



HYBRID AND ADVANCED MACHINE LEARNING PLATFORM FOR
MANUFACTURING, ENGINEERING, AND RESEARCH @ HLRS

D5.1

Training Activities and Portfolio v1

WP5: Training

Grant Agreement Number: 101234027



EuroHPC
Joint Undertaking



Baden-Württemberg
Ministry of Science,
Research and Arts



Bundesministerium
für Forschung, Technologie
und Raumfahrt

Bayerisches Staatsministerium für
Wissenschaft und Kunst



Niedersächsisches Ministerium
für Wissenschaft und Kultur



Co-funded by
the European Union

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101234027. This project is co-funded by the European Commission, the German Federal Ministry of Research, Technology and Space (BMFTR), the Baden-Württemberg Ministry of Science, Research and the Arts, the Bavarian State Ministry of Science and the Arts and the Lower Saxony Ministry of Science and Culture.

Document Information

Work Package	5	Training
Author(s):	Michael Walther	BADW-LRZ
	Nicolay Hammer	BADW-LRZ
	Junghwa Lee	USTUTT/HLRS
	Alexander Studt	KIT
	Till Riedel	KIT
	Kevin Lüdemann	GWDG
Approved by	MGT Board	24.06.2026
Reviewer 1	Hauke Kirchner	GWDG
Reviewer 2	Victoria-Luise Scheible	SICOS
Reviewer 3	Kai Binder	USTUTT/HLRS
Dissemination Level	PU – Public, fully open	

Co-Funded by the European Union. Views and opinions expressed are however those of the author only and do not necessarily reflect those of the European Union or the EuroHPC JU. Neither the European Union nor the granting authority can be held responsible for them.

Document History

Partner	Date	Comments	Version
BADW-LRZ	09.04.2026	Initial Version	0.1
BADW-LRZ	11.06.2026	Ready for first Internal Review	0.2
BADW-LRZ	23.06.2026	Ready for Management Review	1.0
BADW-LRZ	24.06.2026	Management Review Applied	1.1

Executive Summary

This deliverable presents the initial design and structuring of the HammerHAI training programme as part of the EuroHPC AI Factories initiative. The programme aims to support the adoption and effective use of artificial intelligence technologies by enabling users, particularly small and medium-sized enterprises and start-ups, to develop the skills required to access and apply advanced computing resources and services.

The current training catalogue combines structured modules, flexible and on-demand formats, as well as application-oriented activities such as hackathons. These elements are organised into core training tracks reflecting different user roles, levels of expertise, and application scenarios, ranging from foundational knowledge to advanced, infrastructure-oriented use cases. In addition to direct delivery through consortium partners, collaboration with intermediary organisations is explored to extend dissemination and connect users to suitable training and workshop offerings within the wider ecosystem.

Building on this catalogue, the analysis of user requirements identifies gaps between required capabilities and existing skills. The findings highlight the need for accessible entry points, practical and application-oriented training, and support for transitioning from local environments to complex computing infrastructures. At the same time, regulatory aspects and policy priorities play an increasing role, including the importance of AI literacy in line with Article 4 of the EU AI Act [1], as well as broader objectives reflected in the European AI strategy and the Apply AI strategy [2]. These considerations emphasise the need for training approaches that integrate technical, organisational, and compliance-related competencies.

The training programme is therefore designed as a modular and flexible component of the HammerHAI ecosystem, closely aligned with the service portfolio and evolving user journey. Continuous monitoring and feedback support its iterative refinement as infrastructure, services, and user engagement mature. The programme represents a snapshot of an evolving offering, providing a structured foundation for further extension and alignment with emerging needs while serving as an entry point into the broader HammerHAI ecosystem and facilitating the uptake of AI Factory services.

List of Abbreviations

AI	Artificial Intelligence
AIF	AI Factory
AIFA	AI Factory Antenna
BMFTR	Bundesministerium für Forschung, Technologie und Raumfahrt / Federal Ministry for Research, technology and space exploration
CC	Creative Commons
CQF	Competence and Qualification Framework
EDIH	European Digital Innovation Hub
EPCC	Edinburgh Parallel Computing Centre
EU AI Act	EU Artificial Intelligence Act (Regulation 2024/1689)
EuroHPC (JU)	European High Performance Computing Joint Undertaking
EVITA	EuroHPC Virtual Training Academy
GCS	Gauss Center for Supercomputing
GPU	Graphics Processing Unit
GSA	Gap- and Stakeholder Analysis
GWDG	Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen
HLRS	High-Performance Computing Center Stuttgart
HPC	High-Performance Computing
HPC CF	HPC Certification Forum
HPC SPECTRA	High-Performance Computing Skills Platform and European Collaboration for Training
KISSKI	KI-Servicezentrum für sensible und kritische Infrastruktur / AI Service Centre For Sensitive and Critical Infrastructure

KIT	Karlsruhe Institute of Technology
KPI	Key Performance Indicator
LLM	Large Language Model
LRZ	Leibniz Supercomputing Centre
MCML	Munich Center for Machine Learning
ML	Machine Learning
MLOps	Machine Learning Operations
MVP	Minimum Viable Product
NAS	Neural Architecture Search
NLP	Natural Language Processing
PAISE	Process Model for AI Systems Engineering
R&D	Research and Development
RAG	Retrieval-Augmented Generation
SICOS	SICOS BW GmbH
SLURM	Simple Linux Utility for Resource Management
SME	Small and Medium-sized Enterprise
UKAIFA	UK AI Factory Antenna
WP	Work Package

List of Tables

1	Overview of training modules across tracks	23
2	Indicative key performance indicators for the HammerHAI training programme.	48
3	Illustrative measured KPIs for a single course	49

List of Figures

1	Structural overview of the HammerHAI training portfolio.	13
2	Mapping of the core training tracks to their primary target audiences and the skill levels at which they typically operate.	17
3	Broad overview of the HammerHAI Service offer	18
4	An Example view of the HammerHAI training calendar	21
5	The course summaries listed on the HammerHAI webpage	22
6	An overview of the key players in the German AI ecosystem	37
7	Illustration of the four pillars of the AI Factory's talent development and start-up incubation strategy	42

Table of Contents

1	Introduction	10
1.1	Purpose of the Deliverable and Relation to Project Objectives . .	10
1.2	Target Audience	11
1.3	Structure of the Document	12
2	The HammerHAI Training Portfolio	13
2.1	Summary of the Training Portfolio	13
2.2	Core Training Tracks	14
2.2.1	AI Literacy, Ethics, and Administration	15
2.2.2	Research Data Analysis and Scientific Workflows	15
2.2.3	Deep Learning and Scalable AI Development	16
2.2.4	Specialised AI and Domain-Specific Applications	16
2.2.5	SME and Innovation-Oriented Pathway	16
2.2.6	Skill Levels and Progression	17
2.3	Available Training Modules ("Training Catalogue")	18
2.4	Flexible and On-Demand Training Options	26
2.4.1	Web-Seminars	26
2.4.2	Self-Paced and Video-Based Modules	27
2.4.3	Training-on-Demand Service	27
2.4.4	Brokered Training Opportunities	29
2.4.5	Workshop-Based Skill Development	30
2.4.6	Community-Based Learning and Journal Club	31
2.5	Challenge-Based and Intensive Programmes	31
2.5.1	Hackathons	32
2.5.2	Bootcamps	32
2.5.3	Mini-Hackathons and Short Innovation Formats	33

3 Agile Training Strategy and Roadmap Development	34
3.1 Strategic Context	34
3.2 Requirements Analysis	38
3.3 Skill Gap Findings	40
3.4 Structural Alignment of the Training Programme	42
3.5 Insights and KPIs	45
3.5.1 Qualitative Insights	45
3.5.2 Key Performance Indicators	46
3.6 Continuous Improvement and Roadmap Implementation of the Training	49
4 Conclusion	51
A Illustrative Snapshot of the Training Catalogue	54

1 Introduction

1.1 Purpose of the Deliverable and Relation to Project Objectives

Artificial Intelligence (AI) is increasingly becoming a key driver of innovation across private companies, the public sector, and research. At the same time, the adoption of AI technologies requires not only access to advanced infrastructures, but also the development of relevant skills and competences across a diverse range of user groups.

At the European level, strengthening AI capabilities is a central objective supporting digital transformation, industrial competitiveness, and technological sovereignty. In particular, European initiatives aim to enable broad access to advanced technologies while fostering the development of skills required to effectively apply them. Training our customers to acquire and improve AI-competencies is therefore a key enabler for translating infrastructure investments into tangible economic impact, especially for small and medium enterprises (SMEs) as well as start-ups. This ranges from foundational AI literacy to more advanced topics such as scalable model development and domain-specific applications, as reflected in the training tracks introduced in Section 2.2.

HammerHAI is part of this European AI ecosystem, contributing to the provision of AI services, infrastructure, and support mechanisms. Within this context, training plays a central role in enabling users to access and effectively utilise HammerHAI resources. At the same time, aligning training activities with the actual needs of companies, particularly of SMEs and start-ups, remains an ongoing process. Current efforts therefore focus on strengthening engagement with stakeholders from such enterprises to better understand their expectations and to adapt training offerings accordingly.

Initial steps in this direction are reflected in the design of upcoming formats such as the European AI Bootcamp and the European AI Hackathon introduced in Section 2.5, which are explicitly structured to address application-oriented use cases and to involve participants from SMEs, startups, and industry more directly.

In line with European recommendations, the training approach follows a targeted and demand-driven scope, focusing on areas where concrete needs and skill gaps are identified, rather than aiming to cover all possible AI-related topics. At the same time, this approach continues to evolve as additional insights are gathered through interaction with users and partners within our continuous Gap- and Stakeholder Analysis (GSA), see Section 3 for details.

Examples of such targeted training activities include selected modules from the catalogue described in Section 2.3, as well as specialised formats focusing on the practical use of HammerHAI services.

The training activities described in this deliverable are closely connected to the

broader objectives of the HammerHAI project and contribute to the work performed across multiple work packages. In this context, the training programme supports the uptake and effective use of project results by users, facilitates access to technical capabilities, and strengthens the overall impact of the HammerHAI initiative.

This deliverable focuses on the design and structuring of the HammerHAI training catalogue, providing a coherent framework for the development of AI-related skills ranging from foundational literacy to advanced, infrastructure-oriented expertise. At the same time, the training catalogue presented here represents an initial and evolving state of the programme. As both the HammerHAI infrastructure and the understanding of user needs continue to develop, the training offer will be further refined and expanded through ongoing interaction with stakeholders, monitoring of training activities, and iterative adaptation processes.

1.2 Target Audience

This deliverable targets both project reviewers and prospective users of the HammerHAI training programme. For reviewers, it provides a structured overview of the AI factory's training catalogue and its alignment with project objectives. For prospective users, particularly stakeholders from start-ups and SMEs, it offers orientation on available and planned training opportunities.

The HammerHAI training programme addresses a diverse set of users engaging with AI technologies in different roles and application contexts. The primary target audience comprises SMEs and start-ups, particularly in manufacturing and engineering sectors, seeking to explore and integrate AI into their products, services, and processes. Secondary audiences include larger enterprises, research facilities, and public-sector institutions addressing complex, large-scale challenges that can benefit from advanced AI capabilities.

At the level of individual participants, the programme primarily targets developers, researchers, and technical specialists involved in implementing AI solutions, while also addressing non-technical stakeholders such as decision-makers and administrative staff requiring a general understanding of AI capabilities and implications.

While these groups can be clearly identified, the specific needs of especially SMEs and start-ups remain an area of ongoing exploration. The training programme is therefore being developed continuously, combining existing expertise within the consortium with increasing engagement with external stakeholders, including both the target audiences and business support organizations such as the European Digital Innovation Hubs (EDIHs), as well as ongoing iterative refinement of the training portfolio.

This approach is also reflected in the design of application-oriented and collaborative formats, such as bootcamps and hackathons, which aim to strengthen

alignment with real-world use cases and practical needs.

1.3 Structure of the Document

This deliverable is structured as follows.

Section 2 provides an overview of the HammerHAI training programme, including the definition of core training tracks, skill levels, and target user groups. It further presents the current training catalogue, covering structured modules as well as flexible and challenge-based training formats. As such, this section is intended to serve as a point of reference for both reviewers and prospective users of the training programme.

Building on this overview, Section 3 provides the analytical and strategic foundation of the training programme. It examines user requirements and identified skill gaps, outlines the structural alignment of training activities with the HammerHAI service portfolio, and describes how monitoring, feedback, and iterative roadmapping processes are used to support continuous improvement. While this section is primarily intended to document the rationale and development approach for reviewers, it also offers contextual insights into how the training programme adapts to evolving user needs and system capabilities.

Additional supporting material, including course descriptions, is provided in Appendix A and Online.

Note: An updated version of this deliverable will be provided in Fall 2027 as Deliverable D5.2.




Figure 1: Structural overview of the HammerHAI training portfolio. The offer is organised into structured catalogue-based training, flexible and on-demand options, and challenge-based and intensive programmes, complemented by extended dissemination through external partners and the wider AI Factory ecosystem.

2 The HammerHAI Training Portfolio

2.1 Summary of the Training Portfolio

The HammerHAI training programme brings together a range of complementary formats to support users, particularly from SMEs and start-ups, in adopting and applying AI technologies across different contexts and levels of expertise while also enabling the development and continuous improvement of their skills. These formats include structured training modules, flexible and on-demand delivery mechanisms, and application-oriented activities that emphasise hands-on experience and collaboration.

The training portfolio is organised in a programme consisting of core training tracks that reflect different user roles, skill levels, and application scenarios (Section 2.2). Building on this structure, it comprises a catalogue of available and planned training modules contributed by the consortium (Section 2.3), including both foundational and specialised topics.

This catalogue is complemented by additional formats that increase flexibility and practical relevance. These include Training-on-Demand services and self-paced learning options (Section 2.4), as well as challenge-based and intensive programmes such as hackathons and bootcamps (Section 2.5).

In addition to the direct provision of training through consortium partners, the

programme is expected to leverage external dissemination channels to increase its reach. In particular, closer integration with intermediary organisations such as EDIHs is being explored, enabling these organisations to offer HammerHAI training opportunities to their own user base and thereby extending the programme's distribution network beyond its current direct access points.

Together, these elements form a modular and evolving training offer that is designed to accommodate diverse user needs, support the development and application of AI in real-world settings, and extend its impact through both direct provision and collaboration with external dissemination networks. Figure 1 provides a structural overview of the training portfolio, summarising the main categories of offerings and their constituent formats.

2.2 Core Training Tracks

The HammerHAI training catalogue is structured around a set of core training tracks that reflect the main user groups of the HammerHAI facility and the skills required to effectively adopt, use, and scale AI models on our cloud-like, AI-optimized supercomputer. The training programme is designed as a unified offering serving both research-driven and industry-oriented users, allowing flexible entry points and progression paths across different levels of expertise and application contexts.

The core training tracks can be combined and navigated flexibly to address a wide range of organisational settings, including SMEs, start-ups, and academic institutions. While the underlying technical skill requirements are largely shared, different user groups may prioritise specific combinations of tracks depending on their roles, objectives, operational constraints, and levels of AI maturity. Together, the tracks provide a coherent orientation within the training catalogue while supporting modular learning journeys tailored to individual and organisational needs.

The training tracks address three primary target audiences:

1. Administrators and non-technical staff,
2. Researchers and domain scientists,
3. Developers, research software engineers, and technical specialists.

Each track groups complementary courses and activities offered by the HammerHAI consortium partners and serves as an entry point into the wider training catalogue, including regular courses, flexible on-demand formats, and intensive challenge-based programmes.

2.2.1 AI Literacy, Ethics, and Administration

This track focuses on building a shared understanding of AI fundamentals for non-technical users, administrative staff, and decision-makers who interact with AI systems in organisational or professional contexts, including public-sector, enterprise, start-up, and SME environments. It prioritizes conceptual understanding, responsible use, and regulatory awareness rather than software development, drawing on established notions of AI literacy [3] and explicitly supports compliance with Article 4 of the EU AI Act on AI literacy [1].

Training within this track particularly supports participants in:

- understanding core AI concepts and terminology,
- applying AI tools responsibly in everyday workflows,
- recognising ethical, legal, and governance implications, including requirements stemming from the EU AI Act,
- developing basic interaction skills with AI systems in line with AI literacy and compliance expectations.

Representative courses include AI ethics seminars offered by HLRS (see Section 2.4.1), administrative-focused introductions such as *KI in der Verwaltung*, and AI competence and literacy trainings provided by GWDG, partly developed in the context of the AI Service Centre For Sensitive and Critical Infrastructure (KISSKI)¹. Elements of this track are also reflected in selected flexible and on-demand offerings described in Section 2.4.

2.2.2 Research Data Analysis and Scientific Workflows

This track is aimed at data analysts, researchers, and domain scientists who need to analyse data at scale and implement reproducible scientific workflows using HPC and AI-enabled tools. It bridges traditional HPC usage with modern data-driven and AI-supported research practices applicable to both academic research and industrial R&D contexts.

Participants following this track are supported in:

- navigating HPC environments and batch systems,
- applying container technologies for reproducible research,
- performing large-scale data analysis,
- integrating AI methods into scientific workflows.

Courses contributing to this track include *Data Analysis with HPC*, *Natural Language Processing*, *Time Series Forecasting*, *Data Analytics for Engineering Data using Machine Learning*, as well as *Introductions to Container Technologies and AI Systems*.

