabling more accessible, flexible, decentralised, and sustainable Manufacturing-as-a-Service solutions for advanced laser welding applications.

8. Conclusions and Future Work

Overall, WP7 contributes to the development of explainable and interpretable AI/ML within Lasers4MaaS by combining physics-informed XAI/ML, reduced-order models, VAE-based beam shape parametrisation, Bayesian optimisation, process control, and parameter selection. These methods support transparent links between user-defined KPIs, process parameters, beam shaping strategies, process dynamics, sensor signals, and weld quality, thereby going beyond black-box optimisation and enabling more interpretable decision support.

WP7 is closely aligned with the parallel activities in WP9 and WP11. While WP9 focuses on physics-based digital twins and high-fidelity simulation models, WP7 uses such physics-based understanding to support process synthesis, optimisation, and control. Similarly, WP11 provides the sensor and monitoring perspective, including IIoT-enabled data acquisition, feature extraction, data labelling, fusion, and database integration, which are relevant for virtual sensing, process control, remote monitoring, and predictive maintenance. The detailed developments and results of WP9 and WP11 are reported separately.

The outputs of WP7 are further implemented, trained, calibrated, and tested in WP8, supporting their integration into platform-level smart services such as the Process Selector, Process Optimiser, AI-based Decision Module, and Process Controller. Together with the weld data generation and characterisation activities in WP12, these developments provide the basis for model calibration, optimisation, and validation under dynamic beam shaping conditions.

In this way, WP7 contributes to a coherent pathway from explainable AI/ML model development and physics-informed optimisation to industrially relevant digital services. These services support faster process development, improved parameter selection, adaptive process control, reduced defects, shorter development times, and more sustainable laser manufacturing. The integrated models, optimisation

strategies, controllers, and digital services will ultimately be validated in WP15 and demonstrated in relevant industrial environments in WP16 against the project KPIs.

References

[1] Zenz C, Buttazzoni M, Florian T, Crespo Armijos KE, Gómez Vázquez R, Liedl G, et al. A Compressible Multiphase Mass-of-Fluid Model for the Simulation of Laser-Based Manufacturing Processes. Computers & Fluids 2024;268:106109.

[2] Zenz C, Cook PS, Vörös L, Otto A. A Critical Comparison of One- and Two-Fluid Approaches for the Simulation of Laser-Induced Melt Pool Formation and Vaporisation. Discov Mater 2025;5:266.