



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

D6.1

Success Stories v1

WP6: Engagement with the community, networking, collaboration,
communication and outreach

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**

The acquisition and operation of the EuroHPC AI-optimised supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Digital Europe Programme, as well as by the German Federal Ministry of Research, Technology and Space (BMFTR) and the Baden-Württemberg Ministry of Science, Research and the Arts.

Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the EuroHPC JU. This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101234027. It 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	6	Engagement with the community, networking, collaboration, communication and outreach
Author	Sarah Heinisch	USTUTT/HLRS
Approved by	MGT Board	30/09/2026
Reviewer 1	Andreas Wierse	SICOS BW GmbH
Reviewer 2	Michael Walther	BADW-LRZ
Dissemination Level	PU - Public	

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
USTUTT/HLRS	09/09/2026	Initial Version	0.1
BADW-LRZ	14/09/2026	Review Feedback	0.2
SICOS	16/09/2026	Review Feedback	0.3
USTUTT/HLRS	17/09/2026	Adjusted Version ready for MGT Review	1.0
USTUTT/HLRS	29/09/2026	MGT review done	1.1

Executive Summary

This deliverable comprises the collection of user and success stories, their publication in a dedicated communication booklet found in Annex I and on the project website, as well as their dissemination through HammerHAI's communication channels. Together, they constitute the core communication product of Deliverable 6.1, while this document provides the context and approach, outlining the editorial and publication approach, story collection framework and design approach that guided the development of the stories. It explains the broadening of the initial scope of Deliverable 6.1, the selection of contributors of the stories and distinction between the two types of stories. It further presents the editorial, publication and design approaches, and guidelines and story templates supporting the consistent preparation and presentation of each story.

The collection brings together eight stories from six contributing institutions. The success stories highlight demonstrated benefits and impact, illustrating practical and tangible value generated working with HammerHAI and its ecosystem, while the user stories focus on intended value and potential impact envisioned by the contributor through their work and support from HammerHAI.

Taken together, the resulting communication product brings these stories together in a coherent format for publication and dissemination across HammerHAI's communication channels to create awareness and encourage potential users to explore the benefits of HammerHAI. Additionally, it serves a basis for the further development of the upcoming Deliverable 6.2 Success Stories v2, which will build on the stories presented and engage with future contributors.

List of Abbreviations

AI	Artificial Intelligence
CAD	Computer-Aided Design
DPI	Dots Per Inch
KPI	Key Performance Indicator
RAG	Retrieval-augmented generation
SME	Small and Medium-sized Enterprises
WP	Work Package

List of Tables

<i>Table 1: Overview of user and success stories.....</i>	10
<i>Table 2: User Story Template</i>	12
<i>Table 3: Success Story Template.....</i>	14

List of Figures

Figure 1: Excerpts from the booklet illustrating the design, visual layout and general information. ... 12

Table of Contents

Document Information.....	2
Document History	2
Executive Summary	3
List of Abbreviations.....	4
List of Tables.....	5
List of Figures.....	6
Table of Contents	7
1 Introduction	8
1.1 Scope Extension.....	8
2 Editorial and Publication Approach	8
2.1 Overview of Story Types.....	9
2.2 Contributors User and Success Stories.....	10
2.3 Design Approach	11
3 Editorial Guidelines and Story Collection Templates.....	12
3.1 User Story Template.....	12
3.2 Success Story Template.....	13
4 Conclusion.....	15
References.....	15
Appendix.....	16

1 Introduction

Deliverable 6.1 – Success Stories v1 – is a communication product developed as part of HammerHAI’s communication and dissemination activities in Work Package 6 (WP). Its purpose is to showcase its contribution to accelerating the adoption of artificial intelligence (AI) in European startups, SMEs, industry and academia by highlighting concrete stakeholder experiences, and tangible achievements. As such, this document provides the conceptual and editorial structure of the presentation of success and user stories, the publication approach, and the story collection framework that guided the production of the stories and printed booklet. The stories are presented in the communication product in the Annex. They are intended for HammerHAI’s key industrial target audience, including start-ups, SMEs, industrial stakeholders, and research.

The featured stories show how HammerHAI supports the uptake and application of AI in manufacturing and engineering sectors. This strengthens HammerHAI’s visibility, credibility, and relevance by showcasing concrete implementation examples and stakeholder outcomes. In doing so, Deliverable 6.1 supports outreach activities and encourages further engagement with HammerHAI’s services and collaboration opportunities. The approach presented in this deliverable is in line with the WP description and risk mitigation strategies set out in Annex to the Grant Agreement, as well as with the communication channels and target groups identified in Part B [1].

1.1 Scope Extension

During the development of the stories, the HammerHAI system was still in its procurement phase. While the original assumptions anticipated an alignment between the infrastructure and service project timelines, this created an opportunity to broaden the scope beyond success stories by also including user stories as additional content. During the conceptual process, it became clear that the content focus needs to be broadened to create additional value. While the original concept centered on success stories and presenting concrete results and outcomes, the inclusion of user stories was identified as a complementary development to highlight use cases and solutions still being explored or developed with support from HammerHAI; it helps to provide a more comprehensive perspective.