¹<https://kisski.gwdg.de/>

2.2.3 Deep Learning and Scalable AI Development

This track targets developers, data scientists, and research software engineers building and scaling AI models on GPU-accelerated and HPC infrastructures. It focuses on the efficient use of accelerators, parallelisation strategies, performance optimisation, and the scaling of AI workloads from local or single-GPU setups to multi-GPU and multi-node environments suitable for advanced prototyping.

Training in this track supports participants in:

- developing and training deep-learning models on GPU-accelerated systems,
- managing and allocating GPU and accelerator resources,
- applying data-parallel and model-parallel training techniques,
- analysing and optimising the performance and scalability of AI workloads.

Relevant courses include *Deep Learning with GPU Cores* and *Multi-GPU Deep Learning*, advanced modules like *Fundamentals of Deep Learning*, *Data Parallelism* and *Model Parallelism*, as well as performance-oriented offerings such as *Performance Engineering Tools for AI and HPC Workloads* and *Performance Engineering in ML Workloads*.

2.2.4 Specialised AI and Domain-Specific Applications

This track addresses advanced and specialised application areas where AI techniques are tailored to specific domains or require deeper optimisation for HPC infrastructures. It is intended for experienced users who want to apply AI methods beyond generic workflows.

Training within this track enables participants to:

- apply AI techniques to domain-specific and sector-oriented problems,
- work with advanced architectures such as large language models,
- analyse complex temporal and structured data,
- profile and optimise AI software for performance and scalability under real-world constraints,
- engage with current developments and emerging research directions in AI through interactive discussion and critical exchange of recent scientific publications.

Courses contributing to this track include *Natural Language Processing*, *Time Series and Data Analysis*, *Profiling AI Software Bootcamp*, as well as introductory offerings in emerging areas such as neuromorphic computing.

2.2.5 SME and Innovation-Oriented Pathway

This pathway emphasizes how start-ups, SMEs, and innovation-driven users typically engage with the HammerHAI training catalogue by combining elements

Core training track (primary target audience)	Beginner	Intermediate	Advanced
1. AI Literacy, Ethics & Administration <i>Admin & non-technical staff, decision-makers</i>	●		
2. Research Data Analysis & Scientific Workflows <i>Researchers & domain scientists</i>	●	●	
3. Deep Learning & Scalable AI Development <i>Developers, RSEs & technical specialists</i>		●	●
4. Specialised AI & Domain-Specific Applications <i>Experienced / advanced users</i>			●
5. SME & Innovation-Oriented Pathway <i>Start-ups, SMEs (cross-cutting)</i>	●	●	●

● = track typically operates at this skill level (Beginner → Intermediate → Advanced progression)

Figure 2: Mapping of the core training tracks to their primary target audiences and the skill levels at which they typically operate. The SME and innovation-oriented pathway is cross-cutting, drawing on elements from all tracks and skill levels.

from multiple core training tracks. Rather than representing a separate programme, it reflects a characteristic usage pattern focused on rapid adoption, applied development, and scalable deployment of AI technologies.

Training activities relevant to this pathway emphasise:

- effective use of AI tools for business-driven and applied research scenarios,
- rapid prototyping and MVP development,
- cost-aware and scalable deployment strategies,
- transition from experimentation to production-ready solutions.

Participants following this pathway typically draw on training elements from AI literacy and ethics, data analysis and workflows, scalable AI development, and domain-specific applications, complemented by flexible on-demand formats, and selected workshops (see Section 2.4) as well as challenge-based programmes such as hackathons and bootcamps described in Section 2.5.

2.2.6 Skill Levels and Progression

Across all training tracks, the HammerHAI catalogue supports progression across multiple skill levels. This layered design is informed by established digital competence frameworks such as DigComp [4] and the European e-Competence Framework [5], as well as by pedagogical principles of constructive alignment




Figure 3: Broad overview of the HammerHAI Service offer structured by maturity. For trainings, two adjacent maturity levels have been roughly merged to form a training skill level. (Credit: Kerem Kayabay)

and learning-outcome taxonomies [6, 7]:

- **Beginner:** Foundational knowledge, AI literacy, basic Linux and tooling, responsible and trustworthy AI.
- **Intermediate:** Applied skills including container usage, cluster interaction, MLOps concepts, and workflow implementation.
- **Advanced:** Optimisation and scaling, including multi-node training, performance profiling, and advanced architectures.

This layered structure allows participants to enter the training programme at an appropriate level and to expand their skills as their experience, project needs, and AI maturity evolve. Figure 2 summarises how the core training tracks introduced above map onto these skill levels and their primary target audiences.

The three-level model used in the training catalogue represents a simplified abstraction of the more fine-grained maturity model applied within the HammerHAI service framework (see Figure 3). Each training level corresponds to a grouping of two adjacent service maturity levels, enabling a clearer and more accessible representation while maintaining alignment with the overall project structure.

2.3 Available Training Modules ("Training Catalogue")

The HammerHAI training catalogue comprises a set of training modules contributed by the consortium partners. In Table 1, we provide an overview of the currently available and planned training modules, grouped according to the core training tracks and skill levels described in Section 2.2.

Users are invited to combine modules across different tracks and skill levels

to form tailored learning pathways aligned with their specific objectives and organisational context.

The table represents an excerpt of the overall training offer. More detailed course descriptions are available in Appendix A and through the respective providers, an on-line listing of the full catalogue can be found at <https://www.hammerhai.eu/training-catalog/>².

The listed modules are delivered in a variety of formats, including lectures/seminars, hands-on sessions, and hybrid training formats, and may be offered either as recurring modules or as part of scheduled training series and events. Detailed information regarding scheduling, access conditions, and registration procedures is provided through the respective training platforms of the consortium partners, as well as via the central HammerHAI training calendar, which serves as the primary hub for upcoming training activities³. An example snapshot of the training calendar is shown in Figure 4, the corresponding course summaries listed below the calendar are shown in Figure 5. In addition, updates on training activities are planned to be communicated through a dedicated training newsletter, alongside existing communication channels provided by individual consortium partners.

The training modules are delivered by expert instructors from the HammerHAI consortium partners, supported by AI specialists contributing through user support activities and application-oriented formats. Where appropriate, courses may also be provided by or jointly delivered with external contributors, further extending the scope of the training offer.

The catalogue⁴ is continuously evolving, with additional modules being developed and existing offerings being updated based on user feedback, emerging technology trends, and identified skill gaps. Courses are primarily offered in English, with selected courses also available in German, as indicated by the language of the course title. Where appropriate, individual courses may be provided in either language.

In addition, frequent coordination across consortium partners has supported the development of a coherent training portfolio by building on pre-existing course offerings and complementary areas of expertise. An initial consortium-wide inventory confirmed that the combined training programmes provide broad coverage of relevant topics with limited overlap, forming a strong foundation for synergies across partners.

At the same time, further coordination is required to strengthen cross-partner scheduling and to support seamless learning pathways that allow participants to combine modules from different providers. While training activities within individual centres are well established, the coordination of timing and progression

²German page at <https://www.hammerhai.eu/de/trainingskatalog/>

³<https://www.hammerhai.eu/events/>

⁴see for the most up-to-date version

across the consortium remains an area of ongoing development, particularly in light of the integration of training activities into multiple parallel initiatives.

In this context, flexible and on-demand training (see Section 2.4) formats play an important role, enabling more adaptive and user-driven participation independent of fixed schedules. These formats support the bridging of temporal gaps between structured training modules and allow participants to engage with relevant content in a more continuous manner, complementing the further development of coordinated cross-partner learning pathways.

← September 2026

Today

November 2026 →

October 2026

October 2026

MONDAY	TUESDAY	WEDNESDAY	THURSDAY	FRIDAY	SATURDAY	SUNDAY
			1	2	3	4
			Supercomputing Academy: Deployable Data Analysis & AI Pipelines with HPC			
5	6	7	8	9	10	11
Supercomputing Academy: Deployable Data Analysis & AI Pipelines with HPC						
	European AI Hackathon					
12	13	14	15	16	17	18
European AI Hackathon						
19	20	21	22	23	24	25
European AI Hackathon						
26	27	28	29	30	31	
European AI Hackathon						

Figure 4: An Example view of the HammerHAI training calendar, providing access to upcoming training activities across all training tracks, course summaries are listed below the calendar (see Figure 5). The calendar is typically updated every night.

Monday September 7 00:00 — Friday October 9 12:00	Supercomputing Academy: Deployable Data Analysis & AI Pipelines with HPC Contact: Höchstleistungsrechenzentrum Stuttgart (HLRS) https://www.hlrs.de/ This course is designed for practitioners in SMEs, startups, and applied research teams who need to turn data into actionable insights and repeatable, deployable ML workflows using high-performance computing (HPC) resources. Participants learn practical patterns to build reproducible and scalable analytics and ML pipelines that remain reliable under real-world constraints, including messy data, growing volumes, performance bottlenecks, and execution across heterogeneous compute environments. ▼ Read more
Tuesday October 6 10:00 — Thursday October 29 17:00	European AI Hackathon Contact: Leibniz-Rechenzentrum (LRZ) http://www.lrz.de/ Open Hackathons are multi-day, hands-on coding events that help developers and researchers accelerate, optimize, and scale their applications on a variety of data center architectures. Attendees work with dedicated mentors to identify and address data loading bottlenecks, parallelize training, and optimize AI pipelines for faster experimentation to realize performance gains and speedups. ▼ Read more

Figure 5: The course summaries listed on the HammerHAI webpage below the training calendar. Full course descriptions are available via the "Read more" buttons.

Table 1: Overview of training modules across tracks

Legend:

Current Providers: ^H HLRS ^L LRZ ^G GWDG ^K KIT

^V Video-recording available for self-paced learning

* Planned module

Skill levels:

Beginner

Intermediate

Advanced

Track	Course Title
AI Literacy, Ethics, and Administration	AI Competence Training^{G,V} Effectively Utilize AI Tools in Research^{G,V} KI Kompetenz Training^{G,V} KI in der Verwaltung: Eine Einführung in die Nutzung für alle Mitarbeiter*innen^{G,V} KI-Kompetenz für Unternehmen 1: Grundlagen und Begriffe der KI^K KI-Kompetenz für Unternehmen 2: KI-VO erklärt^K KI-Kompetenz für Unternehmen 3: Hands-on generative KI und ihre Risiken^K AI Ethics Lecture: Artificial Intelligence and Ethics^H AI Ethics Lecture: Privacy is not Power, but what is?^H AI Ethics Lecture: Technical Responsibility and Ethical Methods^H Agentic AI Web-Seminar (jointly with UKAIFA) *
Research Data Analysis and Scientific Workflows	Introduction to the AI Systems^L Basic Linux, Ethics & Security Training for HammerHAI Users* Data Analysis with HPC^H Data Analytics for Engineering Data using Machine Learning (in collaboration with Fraunhofer)^H

Table 1: Overview of training modules across tracks (continued)

Track	Course Title
	Data Management Concepts for Efficient and User-Friendly HPC^G Data Management Course[*] Introduction to Container Technology & Application to AI Systems^L Time Series Forecasting^H Using Data Pools for Scientific Data Sharing^G Data management for HammerHAI using piveau[*] Deployable Data Analysis & AI Pipelines with HPC[*] Machine Learning Operations at the HammerHAI infrastructure[*]
Deep Learning and Scalable AI Development	Fundamentals of Deep Learning^L Deep Learning with GPU Cores^G MLOps Training tailored to HammerHAI[*] Data Parallelism – How to Train Deep Learning Models on Multiple GPUs^L Deep Learning Bootcamp: Building and Deploying AI Models^G Model Parallelism – Building and Deploying Large Neural Networks^L Multi-GPU Deep Learning^{H,*} Performance Engineering in ML Workloads Using Score-P and Vampir^G Performance Engineering Tools for AI and HPC Workloads^G Profiling AI Software Bootcamp^{H,L} Scaling Deep Learning – From Single GPU to Clusters^{L,*}
Specialised AI and Domain-Specific Applications	European AI Hackathon^{H,L,*} Introduction to AlphaFold^G

Table 1: Overview of training modules across tracks (continued)

Track	Course Title
SME and Innovation-Oriented Pathway	Natural Language Processing ^H
	PaperHAI Journal Club ^K
	AI for Science Bootcamp ^{H,L}
	European Agentic AI Bootcamp ^{H,L,*}
	Introduction to Neuromorphic Computing and the SpiNNaker Hardware Platform ^G
	Effectively Utilize AI Tools in Research / Business ^{G,V}
	European AI Hackathon ^{H,L,*}
	Introduction to Container Technology & Application to AI Systems ^L
	Selected modules in Natural Language Processing and data analytics ^H
	hub:disrupt Hackathon: KI-Fabrik der Zukunft ^{L,K,*}
	X-ATHON – Künstliche Intelligenz praktisch erlebbar machen. ^{H,L,*}
	European Agentic AI Bootcamp ^{H,L,*}

Note: For this track, please also particularly consider our further services such as Training-on-Demand, brokered trainings, Web-Seminars, and workshop-based formats.

2.4 Flexible and On-Demand Training Options

While the portfolio presented above focuses on structured training modules, the HammerHAI training programme also includes a range of flexible and on-demand learning options that enable more targeted and self-paced engagement with specific topics.

These offerings complement the core training modules by supporting different learning preferences, time constraints, and use cases, and provide opportunities for short-term upskilling, hands-on experimentation, and focused exploration of specific tools, methods, or application domains.

They may take the form of self-paced learning units, guided exercises, applied project modules, or interactive sessions, and can be used both as entry points into the training programme and as complements to the structured modules introduced in Section 2.3. Designed to be modular and accessible, they allow participants to integrate specific learning elements into individual pathways as needed.

2.4.1 Web-Seminars

Web seminars constitute an additional flexible training format within the HammerHAI programme, enabling broad and low-threshold access to relevant topics across different user groups. They primarily target a wide audience, including industry stakeholders, SMEs, and interested newcomers, and are particularly suited for awareness-building and the communication of emerging developments rather than in-depth skill development.

Within the consortium, SICOS contributes extensive experience in organising and delivering web seminar series in the domains of high-performance computing and artificial intelligence, covering a broad range of application-oriented topics. These established series serve as a foundation and can be leveraged to extend the reach of HammerHAI-related topics through existing distribution channels and audiences.

In addition, a series of web seminars focusing on AI ethics is organised at HLRS (e.g. *Artificial Intelligence and Ethics*, *Technical Responsibility and Ethical Methods*, *Privacy and AI*). These activities complement the broader training portfolio by addressing aspects of responsible and trustworthy artificial intelligence, in line with increasing regulatory and societal requirements.

Building on this, additional web seminar activities are being developed in collaboration with partners inside and outside the consortium. Planned formats include topic-specific sessions such as a web seminar on agentic artificial intelligence, to be carried out jointly by HammerHAI and the UK AI Factory Antenna (UKAIFA) at the Edinburgh Parallel Computing Centre (EPCC)⁵. These sessions are expected to combine both theoretical and practical hands-on aspects of the subject.

⁵<https://www.epcc.ed.ac.uk/>

Overall, web seminar series enhance both the reach and visibility of the training programme. They act as an accessible entry point into the HammerHAI ecosystem and can serve as a first step towards more structured and application-oriented training formats.

2.4.2 Self-Paced and Video-Based Modules

A subset of the training modules described in Section 2.3 is also available as video-based, self-paced learning content, as indicated by the corresponding tag in the module listings.

These offerings allow participants to access core material asynchronously and, where applicable, to complement structured training activities with flexible learning options tailored to their individual schedules, including opportunities for repetition and independent study.

Where appropriate, video-based modules may be accompanied by assessment or evaluation mechanisms, allowing participants to validate their understanding and, in some cases, obtain confirmation of completed learning activities within the wider training programme.

Self-paced learning can, for example, be supported by existing online training platforms such as the High-Performance Computing Skills Platform and European Collaboration for Training (HPC SPECTRA) [8]⁶ and its successor the EuroHPC Virtual Training Academy (EVITA) [9]⁷. EVITA provides modularised training content and structured learning pathways in HPC based on the Competence and Qualification Framework (CQF)⁸. The core of the CQF is the Skill Tree⁹ designed and maintained by the HPC Certification Forum (HPC CF) [10]¹⁰ and it already contains AI-based skills, which can also be applied to courses in HammerHAI. HammerHAI training materials are designed to be compatible with such modular learning environments, allowing for integration into existing and future platforms where appropriate. In-kind contributions to the project and the HPC in Europe Portal¹¹ are currently being discussed from both sides. The modules that the EVITA project currently provides and will provide in the future are freely available for HammerHAI to be used.

2.4.3 Training-on-Demand Service

In addition to scheduled training activities, the HammerHAI programme is extending its offer through Training-on-Demand services. These enable users to

⁶<https://www.hpc-spectra.eu/>

⁷<https://www.evitahpc.eu/>

⁸<https://www.evitahpc.eu/cqf/>

⁹<https://www.evitahpc.eu/skill-tree/>

¹⁰<https://www.hpc-certification.org/>

¹¹<https://hpc-portal.eu/training>

request or select training modules that have been pre-developed and can be delivered in a flexible manner based on demand.

This format is particularly suited for organisations, particularly SMEs and start-ups, with specific or time-sensitive training needs. It allows users to access specialised content in a coordinated way, either as standalone training offers or in combination with existing curricula.

Typical Training-on-Demand offerings include:

- **Efficient and robust neural network training:** *Regularisation of Neural Networks* for Improved Efficiency and Robustness in Training and Inference and *Efficient Hyperparameter Tuning and Neural Architecture Search*, focusing on optimisation of neural network models under practical constraints
- **Image Generation:** From Generative Adversarial Networks, over Variational Autoencoder to Diffusion and Flow Matching, how do modern frameworks generate images
- **Bio-inspired and neuromorphic AI:** *Spiking Neural Networks: Principles, Training Methods, and Applications*, introducing event-driven computing paradigms and their potential for energy-efficient AI systems.
- **AI governance and regulatory compliance:** *AI Act AI Literacy: Understanding and Testing AI Literacy Requirements (Article 4)* and *Trustworthy AI: Risk Management for AI and Smart Data Projects*
- **Integration of AI into development processes:** *AI Systems Engineering using the PAISE Process Model* [11], focusing on embedding AI development into existing engineering workflows

The Training-on-Demand service¹² complements the structured training catalogue by introducing a flexible, demand-driven delivery mechanism for specialised, application-oriented training. While the catalogue provides a planned portfolio of recurring and predefined courses, the Training-on-Demand service enables organisations to request targeted training interventions that are directly aligned with their specific technical, sectoral, or organisational needs.

The service builds on existing training assets developed within HammerHAI, its partner institutions, and where relevant, other AI Factory initiatives. This reuse-first approach ensures efficiency, consistency, and alignment across the European AI Factory ecosystem. Where suitable material already exists, it is adapted and delivered in a customised format to the local target audience of HammerHAI (this might include translating the material in some cases). In cases where gaps are identified, HammerHAI may develop new training content on demand. Such newly created material is produced under a Creative Commons licence, ensuring that it can be reused and shared across partners, AI Factories, and the wider ecosystem, thereby contributing to a common European training

¹²Additional up-to-date information can be found at <https://www.hammerhai.eu/2026/06/16/on-demand-training/> (for the German version see <https://www.hammerhai.eu/de/2026/06/16/on-demand-training/>)

knowledge base. In these cases, the service fee is primarily associated with the design effort and the execution of the training delivery.

To operationalise this model, HammerHAI is developing an extended training catalogue that provides a comprehensive overview of available training modules, including not only scheduled offerings, but also modular components that can be combined into tailored training programmes. This catalogue serves as an enabling layer for matching demand with existing content and for identifying where bespoke development is required. It is designed to support dynamic and heterogeneous demand patterns across industries and stakeholder groups, particularly in contexts where standardised training formats are insufficient.

The Training-on-Demand service is designed to address two primary customer segments. The first consists of companies—particularly SMEs, mid-sized enterprises, and industrial organisations—that require training for entire teams or departments. These engagements are typically embedded in broader upskilling or AI adoption initiatives and may be delivered on a commercial basis or, where applicable, through de minimis-supported funding schemes. In these cases, training is positioned as an operational enabler of AI adoption, directly supporting the integration of AI Factory services into organisational workflows and processes.

The second target group comprises multipliers such as industry associations, cluster organisations, innovation agencies, educational institutions, and other ecosystem intermediaries. These actors often seek to extend or enrich their own training portfolios by integrating HammerHAI content into their programmes. In cases where training is delivered openly to support the broader objectives of the AI Factory—such as awareness raising, ecosystem development, or user acquisition—courses may be offered free of charge to open audiences of potential users. In such scenarios, Training-on-Demand functions as a strategic outreach mechanism that supports ecosystem growth and engagement. Alternatively, where training is delivered in a closed or value-added format (e.g., exclusively to members or as a resold service offering), it is provided on a commercial basis covering delivery and operational costs.

Overall, the Training-on-Demand service strengthens the responsiveness of HammerHAI's training portfolio by enabling rapid adaptation to evolving user needs, while ensuring alignment with the broader EuroHPC AI Factory ecosystem. It complements the structured catalogue by adding flexibility, deep customisation capabilities, and a scalable mechanism for transferring knowledge across sectors and organisations.

2.4.4 Brokered Training Opportunities

To further extend the available training offer, HammerHAI also facilitates access to external training opportunities through brokerage. This enables users to engage with additional courses provided outside the core catalogue, complementing the internal training portfolio where appropriate.

This mechanism is particularly relevant for organisations, particularly SMEs and start-ups, with specific or evolving training needs that go beyond the current catalogue. It allows access to specialised expertise and up-to-date content in areas such as emerging technologies, domain-specific applications, or strategic adoption of AI.

Typical areas addressed through brokered training opportunities include:

- specialised tools, platforms, or vendor-specific technologies
- advanced or niche AI applications not yet covered within the core portfolio
- domain-specific or sector-oriented use cases
- strategic and organisational aspects of AI adoption in SMEs and start-ups

An indicative example of such an external training offer is a course on “*Artificial Intelligence in Industrial SMEs: Understanding Economic Potential and Strategic Adoption*”, targeting decision-makers in small and medium-sized industrial enterprises. Such trainings focus on enabling organisations to assess the strategic relevance of AI, identify economically viable use cases, and integrate AI initiatives into existing business and operational contexts.

Brokered training opportunities allow users to access such specialised offerings through HammerHAI-supported mechanisms, contributing to a broader and more flexible training ecosystem while maintaining alignment with user needs and project objectives.

2.4.5 Workshop-Based Skill Development

Beyond classical training, HammerHAI offers a demand-driven and application-customized Workshop program for SMEs¹³. Those services are typically intermediary level and are targeted towards developing necessary specific skills towards the productization of ideas. The workshop format is in contrast to other service offering (except for the Training-on-Demand, see 2.4) offered “as a service”. While one part of such workshops is classical tech consulting, there is also a large training/skill development focus in these workshops. Rather than relying on continuous service delivery, the goal is to enable SMEs to develop the required skills independently. The partners have built specific training material to facilitate this. The focus of workshop based trainings is areas that need address specific needs according to the use case, such as:

- ethics and compliance
- and strategy and business case development

Partially, those workshop-based trainings are also brokered via partners of the AI factories (see Section 2.4.4).

¹³<https://www.hammerhai.eu/resources-solutions/skills-networking/>

2.4.6 Community-Based Learning and Journal Club

In addition to structured training modules and intensive learning formats, HammerHAI aims to foster continuous knowledge exchange through community-based learning activities. Such formats provide participants with opportunities to engage with emerging developments in artificial intelligence beyond the scope of predefined curricula and support the establishment of interdisciplinary learning communities across the HammerHAI ecosystem.

A representative example is the PaperHAI Journal Club, a recurring discussion format centred around recent scientific publications, technological advances, and emerging trends in AI and related fields. In contrast to traditional lecture-based training, the Journal Club emphasizes interactive exchange, critical reflection, and peer learning. Participants are encouraged to discuss the relevance, strengths, limitations, and practical implications of recent research results, thereby strengthening their ability to assess and apply new developments in their own professional contexts.

The Journal Club is designed to be accessible to a broad audience, including researchers, developers, technical specialists, and interested participants from industry. Depending on the selected topic, sessions may focus on methodological advances, application-oriented case studies, novel AI architectures, or cross-cutting themes such as trustworthiness, sustainability, and emerging regulatory considerations. Contributions from participants themselves are explicitly encouraged, enabling the format to adapt dynamically to current interests and technological developments.

By providing a low-threshold and continuously evolving learning opportunity, Journal Clubs complement the structured training portfolio and contribute to the development of an active community of practice within HammerHAI. At the same time, discussions within these formats can provide valuable insights into emerging skill requirements and user interests, supporting the ongoing refinement of the broader training programme.

2.5 Challenge-Based and Intensive Programmes

Beyond structured training modules and flexible learning formats, the HammerHAI training programme is further extended by a set of challenge-based and intensive formats that act as building blocks for applied, collaborative, and use-case-driven learning.

These formats enable participants to apply acquired knowledge in realistic settings, often in interdisciplinary teams and time-constrained environments. They are particularly well suited for SMEs, start-ups, and innovation-driven users who benefit from experiential learning and project-oriented engagement.

The formats described in this subsection are currently in preparation and are planned for implementation within the HammerHAI programme from fall 2026

onwards. They build on established concepts and prior experience from related European initiatives and are being adapted to the HammerHAI context and user community.

2.5.1 Hackathons

Hackathon formats within HammerHAI focus on collaborative, use-case-driven development in structured, time-bound settings, typically including a competitive element with evaluation and potential incentives for participating teams.

A key example is the European AI Hackathon¹⁴, a multi-week programme organised in collaboration with European AI Factory and antenna partners – AI:AT¹⁵, Gaia¹⁶ and Pharos-CY AIFA¹⁷ – as well as OpenACC and NVIDIA. HammerHAI acts as a contributing partner in its execution. The format builds on the established OpenHackathons model and continues the legacy of similar activities carried out under the EuroCC2 initiative [12], enabling participants to collaboratively develop advanced AI solutions in a coordinated European setting.

Complementary regional initiatives include formats such as the the x-athon, a multi-month programme that brings together SMEs, start-ups, industry partners, and technical experts to work on concrete business and technical challenges provided by participating organisations. HammerHAI contributes as a partner in their planning and future implementation, supporting the development of application-driven prototypes in close interaction with industry stakeholders.

2.5.2 Bootcamps

Bootcamps provide intensive, topic-focused training over several consecutive days, combining guided instruction with hands-on work.

The European AI Bootcamp¹⁸, organised in collaboration with the same partners as the European AI Hackathon (see Section 2.5.1) represents a key format in this category. With HammerHAI participating in its delivery, this programme builds on the experience and structure of earlier EuroCC2 training activities and provides an immersive environment for rapid skill development and applied experimentation. The strong interest in this format is reflected in the registration figures for the July 2026 edition: almost 500 applications have been received across EuroHPC member countries, underlining the demand for intensive, hands-on AI training.

¹⁴<https://www.openhackathons.org/s/siteevent/a0CUP00003yKxcX2AS/se000475>

¹⁵<https://ai-at.eu/>

¹⁶<https://gaia.plgrid.pl/>

¹⁷<https://www.cyi.ac.cy/>

¹⁸<https://www.openhackathons.org/s/siteevent/a0CUP00004egTDN2A2/se000495>

2.5.3 Mini-Hackathons and Short Innovation Formats

Short-format innovation activities complement larger programmes by enabling rapid exploration of ideas and initial solution concepts in a highly focused setting. These formats are typically embedded within events or workshops and may include competitive elements and incentives.

An example is the hub:disrupt mini-hackathon¹⁹, a one-day format integrated into the hub:disrupt event. In this setting, participating companies can pitch ideas for the use of HammerHAI infrastructure and services, develop initial concepts, and, where applicable, receive support for further exploration of these use cases. HammerHAI contributes as a partner in the organisation and execution of this format.

These short formats are particularly effective for initiating collaborations, engaging new users, and identifying potential follow-up projects within the HammerHAI ecosystem.

¹⁹<https://www.hubdisrupt.de/>

3 Agile Training Strategy and Roadmap Development

This section outlines the strategic positioning, requirements, and structural design of the HammerHAI training programme. It builds on the training catalogue described in Section 2 and examines how training activities align with user needs, service development, and broader project objectives. The section further analyses existing gaps and describes the mechanisms by which the training programme is structured and continuously refined.

3.1 Strategic Context

Training activities within HammerHAI are embedded within the broader strategic framework of the EuroHPC AI Factories initiative [13], which seeks to strengthen Europe's AI capabilities through a combination of advanced computing infrastructure, AI services, innovation support, as well as skills and talent development. Within this context, training serves a dual purpose: enabling effective utilisation of AI Factory resources and fostering a workforce capable of translating AI technologies into tangible economic and societal value.

The training roadmap is therefore designed not only to support the uptake of HammerHAI services, but also to contribute to wider European objectives related to AI adoption, competitiveness, innovation, and digital transformation. Recent European policy initiatives, including the Apply AI Strategy [2] and the emerging Data Union vision [14], emphasise the need to move beyond the development of AI technologies towards their large-scale deployment in strategic sectors, supported by trusted access to data, scalable computing infrastructure, interoperable services, and strengthened human capabilities. These initiatives highlight that Europe's competitiveness will depend on its ability to integrate AI into real-world industrial, economic, and societal contexts, rather than solely advancing algorithmic or research excellence.

Within this framing, training is a key enabler of AI adoption. It provides users with the knowledge and practical competencies required to access and exploit AI infrastructure, data resources, tools, and expert support. At the same time, it addresses broader challenges such as awareness raising, organisational readiness, change management, governance, and the development of AI literacy across diverse stakeholder groups. This reflects the growing recognition that successful AI deployment requires the co-development of technological, organisational, and human capabilities within coherent innovation ecosystems.

A core principle of the roadmap is agility. As AI technologies, regulatory frameworks, service portfolios, and industrial demands continue to evolve rapidly, training activities must continuously adapt to the changing needs of stakeholders. HammerHAI therefore pursues a demand-driven approach, aligning its training portfolio with feedback from industry, research organisations, educational institutions, start-ups, and public-sector actors. This encompasses collaboration

with mid-sized companies developing workforce upskilling strategies, educational providers seeking to complement theoretical curricula with hands-on and project-based learning formats, and innovation ecosystems such as incubators, accelerators, and venture capital networks supporting fast-growing companies.

Experience across these stakeholder groups consistently shows that the challenge is not the absence of training opportunities, but rather the lack of flexible, targeted, and immediately applicable learning formats. Many organisations—particularly SMEs and scale-ups—require training that can be integrated directly into ongoing innovation and product development activities, rather than traditional classroom-based programmes that demand significant time commitments. Consequently, the HammerHAI training roadmap emphasises practical, application-oriented learning pathways that accelerate the transition from AI awareness to experimentation, implementation, and scalable deployment.

The roadmap is further shaped by the recognition that Europe’s AI ambitions cannot be achieved by individual AI Factories or institutions acting in isolation. The strategic objectives of the EuroHPC AI Factories initiative—including accelerating AI uptake, strengthening European competitiveness, developing AI talent, and increasing innovation capacity—require coordinated action across a distributed and interconnected ecosystem. AI Factories are a central component of this ecosystem, but their effectiveness depends on close collaboration with the EDIHs, research and technology organisations, universities, industry associations, innovation agencies, public authorities, and private-sector partners. Each actor contributes complementary expertise, services, and access channels that together enable the development of a cohesive European AI ecosystem.

This ecosystem perspective is strongly aligned with the Apply AI Strategy and Data Union objectives, which emphasise the importance of integrated innovation systems where data, computing resources, AI services, and skills development are made accessible in a coordinated and interoperable manner. In particular, the Apply AI Strategy [2] identifies manufacturing, engineering, and related industrial sectors as priority domains for AI adoption, reflecting their central role in Europe’s economic base and competitiveness. These sectors are characterised by both high potential for AI-driven transformation and significant barriers to adoption, including fragmented data landscapes, complex legacy systems, and shortages of AI-skilled personnel.

For HammerHAI, manufacturing and engineering therefore represent a primary focus area. Despite strong technological and industrial foundations, many organisations in these sectors still face challenges in moving from experimental AI use cases to production-scale deployment. These challenges include limited access to scalable computing resources, difficulties in integrating AI into existing engineering and production workflows, lack of high-quality and interoperable datasets, and organisational and skills-related constraints. These barriers are particularly pronounced among SMEs and mid-sized industrial companies, which often possess strong domain expertise but limited capacity to independently

develop and operationalise AI solutions.

Within this broader landscape, HammerHAI's training activities are designed not only to serve its direct users but also to complement and reinforce the efforts of other ecosystem actors. The roadmap emphasises cooperation, interoperability, and knowledge exchange across initiatives wherever possible, ensuring alignment with other AI Factories, broader European initiatives such as the EDIHs, and national programmes, including the BMFTR-funded AI Service Centres (e.g. the AI Service Centre for Sensitive and Critical Infrastructure, KISSKI) and Centres of Excellence for AI Research (e.g. the Munich Center for Machine Learning, MCML)²⁰. As summarised in Figure 6, the German AI ecosystem actors relevant to HammerHAI can be grouped into five complementary categories: regional AI ecosystems, national AI competence centres, national AI service centres, European Digital Innovation Hubs, and national and cross-cutting initiatives. By engaging with actors across these groups, HammerHAI can leverage complementary expertise, avoid duplication, and broaden access to its training and service offer for industry and the public sector. Beyond this national ecosystem, HammerHAI also cooperates at the European level with other EuroHPC AI Factories and their antenna partners, for instance through jointly organised formats such as the European AI Hackathon and the European AI Bootcamp (see Section 2.5) and joint web seminars with the UK AI Factory Antenna (UKAIFA). This reflects the understanding that Europe's AI transformation is inherently a collective endeavour, requiring coordinated action across institutions and borders.

Within this context, training cannot be designed as an isolated activity. Instead, it must co-evolve with the AI Factory service roadmap and infrastructure development. As new AI services, computing capabilities, datasets, tools, and sector-specific solutions are introduced, corresponding training activities must be developed to ensure that users can effectively access and exploit these resources. Conversely, insights gained through training and user interaction provide critical feedback for refining and prioritising service development. Training and service evolution therefore form a continuous, bidirectional feedback loop that is essential for ensuring relevance, usability, and impact.

This integration is particularly important in manufacturing and engineering, where AI adoption typically follows a staged pathway from awareness and experimentation to validation, deployment, and scaling within operational environments. Effective support along this pathway requires the seamless combination of training, technical support, infrastructure access, data services, and innovation facilitation. Training should therefore be embedded across the entire AI Factory user journey, rather than positioned as a standalone offering.