2 Editorial and Publication Approach

As Deliverable 6.1 is as a communication and dissemination product, it is published both as digital and printed versions, ensuring consistency across communication channels and facilitating engagement with stakeholder groups (see Annex Booklet PDF).

Dissemination and Publication Channels

The dissemination approach includes publication on the HammerHAI website [2] and promotion through HammerHAI communication channels, including social media (LinkedIn [3] and BlueSky [4]), and events. The stories as well as the digital booklet will be available on the project website at www.hammerhai.eu. Printed booklets support stakeholder engagement at events. Content on the HammerHAI website is updated to support continuous dissemination throughout the project duration, particularly in view of Deliverable 6.2 (Success Stories v2). Story dissemination includes:

- **Online Website version:** To be updated where applicable, e.g., when a user story develops into a success story. The final presentation and categorization of stories on the project website are determined during the implementation.
- **Printed Booklet version:** full-length versions in print booklet (double-page format, estimated 17x24 cm). To be disseminated at events, such as exhibitions, visits to HammerHAI on-site, the system inauguration and other events focused on HammerHAI's target groups.
- **Social media:** dissemination of stories with links to the website as the main entry point for engagement of users to use offered services.

Content has been collected using a full story template (see 3.1 User Story Template and 0 Success Story Template). HammerHAI adapts highlights, teasers, and versions for web publication in the editorial process. By submitting a story, the organizations provided content that HammerHAI may edit, format, and publish, including any provided visuals, across communication channels (including the website, booklet, presentations, and social media channels such as LinkedIn and BlueSky).

2.1 Overview of Story Types

User Stories

User stories capture the goals, needs, and expectations of project users. They provide insight into what users aim to achieve and the type of support required from HammerHAI. The purpose is to document user priorities and requirements in order to inform and reflect HammerHAI's relevance from the user perspective.

Success Stories

Success stories document achieved outcomes resulting from engagement with HammerHAI. They highlight concrete results and demonstrated benefits, showing the value generated for users and organizations through the application of HammerHAI's services and support activities.

Contributor Guidelines

Contributions were developed with the core users in mind, targeting start-ups, SMEs, and industrial stakeholders, particularly in the manufacturing and engineering sectors. Technical content was considered welcome but should be explained and presented in a way that is understandable and relevant for non-specialist industrial audiences. Where possible, visual elements such as graphics or diagrams were included to support key messages. All visual materials and logos, where applicable, were requested in print-quality resolution (minimum 300 DPI). Contributions were asked to be submitted in American English, with both English and German versions provided for digital versions. Where no German version was provided, the editorial team at HLRS will create a German version later for use in the online versions.

2.2 Contributors User and Success Stories

Story contributors have been identified in close collaboration with WP1 and project partners during the editorial guideline phase and expanded as more users have engaged with HammerHAI for support. To ensure a meaningful representation, eight user and success stories have been developed for the deliverable deadline. Further potential contributors have already been identified for possible future stories, supporting the continued development and presentation of HammerHAI's results.

The collected stories showcase how AI and high-performance infrastructure are being explored to address challenges across industry and engineering. They cover applications ranging from real-time shopfloor intelligence, robotics, engineering design to autonomous driving, construction, and AI-assisted access to technical knowledge. They bring together both, emerging use cases being explored within HammerHAI and concrete, successful applications for potential industrial and economic benefits of applying AI.

Table 1 below provides an overview of the institutions contributing stories that were developed with support from HammerHAI. The full stories are presented in the communication product provided as a booklet in PDF format in the Annex.

Table 1: Overview of user and success stories

Contributor	Story Type	Story Title
ARENA2036 [5]	Success story	Accelerating Real-Time Shopfloor Intelligence with HammerHAI
ARENA2036	Success story	Accelerating AI Robot Planning with Large Language Models

ARENA2036	Success story	GenCAD – Foundation Model for Engineering Design Generation
Bosch BGS [6]	Success story	Securing AI across the Machine Learning Operations chain: Piloting cybersecurity and data sovereignty of AI in supercomputing
MPA Universität Stuttgart [7]	Success story	Testing a RAG Prototype with Real Engineering Questions
cogniBIT GmbH [8]	User story	Bridging Simulation and Reality for Safer Autonomous Driving
Ramblr GmbH [9]	User story	JAWS: Anticipating Industrial Operating Conditions from Point-of-View Video
flinq (TTI GmbH, Universität Stuttgart) [10]	User story	Training Europe’s Construction Embedding AI with HammerHAI

The booklet brings together user and success stories developed within HammerHAI and starts with a front cover, followed by editorial information and a funding disclaimer with respective logos of funding bodies of HammerHAI. The booklet is structured to provide readers with a concise introduction to HammerHAI and its key services in an “About” chapter. A section on the project consortium presents each partner through their logo and a brief description. The User and success stories form the core of the booklet and present the experience, outcomes, and intended value emerging from being a HammerHAI user. The booklet additionally includes a list of abbreviations and a section on further information with access to communication channels via QR codes (website, LinkedIn, BlueSky, YouTube [11], and email).