In this setting, training acts as a critical bridge between technological capability and practical deployment. It connects AI infrastructure, data ecosystems, services, expertise, and innovation support mechanisms with the skills and

²⁰<https://mcml.ai/>




Figure 6: An overview of the key players in the German AI ecosystem relevant to HammerHAI, grouped into (1) regional AI ecosystems, (2) national AI competence centres, (3) national AI service centres, (4) European Digital Innovation Hubs, and (5) national and cross-cutting initiatives.

organisational capacities required for successful adoption. By creating coherent and interoperable learning pathways, training contributes to the broader objective of establishing a resilient, competitive, and AI-ready European economy. HammerHAI's training roadmap should therefore be understood not as a standalone programme, but as an evolving component of the EuroHPC AI Factory ecosystem and the wider European effort to build the human capital, organisational capabilities, and innovation systems necessary for the successful deployment of AI at scale.

Against this strategic background, a clearer understanding of the specific requirements of target user groups becomes essential. The following section therefore analyses the capabilities, expectations, and constraints that shape the design of the HammerHAI training programme.

3.2 Requirements Analysis

The identification of training requirements within HammerHAI is currently based on a combination of several input sources, each providing partial and complementary insights. These include direct discussions and indirect interactions with industry stakeholders, existing experience of consortium partners in delivering training and supporting AI adoption, and the analysis of expectations and recommendations formulated at the European level.

Additional input is derived from training activities conducted by the consortium, as well as from related formats such as hackathons and community events. However, these activities have so far been more strongly oriented towards scientific users as the main target audience and therefore provide only limited insight into the needs of the primary target group in industrial contexts consisting of SMEs and start-ups.

As a result, the current understanding of training requirements is based on heterogeneous but complementary sources, providing initial and actionable insights. While these inputs allow the identification of initial patterns and indicative needs, they do not yet constitute a comprehensive and systematically validated picture of user demand.

Based on these inputs, a set of characteristic requirement patterns can be identified, particularly for SMEs and start-ups in early stages of AI adoption.

A first observation is that many organisations are less interested in developing AI solutions from first principles and instead seek ready-to-use or easily adaptable solutions that can be integrated into their existing processes. A common example is the interest in retrieval-augmented generation (RAG) [15] approaches to unify and utilise internal company knowledge. At the same time, the rapid pace of development in AI is already shifting attention towards more advanced paradigms such as agentic AI systems, which extend beyond static information retrieval towards more autonomous and task-oriented behaviour.

At the same time, more advanced user groups, such as those engaged through EDIHs, already exhibit a higher level of AI maturity. Their requirements are more focused on the practical use, adaptation, and optimisation of existing models and tools, including topics such as fine-tuning, efficient deployment, and integration into existing workflows.

Across these user groups, several cross-cutting requirements can be observed. Topics such as regulatory compliance, trustworthy AI, and data handling are consistently relevant, reflecting increasing awareness of the European regulatory landscape. In addition, a subset of users addresses sovereignty-related goals, i.e. reaching or maintaining a suitable level of technological independence from large external providers, although this is not a universal concern.

Beyond these perceived needs, effective use of HammerHAI services requires a set of practical capabilities that are not yet consistently present across target user groups. These include an understanding of the distinction between model usage and model development, the ability to work with AI workflows and development environments, and familiarity with the execution of jobs on remote computing infrastructure. In many cases, users need to transition from local or notebook-based environments to more complex systems, which introduces additional challenges related to system access, configuration, and scaling of workloads.

To ensure user readiness for computing services, training is required to enable effective interaction with platform-level tools and abstractions that support the use of AI services on the underlying infrastructure. This covers operating within abstracted cloud-like environments built on technologies such as Kubernetes, as well as alternative scheduling systems such as SLURM for more HPC-affine users. While lower-level system details are increasingly abstracted through user interfaces and workflow tools, a basic understanding of how to configure and execute tasks, manage data (for example in environments such as the HammerHAI data platform²¹ based on piveau²²), and monitor results remains essential. This also includes awareness of security, data handling, and compliance aspects when working with shared or sensitive resources.

In addition to content-related requirements, recurring preferences regarding training formats can be identified. Short and targeted training sessions are generally more attractive than full-day formats, as they are easier to integrate into daily work schedules. Hands-on examples and use cases that closely reflect the participants' domain are considered essential for effective learning.

Furthermore, there is a strong demand for low-threshold entry points into AI and its associated software and hardware systems. Many users, especially those without a strong technical background, face difficulties with basic concepts such as system access, development environment configuration, or setting up

²¹<https://data.hammerhai.eu/en>

²²<https://piveau.hlr.de>

and running simple workflows. This indicates the need for introductory-level trainings that lower entry barriers and provide practical guidance for initial usage scenarios.

Taken together, these observations define the capabilities and conditions required for effective use of HammerHAI services across different user groups, while also indicating areas where these requirements are not yet consistently met in practice.

3.3 Skill Gap Findings

Building on the requirements analysis described above, several gaps can be identified that arise from mismatches between user expectations, the capabilities required for effective use of HammerHAI services, and the current level of user readiness and training provision.

In particular, a first key gap arises from the mismatch between user expectations and the effort required to implement AI solutions. In interactions with companies, there is a recurring expectation for fully developed solutions addressing specific business problems, which aligns more closely with consultancy or co-development models than with training. Discussions with industry stakeholders further indicate that this expectation is often accompanied by uncertainty regarding the respective roles of training, technical support, and implementation services along the AI adoption lifecycle. This reflects a divergence between perceived needs and the training-oriented approach, influenced in part by rapidly evolving AI technologies and emerging paradigms such as agentic AI.

A second gap relates to the level of technical and operational readiness of users. Effective use of HammerHAI services requires capabilities in areas such as workflow execution, environment configuration, and interaction with platform-level tools. However, many users—particularly SMEs and start-ups at early stages of AI adoption—lack experience in transferring existing workflows efficiently from local or notebook-based environments to distributed systems. This gap can hinder onboarding and early-stage usage.

Differences in user maturity (cf. Figure 3) further affect how training needs are addressed. While parts of the existing catalogue align well with technically experienced users, including developers and researchers familiar with HPC environments, other user groups require additional support for entry-level access and structured onboarding. Corresponding training material is in part already available through partner offerings and is being extended with HammerHAI-specific courses currently under development to better support these use cases.

Uncertainties also remain regarding the effectiveness of training formats. It is not yet clear to what extent target groups are willing to participate in in-person formats such as hackathons or intensive workshops, or whether more flexible and online-oriented approaches are more suitable. At the same time, stakeholder

feedback suggests that demand is shifting from introductory AI topics towards more technically in-depth and implementation-oriented formats, highlighting the need for clearly structured learning pathways that extend beyond initial awareness. This is further compounded by the fact that previous engagement has largely originated from communities outside the primary industrial target groups, resulting in a limited evidence base for industry-specific training needs.

Similarly, participation patterns indicate that initial interest in AI topics does not always translate into engagement in structured training activities. This suggests the presence of barriers related to time constraints, perceived relevance, or alignment with immediate business priorities.

Finally, the current findings remain indicative in nature. They are derived from a combination of heterogeneous input sources, including stakeholder discussions, exchanges with partners such as EDIHs, prior experience from related initiatives (e.g. in the context of EuroCC [12]), and previous consortium activities. These inputs include insights from targeted stakeholder workshops and ongoing exchanges with EDIHs, initially focused on a core group across the German-speaking European region, particularly in manufacturing- and engineering-oriented SME contexts. These interactions highlight recurring challenges such as the need for clearer progression along the AI adoption lifecycle, improved alignment between training and implementation support, and greater transparency regarding follow-up steps after initial engagement. While these inputs provide a solid basis of initial insights, they do not yet constitute a fully comprehensive and systematically validated picture of user demand. This reflects the fact that the training programme is – despite the existence of an established and defined portfolio – still in an active phase of continued requirements analysis and refinement (see also Section 3.6). Validation therefore continues as part of the ongoing improvement and development process. The systematic alignment of both the requirements analysis and the training portfolio will evolve further in response to emerging user needs and HammerHAI's engagement patterns, supported by continued stakeholder engagement, including further targeted workshops planned for upcoming project phases.

To address these identified gaps, several targeted measures are being developed and progressively implemented within the training programme. These include the expansion of low-threshold entry formats, such as web seminars and introductory modules, as well as increased use of flexible delivery mechanisms, including Training-on-Demand services and self-paced learning options.

In addition, application-oriented formats such as workshops, hackathons, and bootcamps are being adapted to better reflect SME-specific use cases and operational constraints, enabling more direct integration of training into ongoing innovation and development activities. Participation of companies in these formats also generates valuable insights into practical needs, which are fed back into the GSA framework to inform the ongoing refinement of the training portfolio. Collaboration with intermediary organisations, in particular the EDIHs and




Figure 7: Illustration of the four pillars of the AI Factory's talent development and start-up incubation strategy (from HammerHAI proposal)

regional partners like bwcon, further supports outreach to SMEs and facilitates access to training through established regional and sectoral networks.

Together, these measures aim to improve accessibility, relevance, and uptake of training among SMEs and early-stage users, while supporting a gradual transition towards more advanced and infrastructure-oriented training activities.

Overall, these observations and the measures outlined above underline that addressing identified skill gaps requires a combination of technical training, accessible onboarding, and appropriate delivery formats, supported by continuous adaptation to evolving user needs and system capabilities. This, in turn, necessitates a structured and flexible training approach aligned with the broader HammerHAI service portfolio and supporting different stages of user engagement.

3.4 Structural Alignment of the Training Programme

Building on the identified requirements and gap analysis, training activities within HammerHAI are structured as an integral component of the overall project architecture, supporting the adoption and effective use of AI technologies across its service portfolio. The programme is designed to anticipate evolving user needs and to provide a coherent framework that combines technical training and flexible delivery formats across different stages of user engagement, complementing broader onboarding and support activities within the project.

In line with the objectives of the work package, the training programme is further aligned with a four-pillar approach (cf. Figure 7) to talent development, comprising academic education, vocational training, lifelong learning, and support for AI-driven start-ups. While the current portfolio already addresses elements of these pillars—particularly through application-oriented formats such as hackathons, bootcamps, and SME-focused pathways—further integration of academic curricula, vocational training initiatives, and structured lifelong learning formats is being developed as part of the ongoing programme expansion.

Initial steps in this direction include ongoing discussions with academic and vocational training partners such as the MCML and the Duale Hochschule Baden-Württemberg, aiming to explore closer alignment with higher education programmes and practice-oriented training formats.

This development builds on the strong ecosystem of the HammerHAI AI Factory, which is embedded in a network of leading technical universities, Excellence Strategy institutions, applied AI research centres, and industry partners, complemented by regional innovation networks. While direct integration of training into formal academic programmes remains limited at this stage, alignment with academic education, research-oriented teaching, and industry-linked training is being progressively strengthened.

Future development includes closer coordination with European and national initiatives, including emerging collaboration frameworks among AIFs/AIFAs and related funding actions supporting cross-factory cooperation, mobility, and exchange. Participation in joint coordination structures, such as training working groups across AIFs/AIFAs, further supports the alignment of training approaches and the development of shared formats. In parallel, vocational training, lifelong learning, and start-up support are reinforced through collaboration with industry partners and regional innovation ecosystems.

From an organisational perspective, the training programme follows a dual-structured trainer model as defined in the project objectives. A core group of training specialists is responsible for the design and delivery of structured training modules and formats, ensuring pedagogical consistency and quality across the programme. This is complemented by a broader group of AI experts contributing through user support activities, application-oriented formats, and the integration of training with the HammerHAI service portfolio. This dual structure ensures that training activities are closely aligned with practical system usage while maintaining a high standard of structured skill development.

A central structural distinction is made between two complementary categories of training activities. On the one hand, general and cross-cutting AI training targets broader audiences such as management-level staff and organisations in early stages of digitalisation. These topics include AI literacy, strategic understanding of AI adoption, and regulatory aspects such as the EU AI Act. While these areas are emphasised at the European level, their delivery within HammerHAI is constrained by limited resources. As a result, such training is primarily supported through brokerage and collaboration with external partners, allowing existing expertise to be leveraged while gradually extending internal capabilities.

On the other hand, HammerHAI-specific training constitutes the core internal contribution of the programme. These activities are closely aligned with the project's infrastructure, tools, and services, focusing on enabling users to access, understand, and effectively use HammerHAI resources. This includes enabling users to effectively access and use the system, the use of AI tools and workflows, and the integration of AI solutions into practical applications. As the underlying

infrastructure is still being developed, training activities evolve in parallel with the availability and maturity of the services.

The alignment of training with the HammerHAI service portfolio is a key design principle²³. Training activities prepare users to work with services ranging from infrastructure access and AI model usage to privacy-oriented AI applications and consulting-related support such as AI adoption guidance, data innovation potential analysis, and trustworthy AI assessment. In this way, training supports the full lifecycle of AI adoption, from initial awareness and onboarding to advanced usage and optimisation. In addition, training activities serve as an accessible entry point into the HammerHAI ecosystem. They provide low-threshold opportunities for potential users to engage with AI technologies and, in doing so, support the identification and onboarding of new users for the broader service portfolio.

This connection also links the training programme to both strategic and technical layers of the project roadmap. At a strategic level, training supports SME engagement and the broader goal of increasing AI adoption, particularly in sectors such as manufacturing and engineering where uptake remains limited. At a technical level, it enables users to apply concrete tools, workflows, and services, thereby bridging the gap between conceptual understanding and operational use.

A key principle enabling this alignment is the modular structure of the training programme. Training content is organised into reusable building blocks that can be adapted to different user groups, formats, and levels of expertise. This modularity supports flexible learning pathways as well as reuse across structured courses, Training-on-Demand services, and application-oriented formats such as bootcamps and hackathons.

These formats serve complementary roles within the overall system. Structured modules provide foundational knowledge, Training-on-Demand services offer targeted and organisation-specific support, and challenge-based formats enable hands-on application and user engagement. Brokered training complements these elements by extending the portfolio in areas where specialised or general-purpose training is more efficiently sourced externally.

The design of the training programme is further shaped by strategic focus areas and current developments, including topics such as agentic AI, secure AI deployment, and trustworthy AI. At the same time, addressing the still limited adoption of AI in key industrial sectors remains a central objective, requiring not only advanced training but also accessible entry points and clear demonstrations of value.

In this context, the training programme is implemented as a coherent component of the HammerHAI ecosystem, with its structure reflecting the interplay between available services, infrastructure, and user engagement mechanisms. It provides

²³<https://www.hammerhai.eu/resources-solutions/>

a framework that connects training activities with the broader service landscape while remaining sufficiently flexible to accommodate future developments and extensions.

3.5 Insights and KPIs

To support and validate this structural approach, the development of the training programme is accompanied by a continuous monitoring framework that combines qualitative insights with a set of quantitative key performance indicators (KPIs). The need for clearly defined success metrics for training activities—including, for example, the number of participants, their organisational background and company type, and gender balance—has been explicitly emphasised in the review process, and is reflected in the monitoring framework outlined here. Given the early stage of both the infrastructure rollout and the engagement with the primary industrial target groups, this framework is designed to be lightweight and adaptive: it aims to capture early signals of relevance, uptake, and impact, while remaining flexible enough to evolve as the evidence base grows. Insights and KPIs are therefore not treated as fixed success criteria, but as instruments that feed directly into the GSA and the continuous improvement process described in Section 3.6.

This deliverable focuses on the definition and structure of the monitoring framework, while the detailed quantitative results—such as participant statistics broken down by organisation type and gender—are reported as part of the dedicated stakeholder analysis, rather than within this catalogue-oriented deliverable. The present section therefore concentrates on which indicators are tracked, how the underlying data is collected, and how the resulting insights feed back into the development of the training programme.

3.5.1 Qualitative Insights

A first set of insights is derived from the training and engagement activities conducted so far, as well as from direct and indirect interactions with users and partners. These observations are necessarily preliminary, as previous engagement has largely originated from scientific and HPC-affine communities, e.g. in the context of the Gauss Centre for Supercomputing (GCS) where both LRZ and HLRS are long-standing close partners, rather than from the primary industrial target groups of SMEs and start-ups.

Within these limitations, several recurring patterns can already be observed. Demand is consistently highest for short, application-oriented formats with hands-on components and domain-relevant examples, whereas longer and more theoretical formats attract comparatively lower participation from industry users. Initial interest expressed at events or fairs does not always translate into participation in structured training, indicating barriers related to time, perceived relevance, and immediate business priorities. Furthermore, entry-level topics

such as system access, environment configuration, and the transition from local to remote execution emerge as recurring obstacles, confirming the need for low-threshold onboarding offerings.

These qualitative insights are systematically captured and consolidated within the GSA, where they are linked to identified gaps and translated into concrete adjustments of the training portfolio. In this way, they complement the quantitative indicators described below and provide essential context for their interpretation.

3.5.2 Key Performance Indicators

To support the monitoring of the training programme in a more structured manner, a set of KPIs is being established along complementary dimensions: reach and participation, engagement quality and satisfaction, target-group penetration, and portfolio development. Together, these dimensions are intended to reflect not only the volume of training activities, but also their relevance to the intended audiences and their contribution to AI adoption.

The monitoring framework builds on established tracking practices within the consortium. In particular, HLRS and LRZ maintain a long-standing and detailed approach to course evaluation and participant statistics—covering, among others, participant numbers, the share of industry participants, gender distribution, and structured feedback on aspects such as course length, depth, and pacing—which draws on established survey frameworks from initiatives such as PRACE [16] and GCS and serves as a reference model for the wider programme. Building on this experience, the consortium is in the process of aligning a common set of KPI definitions and reporting conventions across all partners (LRZ, GWDG, KIT, and HLRS); in particular, registration processes and post-training feedback questionnaires are being progressively harmonised, so that data collected through different partner processes can be consolidated consistently at the programme level.

At the same time, it is recognised that partners operate within different institutional contexts and may contribute to additional initiatives with specific reporting requirements. As a result, harmonisation is being pursued in a pragmatic and incremental manner, aiming to achieve a consistent core set of indicators while maintaining compatibility with partner-specific constraints and external frameworks. This alignment is ongoing and will mature in parallel with the rollout of the training activities.

Reach and participation indicators capture the scale of the training offer, including the number of training activities delivered and the number of participants reached. Engagement-quality and satisfaction indicators assess how training is received, primarily through post-training feedback and satisfaction scores, as well as completion rates for self-paced and on-demand formats. Target-group penetration indicators specifically track the extent to which the programme reaches its primary audience, in particular SMEs and start-ups, which is a central

objective given the manufacturing- and engineering-oriented focus of HammerHAI. Finally, portfolio-development indicators monitor the evolution of the catalogue itself, including the share of HammerHAI-specific courses and the responsiveness of the programme to newly identified needs.

An overview of the indicators currently foreseen is provided in Table 2. The associated target values are to be understood as indicative reference points rather than fixed commitments. As a robust baseline is still being established, targets will be refined iteratively as more representative data from the industrial target groups becomes available.

To illustrate the type of data captured by these indicators at the level of an individual course, Table 3 shows measured values for one course as a concrete example: the Data Analysis with HPC module offered by HLRS. The figures are provided solely to illustrate the kind of metrics that are tracked; comprehensive, programme-wide statistics—including breakdowns by organisation type and gender—are reported as part of the stakeholder analysis rather than within this deliverable. At this stage, such measured data has been consolidated only for HLRS courses, as the corresponding tracking processes are most mature there. The collection of the same set of indicators by all relevant consortium partners (LRZ, GWDG, KIT, and HLRS) is foreseen for future reporting periods, with the aim of enabling the single-course example shown here to be extended into a more consolidated, programme-wide picture.

The same course also illustrates the value of longitudinal tracking: compared with the previous edition, participation increased from 19 to 31 and the average satisfaction improved markedly (from a German-system grade of 2.7 to 1.6, corresponding to roughly 3.3/5 and 4.4/5 respectively). Such year-on-year trends are exactly the kind of signal the monitoring framework is designed to capture and feed back into the continuous improvement process.

Beyond individual course statistics, early signals of programme-level reach and demand are already observable across the broader portfolio, even ahead of the full cross-partner consolidation described above. In particular, the challenge-based and intensive formats (see Section 2.5) provide concrete evidence of strong uptake: the European AI Bootcamp, delivered jointly with European AI Factory partners, attracted more than 380 applications for its July 2026 edition already a week before the application deadline, while established consortium courses such as the HLRS *Data Analysis with HPC* module show both growing participation (from 19 to 31 participants year-on-year) and high satisfaction (4.4/5). Taken together, and although these figures do not yet constitute a single, consistently consolidated programme-wide statistic, they demonstrate substantial demand for the training offer and provide a tangible baseline from which the harmonised reach and participation indicators of Table 2 will be built up over the coming reporting periods.

The collection of KPI data relies on existing mechanisms wherever possible, including registration and feedback processes of the consortium partners, ana-

Table 2: Indicative key performance indicators for the HammerHAI training programme. Targets are preliminary and subject to refinement as the baseline is established.

Indicator	Description	Data source / target
Reach and participation		
Training activities delivered	Number of training activities (courses, sessions, events) offered per reporting period	Training calendar / partner records; target: ≥ 20 per year
Participants reached	Total number of unique participants across all formats	Registration data; target: ≥ 200 per year
Engagement quality and satisfaction		
Participant satisfaction	Average satisfaction rating from post-training feedback	Feedback forms; target: $\geq 4/5$
Completion rate (self-paced)	Share of started self-paced/on-demand modules that are completed	Platform analytics; target: $\geq 60\%$
Target-group penetration and demographics		
Organisation type	Breakdown of participants by organisation type (SME / start-up, large enterprise, academia, public sector)	Registration data; target: $\geq 30\%$ SME/start-up
Industry sector coverage	Reach into priority sectors (manufacturing, engineering)	Registration data; qualitative review
Gender balance	Distribution of participants by gender	Registration / feedback data
Portfolio development		
HammerHAI-specific courses	Number/share of courses tailored to HammerHAI infrastructure and services	Catalogue review; target: growing share
Responsiveness	Number of portfolio adjustments derived from GSA findings per cycle	GSA backlog; qualitative review

Table 3: Illustrative measured KPIs for a single course (HLRS, Data Analysis with HPC), based on the 2025 edition and its post-course feedback.

Indicator	Measured value (2025)
Participants	31
Participant-days	201.5
Participant satisfaction	4.4 / 5
Course length rated adequate	69%
Pace rated adequate	77%
Teaching materials rated useful	92%

lytics from self-paced learning platforms, and the central training calendar. To limit additional overhead, data collection is integrated into ongoing operational activities rather than implemented as a separate reporting layer.

At the current stage, monitoring and evaluation practices remain partly decentralised and rely on a combination of partner-specific processes and manual consolidation of results. Building on this foundation, the programme is progressively developing more structured and semi-automated analysis workflows. This includes the exploration of automated data aggregation and advanced analytical approaches to identify patterns in user needs and derive actionable insights for the continuous refinement of the training portfolio.

KPIs are reviewed as part of the regular work-package cycles described in Section 3.6. Rather than serving purely as accountability metrics, they are used as a steering instrument: deviations from expectations are interpreted within the GSA to identify underlying causes and to inform concrete adjustments to the training portfolio, formats, and outreach. In line with the agile and demand-driven nature of the programme, both the indicators themselves and their target values are expected to evolve over the course of the project as the understanding of user needs matures and the service portfolio develops. In this way, the monitoring framework establishes a direct link between observed outcomes and the ongoing adaptation of the training programme, forming a key input to the continuous improvement process described in the following section.

3.6 Continuous Improvement and Roadmap Implementation of the Training

Building on the monitoring and evaluation framework described in the previous section, training-related activities, as well as their development and maintenance, are integrated into the HammerHAI roadmapping process, which comprises both strategic and operational update cycles to support continuous improvement. This process is supported by regular direct and indirect interactions at both project-wide and work-package levels. At the work package level, we aim for structured internal review cycles of approximately six to eight weeks. These cycles allow for

assessing progress, identifying required adjustments, and ensuring alignment with the evolving service portfolio.

Taking into account the rapid evolution of AI technologies, the ongoing bring-up of the AI Factory's infrastructure, and the growing understanding of user needs, the training catalogue is treated as a living component of the AI Factory rather than a fixed deliverable. At the same time, parts of the monitoring and analysis activities have been developed progressively over the course of the project. As a result, the current framework reflects an early stage of consolidation, with further enhancements expected as dedicated roles in training coordination and user engagement have been established more fully during the recent project phase. This is expected to support a noticeable acceleration in both data collection and the systematic integration of user feedback into the continuous improvement process.

At the operational level, the refinement of training activities follows an iterative approach inspired by agile methodologies. Potential extensions and improvements are systematically identified based on internal discussions, user interactions, and observations from ongoing activities and new developments. These inputs are systematically structured through the GSA framework, which links stakeholder needs, identified gaps, and resulting training actions. This includes also targeted feedback from application-oriented formats involving industrial participants, particularly SMEs, which is incorporated into the continuous refinement of the training portfolio. The outcomes of this process are progressively consolidated into a backlog of training-related topics.

Prioritisation of training activities is guided by the outcomes of the GSA process and focuses on expected user relevance and potential impact, rather than fixed long-term planning. This allows for flexible adjustment to emerging demands, including shifts in technology trends and changes in user engagement patterns. It also supports experimentation with new formats and topics, which can be evaluated and adapted based on observed uptake.

Coordination with external partners and initiatives forms an additional element of this process. Exchanges with other AI Factories and collaboration partners in joint activities, such as hackathons, provide opportunities to share experiences, identify relevant developments, and align training efforts within the broader European ecosystem.

The evolution of the training programme is also closely linked to the implementation timeline of the HammerHAI infrastructure and services. Early user onboarding activities, expected to begin in the second half of 2026, will primarily focus on introducing initial system capabilities, supporting first access, and gathering practical feedback from early users. These early interactions will play an important role in refining system-specific training content and identifying remaining gaps in usability and support.

As the infrastructure matures, training activities will be progressively expanded

towards more advanced and application-oriented topics. Broader service usage is envisioned to commence in 2027, accompanied by training on scalable workflows, service integration, and optimisation.

In parallel, foundational aspects such as security awareness, responsible AI practices, and regulatory considerations—including compliance with the EU AI Act [1]—are incorporated from the outset. This ensures that users are equipped not only with technical skills but also with the necessary understanding of trustworthy and compliant AI usage, thereby reinforcing the structured and evolving approach described above.

Overall, this approach establishes training as an actively managed and continuously adapting component of the HammerHAI ecosystem. The present deliverable should therefore be understood as an inventory of the initial phase embedded into this ongoing process, reflecting the current state of development while anticipating further refinement as both the project and its user engagement evolve.

4 Conclusion

This deliverable presents the initial design and structuring of the HammerHAI training catalogue, together with the underlying strategic considerations guiding its development. Section 2 provided an overview of the current training portfolio, including core training tracks, target groups, and available formats. This catalogue reflects the existing expertise within the consortium and establishes a first set of offerings addressing both general AI topics and HammerHAI-specific capabilities. At the same time, it highlights that parts of the portfolio, in particular those closely tied to the emerging infrastructure and services, are still under development.

Building on this foundation, Section 3 examined the strategic context, analysed user requirements, and identified existing skill and alignment gaps. It further described the structural design of the training programme, its integration with the HammerHAI service portfolio, and the mechanisms for continuous monitoring and improvement. The analysis shows that while the current portfolio already addresses a number of relevant topics and user groups, it does not yet fully reflect the diversity of requirements, particularly with respect to SMEs and early-stage users, nor the full range of capabilities required for effective system usage.

The comparison between the existing portfolio and the identified strategic requirements therefore reveals both alignment and remaining gaps. On the one hand, the modular structure, the availability of application-oriented formats, and the integration into the broader service ecosystem provide a solid basis for enabling AI adoption. On the other hand, further work is required to expand HammerHAI-specific training, strengthen low-threshold entry points, and ensure closer alignment with actual user needs and engagement patterns.

To address these gaps, the training programme is explicitly designed as a continuously evolving component of the HammerHAI ecosystem. Through the combined use of monitoring tools, including qualitative insights and KPIs, and structured processes such as the GSA, the training portfolio will be iteratively refined. Early user onboarding, initial system usage phases, as well as application-oriented training activities and events will provide valuable feedback, enabling the progressive development of more targeted and advanced training content aligned with the maturation of the infrastructure and services.

In this sense, the programme described in this deliverable should be understood as a snapshot of an ongoing development process rather than a finalised training offering. As HammerHAI progresses and user engagement intensifies, the training programme will continue to adapt, ensuring its relevance and effectiveness in supporting the broader objective of accelerating AI adoption and strengthening AI capabilities across the European economy and society.

References

- [1] European Parliament and Council of the European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L series, 12 July 2024, 2024. Article 4 on AI literacy applicable from 2 February 2025. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>.
- [2] European Commission. Communication from the Commission to the European Parliament and the Council: Apply AI Strategy. COM(2025) 723 final, Brussels, 8 October 2025, 2025. <https://digital-strategy.ec.europa.eu/en/library/commission-communication-apply-ai-strategy>.
- [3] Duri Long and Brian Magerko. What is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*, pages 1–16. ACM, 2020.
- [4] Riina Vuorikari, Stefano Kluzer, and Yves Punie. DigComp 2.2: The Digital Competence Framework for Citizens. Technical Report EUR 31006 EN, JRC128415, Publications Office of the European Union, Joint Research Centre (JRC), Luxembourg, 2022.
- [5] European Committee for Standardization (CEN). European e-Competence Framework (e-CF) – A common European Framework for ICT Professionals (EN 16234-1:2019). CEN, Brussels, 2019. <https://www.ecompetences.eu/>.
- [6] John Biggs. Enhancing teaching through constructive alignment. *Higher Education*, 32(3):347–364, 1996.

- [7] Lorin W. Anderson and David R. Krathwohl. *A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives*. Longman, New York, 2001.
- [8] HPC SPECTRA. High-Performance Computing Skills Platform and European Collaboration for Training (HPC SPECTRA). EuroHPC JU, 2023. <https://www.hpc-spectra.eu/>.
- [9] EuroHPC Virtual Training Academy (EVITA). EuroHPC Virtual Training Academy (EVITA) and the Competence and Qualification Framework (CQF). EuroHPC JU, 2024. <https://www.evitahpc.eu/>.
- [10] Julian M. Kunkel, Weronika Filinger, Christian Meesters, and Anja Gerbes. The HPC Certification Forum: Toward a Globally Acknowledged HPC Certification. *Computing in Science & Engineering*, 22(4):110–114, 2020.
- [11] Constanze Hasterok and Janina Stompe. PAISE[®] – process model for AI systems engineering. *at – Automatisierungstechnik*, 70(9):777–786, 2022.
- [12] EuroCC / CASTIEL Consortium. EuroCC: National Competence Centres in High-Performance Computing. EuroHPC JU EuroCC programme, 2024. <https://www.eurocc-access.eu/>.
- [13] EuroHPC Joint Undertaking. AI Factories. European High Performance Computing Joint Undertaking (EuroHPC JU), 2024. Established under Council Regulation (EU) 2021/1173 as amended by Council Regulation (EU) 2024/1732. https://www.eurohpc-ju.europa.eu/ai-factories_en.
- [14] European Commission. Communication from the Commission to the European Parliament and the Council: Data Union Strategy – Unlocking Data for AI. COM(2025) 835 final, Brussels, 19 November 2025, 2025. <https://digital-strategy.ec.europa.eu/en/library/data-union-strategy-unlocking-data-ai>.
- [15] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 9459–9474, 2020.
- [16] PRACE. PRACE Training Portal and PRACE Advanced Training Centres (PATC). Partnership for Advanced Computing in Europe, 2024. <https://training.prace-ri.eu/>.

A Illustrative Snapshot of the Training Catalogue

This preliminary listing provides an illustrative snapshot of course descriptions underlying the HammerHAI online training catalogue. It reflects a selection of courses delivered across different phases of the project and is intended to give an overview of the structure, level of detail, and scope of the training offer. The listing is automatically generated from underlying data sources and may therefore still contain minor inconsistencies, such as formatting issues or translation artefacts, which are subject to ongoing refinement.

At the current stage, this listing should be understood as a representative example rather than a complete or final catalogue. The training portfolio is under continuous development, and both the structure and content of the catalogue are expected to evolve. In addition, a live version of upcoming events is accessible via the HammerHAI course calendar, while a full catalogue will be provided via <https://www.hammerhai.eu/training-catalog/>, offering access to the most up-to-date information. Future versions of this deliverable may place increased emphasis on such live references or include updated automatically generated snapshots once the underlying data and integration mechanisms are fully established. As such, the material in this appendix is provided for illustrative purposes only.

Note on language: The HammerHAI training programme is delivered in both English and German. To remain faithful to the original offerings, course titles and descriptions in this listing are reproduced in the language in which they are provided by the respective consortium partner. Consequently, some entries in this appendix appear in German rather than English.

Courses

- AI Ethics Lecture - Privacy is not Power, but what is?
- AI Training Series - Data Parallelism - How to Train Deep Learning Models on Multiple GPUs
- AI Training Series - Fundamentals of Deep Learning
- AI Training Series - Introduction to Container Technology & Application to AI at LRZ
- AI Training Series - Introduction to the LRZ AI Systems
- AI Training Series - Model Parallelism - Building and Deploying Large Neural Networks
- AI Training Series - Orientation Session
- AI Training Series - Scaling Deep Learning - From Single GPU to Clusters
- Artificial Intelligence and Ethics
- Data Analytics for Engineering Data using Machine Learning
- Data Management Concepts for Efficient and User - Friendly HPC
- Deep Learning Bootcamp - Building and Deploying AI Models
- Deep Learning with GPU Cores
- Effectively Utilize AI Tools in Research
- European Agentic AI Bootcamp
- European AI Hackathon
- Getting Started with Linux Bash
- Getting Started with the AI Training Platform
- GPU Programming with CUDA
- Introduction to AlphaFold
- Introduction to Neuromorphic Computing and the SpiNNaker Hardware Platform
- KI in der Verwaltung - Eine Einführung in die Nutzung für alle Mitarbeiter*innen
- Multi - GPU Deep Learning
- Performance Engineering in ML Workloads Using Score - P and Vampir
- Performance Engineering Tools for AI and HPC Workloads
- Secure HPC - Parallel Computing with Highest Security
- Supercomputing Academy - Deployable Data Analysis & AI Pipelines with HPC
- Technical responsibility and ethical methods
- Time Series Forecasting
- Using the GWDG Data Pools for Scientific Data Sharing

A.1 AI Ethics Lecture - Privacy is not Power, but what is?

Keywords: AI

Course page: <https://www.hlrs.de/event/2026/privacy-is-not-power-but-what-is>

Abstract

Data is the fuel of many AI-systems, but at what cost? In her book, *Privacy is Power*, Carissa Véliz argues that we must provide absolute privacy protections because, otherwise, we will suffer both individually and collectively from power asymmetries in politics, economics, and social life. In particular, Véliz is concerned with maintaining a functioning liberal democracy. In this talk, it is argued that the worries raised by Véliz cannot be resolved by protecting privacy alone; instead, these worries are due to diminishments of the closely related concept of “informational autonomy”, which will be qualified in this talk. Moreover, it is shown that the problems identified by Véliz have no simple overarching solution—such as absolute protection of privacy or informational autonomy—because we face a set of contextually sensitive tradeoffs. Björn will continue our AI ethics lecture series with a talk on privacy in the age of AI.

Philosopher Björn Lundgren from the Centre for Philosophy and AI Research (FAU Erlangen) will join us in person on May 7th, 2026 – 11am at HLRS Room Berkeley.

Online link:

A stream will be available at <https://unistuttgart.webex.com/unistuttgart/j.php?MTID=m6e7fcf16a09d5ce65361e3681bb5d21e>

A.2 AI Training Series - Data Parallelism - How to Train Deep Learning Models on Multiple GPUs

Keywords: AI, Deep Learning, GPU

Course page: https://app1.edooobox.com/en/LRZ/Onsite%20Courses/Course.ed.5b29f1dbbc34_9721639477.AI%20Training%20Series%20-%20Data%20Parallelism%20-%20How%20to%20Train%20Deep%20Learning%20Models%20on%20Multiple%20GPUs

This course is part of the “LRZ AI Training Series”, a series of courses aiming at the needs and expectations of data analytics, big data & AI users at LRZ.

The course is organised as an on-site even at LRZ in Garching near Munich. There will be no possibility to join online remotely via video conference. Participants are expected to bring their own laptops running the latest version of Chrome or Firefox. There are no PCs installed in the course room!

Contents:

Modern deep learning challenges leverage increasingly larger datasets and more complex models. As a result, significant computational power is required to train models effectively and efficiently. Learning to distribute data across multiple GPUs during deep learning model training makes possible an incredible wealth of new applications utilizing deep learning.

Additionally, the effective use of systems with multiple GPUs reduces training time, allowing for faster application development and much faster iteration cycles. Teams who are able to perform training using multiple GPUs will have an edge, building models trained on more data in shorter periods of time and with greater engineer productivity.

This course teaches you techniques for data-parallel deep learning training on multiple GPUs to shorten the training time required for data-intensive applications. Working with deep learning tools, frameworks, and workflows to perform neural network training, you'll learn how to decrease model training time by distributing data to multiple GPUs, while retaining the accuracy of training on a single GPU.

The course is co-organised by LRZ and NVIDIA Deep Learning Institute (DLI). All instructors are NVIDIA certified University Ambassadors.

Learning Objectives:

By participating in this course, you'll:

- Understand how data parallel deep learning training is performed using multiple GPUs
- Achieve maximum throughput when training, for the best use of multiple GPUs
- Distribute training to multiple GPUs using Pytorch Distributed Data Parallel
- Understand and utilize algorithmic considerations specific to multi-GPU training performance and accuracy

A.3 AI Training Series - Fundamentals of Deep Learning

Keywords: AI, Deep Learning

Course page: https://app1.edoobox.com/en/LRZ/Onsite%20Courses/Course.ed.27b7bd983720_9721630234.AI%20Training%20Series%20-%20Fundamentals%20of%20Deep%20Learning

This course is part of the "LRZ AI Training Series", a series of courses aiming at the needs and expectations of data analytics, big data & AI users at LRZ.

The course is organised as an on-site even at LRZ in Garching near Munich. There will be no possibility to join online remotely via video conference. Participants

are expected to bring their own laptops running the latest version of Chrome or Firefox. There are no PCs installed in the course room!

Contents:

Businesses worldwide are using artificial intelligence to solve their greatest challenges. Healthcare professionals use AI to enable more accurate, faster diagnoses in patients. Retail businesses use it to offer personalised customer shopping experiences. Automakers use it to make personal vehicles, shared mobility, and delivery services safer and more efficient. Deep learning is a powerful AI approach that uses multi-layered artificial neural networks to deliver state-of-the-art accuracy in tasks such as object detection, speech recognition, and language translation. Using deep learning, computers can learn and recognise patterns from data that are considered too complex or subtle for expert-written software.

In this course, you'll learn how deep learning works through hands-on exercises in computer vision and natural language processing. You'll train deep learning models from scratch, learning tools and tricks to achieve highly accurate results. You'll also learn to leverage freely available, state-of-the-art pre-trained models to save time and get your deep learning application up and running quickly.

The course is co-organised by LRZ and NVIDIA Deep Learning Institute (DLI). All instructors are NVIDIA certified University Ambassadors.

Learning Objectives:

By participating in this course, you will:

- Learn the fundamental techniques and tools required to train a deep learning model
- Gain experience with common deep learning data types and model architectures
- Enhance datasets through data augmentation to improve model accuracy
- Leverage transfer learning between models to achieve efficient results with less data and computation
- Build confidence to take on your own project with a modern deep learning framework

A.4 AI Training Series - Introduction to Container Technology & Application to AI at LRZ

Keywords: AI, Container, Infrastructure

Course page: https://app1.edoobox.com/en/LRZ/Online%20Courses/Online%20Course.ed.8ef59efc55a4_10146078037

This course is part of the "LRZ AI Training Series", a series of courses aiming at the needs and expectations of data analytics, big data & AI users at LRZ.

While tailored to LRZ users, the course is also open to other interested academic participants from Germany.

Contents Since the introduction of docker back in 2013, container technology has become the industry standard for software packaging, distribution, and deployment. Creating a container consists in bundling an application, its dependencies and runtime in a single unit that can later run independently of the underlying infrastructure. Unlike virtual machines, containers are lightweight and yield higher performances while providing greater versatility and interoperability. As containers accommodate an easy, safe, reliable, and scalable way to run applications and pipelines, they are an attractive candidate for high performance computing and artificial intelligence workloads. With this module, we will showcase the most enticing features and niceties offered by containers. Not only will we explore their history and implementations, but we will also dive into actual and cutting edge uses with a particular emphasis on artificial intelligence tasks, reproducible biomedical pipelines, and automated workflows. Participants will roll up their sleeves and get their hands on the compute cloud of LRZ to set containers in action. By the end of the course, participants will be able to transfer their experience and knowledge to their specific use-cases and requirements.

Three content blocks of roughly one hour each:

1 A Tour of Containers and Kickstart 1.1 You Are Here 1.2 Basic Concepts for Containers 1.3 Containers vs. the World 1.4 Kickstarter: Your First Steps With Containers

2 Under the Hood of Containers: Warm-up Lap to Flat-out 2.1 The Bolts and Nuts of Containers 2.1.1 namespaces, cgroups, and copy-on-write Storage 2.1.2 Volumes, Networking, and More 2.1.3 Configuration and Image Creation 2.2 Warm-up lap: Make an ANN Dream in a Container 2.3 Containers and Security 2.3.1 Security and User Namespace 2.3.2 Security and Docker Socket Exploit 2.4 Containers and HPC 2.5 Flat-out: Containers and HPC With Enroot and Charliecloud

3 Containers on Nitromethane 3.1 Abstract Your Containers 3.1.1 Continuous Integration / Continuous Deployment 3.1.2 Reproducible Scientific Pipelines 3.2 Hardware Accelerated Containers 3.3 Orchestration and Scaling Across a Compute Cluster 3.4 Last Lap: Reproducible Transcriptomic Workflow With Containers and Nextflow

End: Wrap-up, take home message, last questions and book / website / podcast / events recommendations

Prerequisites

- AI Training Series: Orientation Session (or comparable previous knowledge): access the remote work environment, navigate the work environ-

ment, edit files, pull images, spin containers...

Hands-On

- Hands-on #1: Your first steps with containers
- Hands-on #2: Make an artificial neural network dream in a container
- Hands-on #3: Container and HPC with Charliecloud
- Hands-on #4: Enable transcriptomic reproducible workflow with containers and Nextflow

Content Level The content level of the course is broken down as:

Language English

Lecturers Navdar Karabulut and Sérgio Tafula (LRZ)

Prices and Eligibility The course is open and free of charge for people from academia and industry from the Member States of the European Union and Associated/Other Countries to the Horizon 2020 programme.

Registration Please register with your official e-mail address to prove your affiliation.

A.5 AI Training Series - Introduction to the LRZ AI Systems

Keywords: AI, Deep Learning, GPU, Infrastructure

Course page: https://app1.edoobox.com/en/LRZ/Online%20Courses/Course.ed.b071919747ca_10146087280

This course is part of the "LRZ AI Training Series", a series of courses aiming at the needs and expectations of data analytics, big data & AI users at LRZ. While tailored to LRZ users, the course is also open to other interested academic participants from Germany.

Contents The aim of this course is to give an overview of the LRZ AI Systems, and provide participants with the knowledge and skills necessary to efficiently utilise them. The course consists of mini lectures, demos and hands on sessions (breaks included).

By participating in this lecture, you will be able to:

- Understand the resources that the LRZ AI System provides
- How to allocate resources on the LRZ AI System and provision them with the needed software stack
- How to interactively work with the LRZ AI Systems via the terminal (also Jupyter Notebooks for single GPU workload)

Upon completion, you will be able to effectively use the LRZ AI Systems to run Deep Learning workflows.

Prerequisites

- AI Training Series: Orientation Session (or comparable previous knowledge)
- AI Training Series: Introduction to Container Technology & Application to AI at LRZ (or comparable previous knowledge)
- Good understanding of Deep Learning and Classical Machine Learning (courses such as Introduction to Deep Learning (I2DL) (IN2346) are provided by TUM - material also available for non-TUM students)

Hands-On During the course a live demo on how to access and operate the LRZ AI system will be showcased. Exercises will be conducted on the LRZ AI Systems. In addition, the parallelisation of the training of a ML model will be also demonstrated.

Content Level The content level of the course is broken down as:

Language English

Lecturers Dr. Mares Barekzai and Benjamin Geißler (both LRZ)

Prices and Eligibility The course is open and free of charge for people from academia and industry from the Member States of the European Union and Associated/Other Countries to the Horizon 2020 programme.

Registration Please apply with your official email address to prove your affiliation. The final participants will be selected and informed after the registration deadline has passed. Priority will be given to users of the LRZ AI systems, who are kindly requested to provide their Project ID associated with the LRZ AI Systems in the registration form.

Withdrawal Policy See Withdrawal

A.6 AI Training Series - Model Parallelism - Building and Deploying Large Neural Networks

Keywords: AI, Deep Learning, GPU

Course page: https://app1.edoobox.com/en/LRZ/Onsite%20Courses/course.ed.f0be145585cd_9721657963.AI%20Training%20Series%20-%20Model%20Parallelism%20-%20Building%20and%20Deploying%20Large%20Neural%20Networks

This course is part of the "LRZ AI Training Series", a series of courses aiming at the needs and expectations of data analytics, big data & AI users at LRZ.

The course is organised as an on-site even at LRZ in Garching near Munich. There will be no possibility to join online remotely via video conference. Participants are expected to bring their own laptops running the latest version of Chrome or Firefox. There are no PCs installed in the course room!

Contents:

Large language models (LLMs) and deep neural networks (DNNs), whether applied to natural language processing (e.g., GPT-3), computer vision (e.g., huge Vision Transformers), or speech AI (e.g., Wave2Vec 2), have certain properties that set them apart from their smaller counterparts. As LLMs and DNNs become larger and are trained on progressively larger datasets, they can adapt to new tasks with just a handful of training examples, accelerating the route toward general artificial intelligence. Training models that contain tens to hundreds of billions of parameters on vast datasets isn't trivial and requires a unique combination of AI, high-performance computing (HPC), and systems knowledge. The goal of this course is to demonstrate how to train the largest of neural networks and deploy them to production.

The course is co-organised by LRZ and NVIDIA Deep Learning Institute (DLI). All instructors are NVIDIA certified University Ambassadors.

Learning Objectives:

By participating in this workshop, you'll learn how to:

- Scale training and deployment of LLMs and neural networks across multiple nodes.
- Use techniques such as activation checkpointing, gradient accumulation, and various forms of model parallelism to overcome the challenges associated with large-model memory footprint.
- Capture and understand training performance characteristics to optimize model architecture.
- Deploy very large multi-GPU, multi-node models to production using NVIDIA Triton™ Inference Server.

A.7 AI Training Series - Orientation Session

Keywords: AI, Infrastructure, Linux, MPI, OpenMP

Course page: https://app1.edoobox.com/en/LRZ/Online%20Courses/Online%20Course.ed.d3c27a749323_10146068794

This two-day online "Orientation Session" is the kick-off event of the "LRZ AI Training Series", a series of courses aiming at the needs and expectations of data analytics, big data & AI users at LRZ. While tailored to LRZ users, the course is also open to other interested academic participants from Germany.

Contents By joining this session, participants will learn about the landscape of LRZ systems and services relevant to data analytics & AI users. Further, they will be provided with the initial training of the most essential skills needed to work with these systems (i.e. fundamentals of GNU/Linux, SSH). Participation in this Orientation Session should be considered a mandatory requirement for anyone interested in other sessions of the series.

It has been about 350 years that separate the original 1673 Leibniz mechanical calculator from today's Leibniz Supercomputing Centre's (LRZ) facilities located on the Campus Garching. And yet, the spirit has not changed. To paraphrase the German mathematician: "It is beneath the dignity of excellent people to waste their time in calculation when the work could just as accurately be done with the aid of a machine."

Participants joining this course will be provided with a comprehensive walk-through and usage guides to contemporary types of such computing machines which can be found at LRZ. In a general overview, historical and current developments and trends in the field of scientific computing and cluster systems will be presented. This will, amongst others, address the following questions: How are modern cluster systems designed and how do they work? How did we come to High Performance Computing, High Performance Data Analytics and High Performance AI? What makes a system adequate to a specific workload? How are these systems operated and how are they made available to their users?

In addition, an introduction to GNU/Linux and the Unix Shell as well as the use of remote systems via Secure Shell (SSH) will be provided. GNU/Linux is a family of open source operating systems, powering all different kinds of computing hardware (from mobile devices to high performance computing clusters... and pretty much everything in-between). The typical command line interface is a Unix-like shell. It serves as interactive command environment and scripting language, allowing users to control the system, and to automate tasks of varying complexity. SSH is a cryptographic network protocol, which is typically used to login and execute commands on remote (GNU/Linux) systems.

The introduction to GNU/Linux opens with a short historical overview and some common concepts & terminology will be explained. Then the focus is directed

toward working with the Unix Shell. First, it will be used to navigate the file system and directories of a system, then the mechanisms of file manipulation and ownership will be explored. Additional commands and concepts will be explained, and characteristics as well as methods of manipulating the shell environment will be discussed.

Participants will be guided to install and configure a SSH client on their personal systems. Different applications for remote access and file transfer will be introduced. A conceptual and practical introduction to SSH keys will be given, followed by a discussion of more advanced features like SSH tunneling and custom configuration options.

These skills will then be combined to further explore typical interaction methods and usage patterns of multiuser cluster systems, including various possibilities of setting up user environments (e.g. environment variables and modules, user space package managers, containers) as well as tools for resource allocation (i.e. Slurm Workload Manager) and efficient parallelization (MPI, OpenMP, ...). Finally, an overview of different compute clusters as well as their storage systems operated by LRZ will be provided. The requirements for gaining access to these systems will be covered as well.

Participants will obtain a good understanding of the characteristics of multiuser cluster systems in general and will practise basic methods of typical interaction. They will acquire essential knowledge and skills necessary to successfully interact with the command line interface of a GNU/Linux system. They will learn to use SSH in order to connect to and interact with different remote systems. They will familiarize themselves with the landscape of cluster systems available at LRZ and this will allow them to choose the right system for their own compute projects.

The material will be presented as a combination of lectures, demos and hands-on sessions, with a focus on the latter. There will be breaks during the session.

Preliminary Agenda Day 1:

- History of computing
- Building intelligent machines
- Panorama of systems at LRZ
- The SLURM scheduler

Day 2:

- Basics of Linux
- Intro to SSH
- SSH tunnels

Prerequisites

- Some training material will be provided prior to the actual course dates and participants are encouraged to study this by themselves in preparation of the event.

Content Level The content level of the course is broken down as:

Language English

Lecturers Dr. Jophin John, Florian Kaspar, Dr. Martin Ohlerich, and Sérgio Tafula (all LRZ)

Prices and Eligibility The course is open and free of charge for people from academia and industry from the Member States of the European Union and Associated/Other Countries to the Horizon 2020 programme.

Registration Please register with your official e-mail address to prove your affiliation.

A.8 AI Training Series - Scaling Deep Learning - From Single GPU to Clusters

Keywords: AI, Data Analysis, Deep Learning, Machine Learning

Course page: https://app1.edoobox.com/en/LRZ/Onsite%20Courses/Course.ed.a8d6dced45cd_10146151981

We are planning to reschedule the course later this year. |

This course is part of the "LRZ AI Training Series", a series of courses aiming at the needs and expectations of data analytics, big data & AI users at LRZ. While tailored to LRZ users, the course is also open to other interested academic participants from Europe.

The course is organised as an on-site event at LRZ in Garching near Munich. There will be no possibility to join remotely via video conference. Participants are expected to bring their own laptops running the latest version of Chrome or Firefox. There are no computers provided in the course room!

Contents This 3-days block course combines lectures about Fundamentals of Deep Learning, Data Parallelism and Model Parallelism.

The lectures are interleaved with many demos and hands-on sessions using Jupyter Notebooks.

The course is co-organised by LRZ and NVIDIA Deep Learning Institute (DLI). All instructors are NVIDIA certified University Ambassadors.

Day 1: Fundamentals of Deep Learning

Businesses worldwide are using artificial intelligence to solve their greatest challenges. Healthcare professionals use AI to enable more accurate, faster diagnoses in patients. Retail businesses use it to offer personalised customer shopping experiences. Automakers use it to make personal vehicles, shared mobility, and delivery services safer and more efficient. Deep learning is a powerful AI approach that uses multi-layered artificial neural networks to deliver state-of-the-art accuracy in tasks such as object detection, speech recognition, and language translation. Using deep learning, computers can learn and recognise patterns from data that are considered too complex or subtle for expert-written software.

In this lecture, you'll learn how deep learning works through hands-on exercises in computer vision and natural language processing. You'll train deep learning models from scratch, learning tools and tricks to achieve highly accurate results. You'll also learn to leverage freely available, state-of-the-art pre-trained models to save time and get your deep learning application up and running quickly.

By participating in this lecture, you will:

- Learn the fundamental techniques and tools required to train a deep learning model
- Gain experience with common deep learning data types and model architectures
- Enhance datasets through data augmentation to improve model accuracy
- Leverage transfer learning between models to achieve efficient results with less data and computation
- Build confidence to take on your own project with a modern deep learning framework

Day 2: Data Parallelism

Modern deep learning challenges leverage increasingly larger datasets and more complex models. As a result, significant computational power is required to train models effectively and efficiently. Learning to distribute data across multiple GPUs during deep learning model training makes possible an incredible wealth of new applications utilizing deep learning.

Additionally, the effective use of systems with multiple GPUs reduces training time, allowing for faster application development and much faster iteration cycles. Teams who are able to perform training using multiple GPUs will have an edge, building models trained on more data in shorter periods of time and with greater engineer productivity.

This lecture teaches you techniques for data-parallel deep learning training on multiple GPUs to shorten the training time required for data-intensive applica-

tions. Working with deep learning tools, frameworks, and workflows to perform neural network training, you'll learn how to decrease model training time by distributing data to multiple GPUs, while retaining the accuracy of training on a single GPU.

By participating in this lecture, you'll:

- Understand how data parallel deep learning training is performed using multiple GPUs
- Achieve maximum throughput when training, for the best use of multiple GPUs
- Distribute training to multiple GPUs using Pytorch Distributed Data Parallel
- Understand and utilize algorithmic considerations specific to multi-GPU training performance and accuracy

Day 3: Model Parallelism

Large language models (LLMs) and deep neural networks (DNNs), whether applied to natural language processing (e.g., GPT-3), computer vision (e.g., huge Vision Transformers), or speech AI (e.g., Wave2Vec 2), have certain properties that set them apart from their smaller counterparts. As LLMs and DNNs become larger and are trained on progressively larger datasets, they can adapt to new tasks with just a handful of training examples, accelerating the route toward general artificial intelligence. Training models that contain tens to hundreds of billions of parameters on vast datasets isn't trivial and requires a unique combination of AI, high-performance computing (HPC), and systems knowledge. The goal of this course is to demonstrate how to train the largest of neural networks and deploy them to production.

By participating in this lecture, you'll learn how to:

- Scale training and deployment of LLMs and neural networks across multiple nodes.
- Use techniques such as activation checkpointing, gradient accumulation, and various forms of model parallelism to overcome the challenges associated with large-model memory footprint.
- Capture and understand training performance characteristics to optimize model architecture.
- Deploy very large multi-GPU, multi-node models to production using NVIDIA Triton™ Inference Server.

Important information After you are accepted, please create an account under <https://learn.nvidia.com/join>.

Ensure your laptop / PC will run smoothly by going to <http://websocketstest.com/> Make sure that WebSockets work for you by seeing under Environment, WebSockets is supported and Data Receive, Send and Echo Test all check Yes under WebSockets (Port 80). If there are issues with WebSockets,

try updating your browser.

NVIDIA Deep Learning Institute The NVIDIA Deep Learning Institute delivers hands-on training for developers, data scientists, and engineers. The program is designed to help you get started with training, optimising, and deploying neural networks to solve real-world problems across diverse industries such as self-driving cars, healthcare, online services, and robotics.

Prerequisites

- Day 1: An understanding of fundamental programming concepts in Python 3, such as functions, loops, dictionaries, and arrays; familiarity with Pandas data structures; and an understanding of how to compute a regression line.
- Day 2: Experience with deep learning training using Python.
- Day 3: Good understanding of PyTorch, good understanding of deep learning and data parallel training concepts, practice with natural language processing are useful, but optional

Hands-On The lectures are interleaved with many hands-on sessions using Jupyter Notebooks. The exercises will be done on a fully configured GPU-accelerated workstation in the cloud.

Language English

Lecturers Ajay Navilarekal Rajgopal (LRZ, NVIDIA certified University Ambassador)

Prices and Eligibility The course is open and free of charge for people from academia from the Member States of the European Union (EU) and Associated Countries to the Horizon 2020 programme.

Registration Please register with your official e-mail address to prove your affiliation.

Withdrawal Policy See Withdrawal

A.9 Artificial Intelligence and Ethics

Keywords: AI

Course page: <https://www.hlr.de/de/event/artificial-intelligence-and-ethics>

Cultural techniques for processing information have been around for a very long time, from slide rules to calendars. Cuneiform writing on clay tablets and printing with movable type are examples of the processing of finite alphabets using discrete, i.e. ‘spaced’, digital state sequences. Today, we associate digital information processing primarily with Turing-powerful electronic digital computers and big data. Ever since ChatGPT was launched at the end of 2022, ‘artificial intelligence’ (AI) has been on everyone’s lips. This lecture begins in the 1940s and spans the arc to today’s large language models. Using selected episodes, authors and concepts, it tells a dual story: the emergence of today’s information processing in conjunction with its social, philosophical and ethical reflection. In fact, AI ethics began in the 1940s at the latest, was enriched by developments in applied ethics in the 1970s, and today has achieved particular prominence in the areas of regulation, technology assessment and ethical guidelines. Against this backdrop, the stories of technology ethics behind current buzzwords such as ‘trustworthy and transparent AI’ or ‘bias’ become visible.

Philosopher of technology Michael Funk (U Vienna) will join us on February 19th, 4pm for an introductory lecture on ‘Artificial Intelligence’ and Ethics. This lecture is the first part of a two part lecture (second part on February 20th) especially aimed at computer scientists, engineers and technologists.

Join us online at: <https://unistuttgart.webex.com/unistuttgart/j.php?MTID=mebb62414e14ee39f97f40b314f745d09>

A.10 Data Analytics for Engineering Data using Machine Learning

Keywords: AI, Machine Learning

Course page: <https://www.hlr.de/training/2026/ml4sim>

Fraunhofer SCAI in cooperation with HLRS offers a three-day workshop on data analytics for simulation data using machine learning.

This three-day online workshop addresses the preparation, analysis and interpretation of numerical simulation data by machine learning methods. Besides the introduction of the most important concepts like clustering, dimensionality reduction, visualization and prediction, this course provides several practical hands-on tutorials using the python libraries numpy, scikit-learn and pytorch as well as the SCAI DataViewer.

A.11 Data Management Concepts for Efficient and User - Friendly HPC

Keywords: Data Analysis, HPC, Linux

Course page: <https://academy.gwdg.de/p/event.xhtml?id=6916d81f5067a75055a4c0bd>

Data management is generally challenging, particularly on HPC systems. Modern HPC systems offer different storage tiers with different characteristics. Some of these characteristics are for instance the availability of backups, the storage capacity, the IO performance, the difference between node local and globally available access, the semantics of the storage system, and the duration for which the storage endpoint is available, ranging from years to quarters, and sometimes only hours. This is confusing and entails different challenges and risks. First of all, users have to be aware of the different storage tiers and their performance profiles to optimize their job runtimes and not leave their jobs starving for data or wait for minutes that a Python environment has been loaded. However, users then need to move their results back to a storage tier with enough space and durability, to not lose their results at the end of a computation or soon after. While moving input and output data around users have to keep oversight over the data provenance to ensure the reproducibility and retrospective comprehensibility of their research. In addition, sometimes users don't just want to copy an entire data set but want to explore only a concise subset. For this, a data catalog can be used where all available data is indexed with respect to some domain-specific metadata. Once this data catalog is filled with all the data sets of a user, concise queries can be used to select the input data, and ideally, stage it to the correct storage tier as part of the job submission process. This data catalog can also be used to keep the oversight of all data that are distributed over the different storage tiers.

This course will provide an introduction to the different storage tiers available at GWDG and for what workloads they should be used. Then the concept of a data catalog and its usage of will be covered. Both parts will offer hands-on exercises on our HPC system.

Learning goal 1. Learn the concept of storage tiers and how to properly use them for best performance and data durability 2. Learn the concept of a data catalog and how to use it to select input data for an HPC job based on domain specific metadata

A.12 Deep Learning Bootcamp - Building and Deploying AI Models

Keywords: AI, Deep Learning, Jupyter Notebook, Machine Learning, Python

Course page: <https://academy.gwdg.de/p/event.xhtml?id=6916fd5e5067a75>

055a4c0dc

This bootcamp is designed to provide an introduction to deep learning. The course will cover the process of building deep learning models using popular frameworks like TensorFlow and PyTorch. Additionally, participants will be introduced to the basics of deploying deep learning models, including the use of web interfaces for model interaction. The course will also include practical exercises where participants will apply their learning to build and deploy a simple AI model.

Learning goal

- Deep Learning Fundamentals: Understanding the core concepts of neural networks
- Model Building with TensorFlow and PyTorch
- Deployment of AI Models

A.13 Deep Learning with GPU Cores

Keywords: AI, CUDA, Deep Learning, GPU

Course page: <https://academy.gwdg.de/p/event.xhtml?id=6916ce4a5067a75055a4c0b0>

Event content

Graphic processors enable efficient deep learning on a massive scale. This course will explore how to use deep learning frameworks with GPU cores. Each participant gets hands-on experience with an example workflow. As the GWDG will offer several high-end GPU nodes, this course will help our attendees to get the most out of our new hardware. Moreover, theoretical concepts such as hardware architecture will be covered.

Learning Goal After attending the course, participants

- know basic concepts of GPU cores and when to use them
- can migrate their deep learning workflow to the GPU
- have an overview of hardware accelerators available on our clusters

Requirements

- Participants should have a basic knowledge of Linux and experience in one programming language (preferably Python).
- In addition, basic deep-learning understanding is recommended.

A.14 Effectively Utilize AI Tools in Research

Keywords: AI, Deep Learning, Machine Learning

Course page: <https://academy.gwdg.de/p/event.xhtml?id=6916fbc15067a75055a4c0d9>

This course is designed for researchers and scientists from the University of Göttingen, the University Medical Center Göttingen (UMG), and the Max Planck Society who are interested in enhancing their research capabilities through the application of artificial intelligence (AI). Participants will explore how AI can assist in analyzing large datasets, automating routine tasks, and improving literature research and organization. The course also addresses the legal and ethical considerations surrounding the use of AI in research, ensuring that participants are equipped to use these tools responsibly and in compliance with relevant standards.

Learning goal

- Gain an understanding of how AI can support and enhance research efforts
- Develop practical skills in using AI tools for data analysis and literature research
- Learn about the legal and ethical frameworks governing the use of AI in research
- Explore specific use cases of AI-enhanced research processes
- Master efficient prompting techniques for AI tools and strategies for integrating these tools into research workflows

A.15 European Agentic AI Bootcamp

Keywords: AI, HPC

Course page: https://app1.edoobox.com/en/LRZ/Bootcamps%20and%20Hackathons/Course.ed.0868b75e015c_10384815484.European%20Agentic%20AI%20Bootcamp

General Information

This Bootcamp is co-organized by Austrian Scientific Computing (ASC), the Leibniz Supercomputing Centre (LRZ), the High-Performance Computing Center Stuttgart (HLRS), the Academic Computer Centre Cyfronet AGH (Cyfronet), OpenACC organization, and NVIDIA. It will host a virtual European Agentic AI Bootcamp starting 20 July 2026 and concluding 22 July 2026.

Agentic AI employs sophisticated reasoning to solve complex problems, utilizing data from multiple sources to analyze, strategize, and complete tasks independently. This bootcamp introduces powerful building blocks to accelerate the

development and deployment of agentic AI workflows.

The program covers tools such as Agent Skills, Model Context Protocol (MCP), NVIDIA® NIM™, LangGraph, and NeMo Agent Toolkit. Attendees will be guided through creating Agent Skills, building MCP clients/servers, and integrating these into LLM workflows orchestrated with LangGraph and NeMo, powered by NVIDIA NIM microservices. A hands-on challenge will allow participants to apply these skills to modify an existing workflow.

Why join this bootcamp

- Build a practical AI agent for a real use case
- Learn reusable patterns your team can adapt quickly
- Reduce the risk of investing in the wrong architecture
- Understand how to evaluate quality, latency, and performance early
- Get hands-on guidance

Who should attend SMEs, startups, developers, and researchers interested in accelerating AI applications. Prerequisite: Basic Python knowledge (no GPU programming required)

Key details

- Dates: 20–22 July 2026
- Format: Virtual
- Language: English
- Cost: Free of charge
- Eligibility: Participants affiliated with institutions in EU / EuroHPC member countries

How to apply Apply via the Openhackathons.org registration page using your official email address. Selected participants will be informed after the registration deadline.

A.16 European AI Hackathon

Keywords: AI

Course page: https://app1.edoobox.com/en/LRZ/Bootcamps%20and%20Hackathons/Online%20Course.ed.dce962f3cf35_10327924819.European%20AI%20Hackathon%20register%20via%20openhackathons

Contents Open Hackathons are multi-day, hands-on coding events that help developers and researchers accelerate, optimize, and scale their applications on a variety of data center architectures. Attendees work with dedicated mentors

to identify and address data loading bottlenecks, parallelize training, and optimize AI pipelines for faster experimentation to realize performance gains and speedups.

The European AI Hackathon, starting 6 October and concluding 29 October 2026, targets small and medium-sized enterprise (SMEs), startups, industry practitioners, and researchers across Europe building AI-driven products and services. Teams benefit from personalized, expert guidance to take their AI workloads to the next level, and access to leadership-class compute resources .

Examples of AI projects include:

Agentic Contract Intelligence: Multi-agent pipelines that parse, risk-score, and redline contracts across languages and jurisdictions

Real-Time Energy Grid Balancing: Agentic systems running thousands of parallel simulations to autonomously balance distributed energy resources within live decision windows

Legacy Code Modernization: Fine-tuning open-weight code LLMs on COBOL/Fortran codebases with long-context windows to enable enterprise migration

Autonomous Quality Inspection: Vision models running high-FPS inference to detect defects on production lines

Predictive Maintenance at Scale: Multimodal LLMs that integrate sensor logs, manuals, and historical failures to recommend interventions

These projects are illustrations, not constraints. Any AI workload that benefits from scaling, speed-up, and hands-on expert guidance is welcome.

Applying teams should have a clear project that is ready to be further developed and scaled. Participating teams must be based in the EU and/or EuroHPC JU member countries. Please ensure that your team meets all prerequisites before you apply and use your company/institute email address in the application form.

This hackathon is co-organized by Austrian Scientific Computing (ASC), the Leibniz Supercomputing Centre (LRZ), the High-Performance Computing Center Stuttgart (HLRS), the Academic Computer Centre Cyfronet AGH (Cyfronet), OpenACC organization, and NVIDIA for AI Factory Austria AI:AT, HammerHAI, and Gaia AI Factory.

Prerequisites

- Teams are expected to be fluent with the code or project they bring to the event and motivated to make progress during the hackathon.
- A minimum of 3 team members must participate throughout the entire event. Lack of complete participation will result in the entire team being excused.

- All team members must be based in the EU or EuroHPC JU member countries.
- Projects brought to the event are required to have a license attached and detailed in the application. Permissive-style Open Source License (e.g. BSD, MIT or Apache 2.0 license) are preferred as it is much easier for the hackathon organizer to pair your team with Mentors.

Details of the course can be found [here](#)

A.17 Getting Started with Linux Bash

Keywords: Linux

Course page: <https://academy.gwdg.de/p/event.xhtml?id=6916cd385067a75055a4c0ae>

Are you afraid of the black box, the one with the \$ sign in the front, the terminal? Become acquainted with the Linux Bash in order to harness the megawatts of power the HPC offers.

Want power? We Bash!

Using Linux and in particular the Bash is new for many first-time HPC users. This course is specifically designed to make the transition from laptop to Data Center easy. We will slowly and gently get you used to running a Bash, navigating the file tree to find folders and edit files. Most importantly, we will show you how to use the SSH protocol to access the front-end servers of our HPC cluster.

Note on the language of the course: We prepare the slides for the course in English, but we can also present in German if all participants understand German.

Learning goal

- Using a Bash for work and improve your efficiency
- Exploring and understanding the Linux file tree
- Navigating folders and editing files without IDE
- Using SSH for remote login into the HPC front-end servers

A.18 Getting Started with the AI Training Platform

Keywords: AI, HPC, Linux

Course page: <https://academy.gwdg.de/p/event.xhtml?id=6916df545067a75055a4c0cb>

Using the KISSKI services can be a bit overwhelming without knowledge about how to use HPC resources. We will teach how to get onto the front end systems and start batch jobs using SLURM.

Additionally, we will work on using other resources such as the chat-ai services and jupyter hub to get you acquainted with running these services.

It is highly suggested to have experience with the Bash. A course covering Linux basics is offered before this course. You will be able to use the resources, but optimal utilization is currently only achievable using a Bash.

Note on the language of the course: We prepare the slides for the course in English, but we can also present in German if all participants understand German.

Learning goal

- Getting to the front-end servers
- Using SLURM
- Resource selection
- Using the module system
- Using non SLURM KISSKI resources

A.19 GPU Programming with CUDA

Keywords: AI, C, C++, CUDA, Fortran, GPU

Course page: <https://academy.gwdg.de/p/event.xhtml?id=6916ded85067a75055a4c0ca>

Graphic processors with massive parallelism (GPU) are increasingly used as computational accelerators suitable for highly parallel applications. The latest upgrade of the GWDG Compute Cluster for High Performance Computing has added nodes with a total of 50 GPU accelerator cards. CUDA is a widely used programming environment for GPUs. This course introduces hardware and parallelization concepts for GPUs. The CUDA programming environment is described in detail, both for C and Fortran, including the language elements for controlling the processor parallelism and for accessing the various levels of memory. The use of GPU accelerated libraries (cuBLAS, cuFFT) is demonstrated. All topics are explained by means of examples in practical exercises.

Learning goal Use of CUDA to optimize algorithms on GPU

Requirements

- Using the GWDG Scientific Compute Cluster - An Introduction, or equivalent knowledge
- Practical experience with C
- For the practical exercises: GWDG account (preferable) or course account (available upon request), own notebook

A.20 Introduction to AlphaFold

Keywords: AI, Data Analysis, Machine Learning

Course page: <https://academy.gwdg.de/p/event.xhtml?id=6916ddbe5067a75055a4c0c8>

AlphaFold is a groundbreaking machine learning tool in the complex field of protein folding. Simple to use, but hard to master and understand, AlphaFold provides a vital step in many bioinformatics and molecular simulation workflows. In this tutorial, we will cover the theoretical background of AlphaFold, show examples of incorporating it into a research workflow, provide an opportunity to perform hands-on AlphaFold simulations, and explore advanced techniques and alternative algorithms.

Learning goal

- Basics of protein folding with AlphaFold
- Incorporating AlphaFold into your workflow
- Hands on folding and simulation experience

Requirements

- Knowledge of HPC usage in general, working with a terminal
- Background knowledge of basic biological concepts (DNA, proteins, protein synthesis, etc)
- Not necessary/will be covered, but nice to have: AI/ML, molecular simulations, GPUs

A.21 Introduction to Neuromorphic Computing and the SpiNNaker Hardware Platform

Keywords: AI, HPC, Jupyter Notebook

Course page: <https://academy.gwdg.de/p/event.xhtml?id=6916f7b15067a75055a4c0d3>

Neuromorphic computing is a computing paradigm that combines brain-inspired algorithms and associated hardware architectures. While it has been developed for decades, the field is recently gaining strong traction due to its possible applications in computational neuroscience and AI. Neuromorphic computing is based on the behavior of biological spiking neurons, and is well suited to implement and study the dynamics of neuronal networks, graph and optimization problems, and specific AI/machine learning applications. A key property of neuromorphic hardware is that it often exhibits tremendously lower power consumption when compared to traditional hardware and algorithms, which is a crucial concern in today's power hungry computational workloads.

In this course, you will learn about the basic concepts of neuromorphic computing, as well as the different libraries available to perform neuromorphic simulations. Furthermore, you will be among the first to gain hands-on access to the novel SpiNNaker-2 hardware architecture, and you may get to apply neuromorphic approaches to a series of problems including simulating neuronal networks and machine learning tasks.

This course should be of interest to those working with neuron and network models in neuroscience, AI/machine learning, optimization problems, and to those interested in developing new algorithms and working with novel hardware in general. Only basic knowledge of Python is required, and an idea of how biological neurons operate may be helpful.

Learning goal

- Learn what neuromorphic computing is, and what sort of problems it can solve
- Get to know different hardware systems used in neuromorphic computing
- Gain access to the novel neuromorphic computing platform SpiNNaker-2, and learn how to perform basic neuromorphic simulations with it

Requirements

- Basic Python knowledge
- Knowledge of how biological neurons function is an advantage but not necessary

A.22 KI in der Verwaltung - Eine Einführung in die Nutzung für alle Mitarbeiter*innen

Keywords: AI, Machine Learning

Course page: <https://academy.gwdg.de/p/event.xhtml?id=6916f9815067a75055a4c0d6>

Dieser Kurs richtet sich an alle Verwaltungsmitarbeiterinnen der Universität Göttingen, der Universitätsmedizin Göttingen (UMG) und der Max-Planck-Gesellschaft, die daran interessiert sind, die Effizienz ihrer Arbeitsabläufe durch den Einsatz von Künstlicher Intelligenz (KI) zu steigern. Insbesondere sind damit Mitarbeiterinnen in der zentralen und dezentralen Verwaltung, in den Sekretariaten und Antragsberatungen, Anfängerinnen und aktive Nutzerinnen gemeint, die ihre bestehenden Kenntnisse erweitern wollen. Der Kurs vermittelt grundlegende Konzepte und Anwendungen von KI, die speziell auf die Anforderungen der Verwaltungsprozesse zugeschnitten sind. Teilnehmer*innen lernen, wie KI-Tools praktisch eingesetzt werden können, um Routineanfragen zu automatisieren und Daten optimal zu verwalten. Darüber hinaus werden die rechtlichen Rahmenbedingungen und Datenschutzrichtlinien

behandelt, um eine verantwortungsvolle Nutzung von KI in der Verwaltung sicherzustellen.

Lernziel

- Verständnis erlangen, wie KI die Effizienz in Verwaltungsprozessen steigern kann
- Praktische Fähigkeiten im Umgang mit KI-Tools zur Automatisierung und Datenanalyse entwickeln
- Kenntnisse über die rechtlichen Rahmenbedingungen und Datenschutzrichtlinien im Zusammenhang mit der Nutzung von KI in der Verwaltung erwerben
- Einblicke in konkrete Anwendungsfälle von KI in Verwaltungsprozessen gewinnen
- Strategien zur effektiven Formulierung von Anfragen an KI-Tools erlernen und die Integration dieser Tools in den Verwaltungsalltag optimieren

A.23 Multi - GPU Deep Learning

Keywords: CUDA, Deep Learning, HPC, Jupyter Notebook

Course page: <https://www.hlr.de/training/2026/DL-Multinode>

This course has been canceled due to a pause of the NVIDIA Deep Learning Institute program! We plan to offer it again once the program resumes.

NVIDIA Deep Learning Institute (DLI) offers hands-on training for developers, data scientists, and researchers looking to solve challenging problems with deep learning.

Learn how to train and deploy a neural network to solve real-world problems, and how to effectively parallelize training of deep neural networks on Multi-GPUs.

The workshop combines the following DLI courses:

- Fundamentals of Deep Learning
- Data Parallelism: How To Train Deep Learning Models on Multiple GPUs
- Model Parallelism: Building and Deploying Large Neural Networks.

The lectures are interleaved with many hands-on sessions using Jupyter Notebooks. In addition, a short introduction to multi-GPU programming will provide a first look at fundamental concepts for parallelizing machine learning and deep learning across multiple GPUs.

This course is organized in cooperation with Leibniz Supercomputing Centre (Germany). The instructor is an NVIDIA certified University Ambassador.

A.24 Performance Engineering in ML Workloads Using Score - P and Vampir

Keywords: AI, Machine Learning, Score-P, Vampir

Course page: <https://academy.gwdg.de/p/event.xhtml?id=691ca4525067a75055a4c0f4>

Performance analysis is essential for understanding and optimizing the efficiency of ML/AI workloads, where computational and I/O bottlenecks can significantly impact training time, scalability, and resource utilization. By systematically identifying hotspots and inefficiencies, researchers and engineers can make informed decisions to improve both performance and energy efficiency. In this workshop, we introduce Score-P and Vampir, two powerful and complementary tools for performance measurement and visualization. Score-P provides detailed profiling and tracing of program behavior, while Vampir enables intuitive exploration of these traces through interactive visual analysis, helping users uncover complex performance patterns in ML workflows. The workshop consists of an introductory talk covering fundamental concepts, followed by a hands-on session where participants will instrument an example ML workload, collect performance data, and analyze the results using Vampir.

Learning goal

- Having a good understanding of performance analysis and its impact on AI workloads
- An understanding of the difference between I/O and computation bottlenecks
- Getting familiar with Profiling tool Score-P and visualization tool Vampir
- Getting knowledge on how to use the tools in analyzing the performance of a code

Requirements

- Basic Linux command-line skills
- Foundational understanding of ML workloads and performance concepts
- Basic understanding of Python codes

A.25 Performance Engineering Tools for AI and HPC Workloads

Keywords: AI, Data Analysis, GPU, HPC, I/O, Linux, Score-P, Vampir

Course page: <https://academy.gwdg.de/p/event.xhtml?id=691703345067a75055a4c0df>

With the increasing adoption of AI technologies, evaluating computational performance of AI applications in HPC systems has become critical for results im-

provement. In this course we shall use HPC performance tools to profile and evaluate performance bottlenecks in deep neural networks and learn tips for efficient training and deployment of deep learning models in HPC systems. Tools to be covered in this course include: Nvidia Nsight System, Score-P and Vampir.

Learning goal The course is intended to equip participants with fundamental knowledge on how to efficiently use HPC systems to run AI applications.

Requirements

- Practical knowledge of deep learning frameworks, Tensorflow or Pytorch
- Programming skills in Python
- Knowledge in Linux

A.26 Secure HPC - Parallel Computing with Highest Security

Keywords: Container, HPC

Course page: <https://academy.gwdg.de/p/event.xhtml?id=6916cf095067a75055a4c0b1>

Learning goal

- Understanding the importance of Secure HPC to process workflows that involve sensitive data
- Familiarize with Secure HPC main steps
- Understanding the execution with an automatic script

Requirements

- Fundamental proficiency in Linux commands (e.g., cd, mkdir, ...)
- Initial exposure to SSH for remote access
- First experience submitting Slurm jobs

A.27 Supercomputing Academy - Deployable Data Analysis & AI Pipelines with HPC

Keywords: AI, Data Analysis, Deep Learning, Machine Learning

Course page: <https://www.hlr.de/training/2026/sca-da-an>

This course is designed for practitioners in SMEs, startups, and applied research teams who need to turn data into actionable insights and repeatable, deployable ML workflows using high-performance computing (HPC) resources. Participants learn practical patterns to build reproducible and scalable analytics and ML pipelines that remain reliable under real-world constraints, including messy data,

growing volumes, performance bottlenecks, and execution across heterogeneous compute environments.

For teams preparing to use HPC resources, the course provides a clear path from notebook prototypes to reproducible, job-based workflows. Participants develop solutions in Jupyter Notebooks for rapid iteration and then operationalize them as script-based jobs suitable for production-style execution on high-performance systems. The hands-on curriculum covers scalable data processing with Apache Spark, performance-aware execution, and portable environments that help teams turn allocated compute time into measurable progress.

By the end of the course, participants will take back reusable assets they can apply immediately, including a project template (notebooks and job scripts), evaluation reports, scalable preprocessing pipelines, and runnable workflows that scale from laptop to HPC environments.

Prerequisites and content levels

Prerequisites

- Working knowledge of Linux/Unix: shell commands, SSH, file permissions, directory structure, and basic script editing (nano/vim/emacs).
- Basic Python skills (or strong experience in another language such as C/C++/Fortran/Java and the ability to ramp up quickly).

Recommended (not required):

- Basic Git usage (clone/commit/branch)
- Familiarity with typical data/ML libraries (numpy, pandas, matplotlib, seaborn) and Jupyter Notebooks.
- Familiarity with virtual environments (venv/conda) or containers

Content levels

- Beginners: 20 hours
- Intermediate: 20 hours
- Advanced: 10 hours

Target audience

This course is intended for, but not limited to:

- HPC users who want to run modern data analytics and ML workflows efficiently on shared compute infrastructure.
- Engineers and researchers who need practical skills in data preparation, supervised ML, deep learning fundamentals, and scalable processing.
- Industry data roles, including data analysts/scientists, data engineers, ML engineers, and platform/AI infrastructure engineers who want to move from prototypes to repeatable, scalable pipelines.

- Professionals who want to stay up-to-date with Big Data, HPC, and modern AI trends.

Agenda

- Week 1: Introduction to Data Analysis and AI within HPC
 - Week 2: Machine Learning (ML) and Deep Learning (DL) – from fundamentals to first runnable workflows
 - Week 3–4: Practical exercises in data pre-processing, feature engineering, and machine learning applications
 - Week 5: Emerging HPC/AI trends – agentic AI, containers, and cost/performance scaling
- Registration closes on Sunday, August 23, 2026.

Learning outcomes

After completing this course, participants will be able to:

- Build end-to-end analytics and ML workflows from raw data to evaluated models.
- Perform EDA, data profiling, preprocessing, and feature engineering with practical patterns that reduce downstream issues.
- Implement and compare supervised learning approaches (classification/regression) and select metrics that reflect real objectives.
- Avoid common real-world pitfalls (e.g., data leakage, improper splitting, misleading metrics) and document assumptions clearly.
- Use Apache Spark to scale preprocessing and feature engineering workflows.
- Train and run deep learning models using PyTorch, with practical GPU-aware considerations.
- Understand responsible AI principles and key ethical/legal considerations relevant to real-world usage.
- Gain exposure to emerging HPC/AI trends relevant to modern data and ML workloads.

Instructors

Junghwa Lee (HLRS) [junghwa.lee\(at\)hlrs.de](mailto:junghwa.lee(at)hlrs.de), and

Dr. -Ing. Lorenzo Zanon (HLRS) [lorenzo.zanon\(at\)hlrs.de](mailto:lorenzo.zanon(at)hlrs.de)

Registration closes on Sunday, August 23, 2026.

A.28 Technical responsibility and ethical methods

Keywords: AI

Course page: <https://www.hlr.de/de/event/technical-responsibility-and-ethical-methods>

Responsibility is probably the central concept of technology ethics. In any case, it is one of the most prominent buzzwords when it comes to regulation, technology assessment or ethical guidelines. Organisations such as ACM, IEEE and FEANI have published guidelines for responsible behaviour by engineers. However, responsibility cannot be reduced to printed paper. Applied technology ethics develops concepts and methods for analysing and resolving ethical conflicts. These include specifics and principles of technical responsibility such as the principle of condition preservation, the understanding of collective long-term responsibility, and the distinction between six relata for ethical responsibility analysis. These concepts will be explained in the lecture using selected examples, and insights into the methodological practice of professional engineering ethics will be offered.

Philosopher of technology Michael Funk (U Vienna) will join us on February 20th, 4pm for an introductory lecture on Technical responsibility and ethical methods. This lecture is the second part of a two part lecture (first part on February 19th) especially aimed at computer scientists, engineers and technologists.

Join us online at: <https://unistuttgart.webex.com/unistuttgart/j.php?MTID=m6128876942256d6b35c1f595d4b776c6>

A.29 Time Series Forecasting

Keywords: AI, Deep Learning

Course page: <https://www.hlr.de/training/2025/SCA-TSF/>

Time series forecasting has gained particular attention in recent years because it can provide a fundamental insight into the future by analyzing historical data, assuming that future trends are similar to those that occurred in the past. The primary value of using time series for forecasting is that at business time, future trends and outcomes are not fully available and can only be estimated using historical forecasts. This technology is vital in many fields to understand current seasonal trends and fit their data to the flow of time. Time series forecasting has important applications in weather forecasting, business, healthcare, and finance.

A.30 Using the GWDG Data Pools for Scientific Data Sharing

Keywords: Data Analysis, I/O, Linux

Course page: <https://academy.gwdg.de/p/event.xhtml?id=6916f8dd5067a75>

055a4c0d5

The scientific community relies heavily on the sharing of standardized datasets, such as ImageNet or Sentinel-2 imagery. To host these popular datasets in a central store, the GWDG offers the Data Pools service. Compared to conventional cloud-based approaches, we achieve significantly higher performance with Data Pools when running on our HPC systems. Additionally, the GWDG provides a number of standard datasets and derived data products, such as machine learning models. This service is not only for users to consume data but also allows them to share and host their own versioned datasets within our HPC systems. Other users of our systems can then use your dataset or data products to conduct their own research. Data Pools are specifically designed for the scientific community, providing versioned datasets that are citable.

Within this course, we will teach you how to discover existing Data Pools and how to publish your own dataset as a Data Pool to share it with others.

Learning goal

- Understand the concept of data pools
- Learn how to discover existing data pools
- Learn to publish your own dataset as data pools

Requirements

Basic Linux and HPC experience