2.3 Design Approach

The design of the booklet has been developed in cooperation with SICOS BW and in alignment with HammerHAI’s established visual and brand identity. This includes consistency with HammerHAI’s color scheme, typography, logo usage, and other branding elements already applied across project communication materials, including the website. A style guide outlining these elements, as well as the required funding acknowledgement, is internally available to all project partners to support consistent and coherent communication across project-related materials. An example of the resulting booklet design and visual layout is provided below. Each

submitted story was individually laid out in the design, as the provided text lengths and graphics varied throughout contributors.



Figure 1: Excerpts from the booklet illustrating the design, visual layout and general information.

3 Editorial Guidelines and Story Collection Templates

As part of the editorial preparation, structured content collection templates have been established and shared with WP Leads and WP6 partners. The templates defined the proposed layout and content requirements for success stories and user stories and served as the basis for collecting content from contributors. They provided guidance on the expected narrative structure, key information to be supplied, and editorial requirements, thereby supporting a coherent and consistent presentation of project achievements. The templates will also serve as a basis for the collection and development of future stories.

3.1 User Story Template

This template defines the structure and content requirements for user stories to be submitted by contributors for inclusion in the booklet.

Table 2: User Story Template

1. Catchy title (about 50-80 Characters, incl. spaces)
How would you describe the goal of working with HammerHAI in one sentence?
2. User Goals

Short bullet list (max. 3 lines total). Keep this section concise and impactful. What are the core business, scientific, or operational goals? What do you aim to achieve in the short to medium term? Why are AI or data-driven methods relevant to these goals?
3. Current State and Key Challenges (1000-1100 characters, incl. spaces)
Describe your organization (e.g., startup, SME, research group, industrial partner) and your current role in the sector. What workflows, tools, or methodologies are currently in use? What data sources and infrastructure are available? Why is the challenge relevant for industrial users? Explain where limitations arise in practice and what technical, scientific, operational, or data-related challenges (e.g., quality, access, interoperability), skills gaps, scalability constraints, or integration barriers restrict further progress. Why do these require support from HammerHAI?
4. Support from HammerHAI (300-400 characters, incl. spaces)
This describes your expectations, not the outcomes. What type of support do you seek from HammerHAI? What are desired capabilities (e.g. AI models, data pipelines, validation, deployment support)? Are there collaboration needs (expertise, infrastructure, ecosystem access)?
5. Intended Value and Impact (400-550 characters, incl. spaces)
This explains the targets and aspirations, not measured results. What are e.g., intended improvements in performance or efficiency? Are there expected qualitative benefits (better decisions, faster development, increased reliability)? If available: Are there indicative or aspirational targets? What is the potential impact you are aiming for?
6. Current Status within HammerHAI and Next Steps (300–400 characters, incl. spaces)
As this may become outdated quickly, this section will only be included in the online version, and not the printed version. This shows the engagement status with HammerHAI. What is the stage of involvement (e.g. onboarding, requirements analysis, pilot preparation)? What are activities already initiated or completed? Are there open questions, dependencies, or constraints? This shows planned activities or milestones within the project. What are decision points or evaluation phases? Are there dependencies, risks, or prerequisites?
7. Contact
Name, Institution, Email.

3.2 Success Story Template

This template defines the structure and content requirements for success stories to be submitted by contributors for inclusion in the booklet.

Table 3: Success Story Template

1. Catchy Title (about 50-80 Characters, incl. spaces)
How would you describe the success of your engagement with HammerHAI in one sentence? E.g., "From Data Chaos to Insight: AI Cuts Production Time by 15 % in Automotive Manufacturing"
2. Highlights
Short bullet list (max. 3 lines total). Keep this section concise and impactful.
1. E.g., key technical AI solutions applied
2. E.g., improved KPIs that demonstrate the solutions success or benefit
3. E.g., integration of the AI solution into existing industrial workflows or systems
3. Challenge (800-1000 characters, incl. spaces)
What type of organization are you (startup, SME, research group, industrial partner, etc.)? Describe the core technical or industrial challenge addressed. What were fundamental challenges or bottlenecks before the solution was implemented (e.g., energy inefficiency, data management)? Why is this important for industrial users of AI or supercomputing infrastructure and other organizations? How did the existing challenges affect e.g., industrial processes, energy, computing time or costs and the business impact?
4. Solution (600-750 characters, incl. spaces)
Summarize the approach or innovation used to solve the challenge. What kind of AI technologies or tools were developed to address the challenge? How was the outcome different or better than your starting point? How did you overcome previous challenges? How was the solution e.g., integrated into existing industrial workflows or was the solution customized? How was the solution deployed?
5. Benefits & Impact (800-1000 characters, incl. spaces)
Outline the concrete results and the wider impact. Include performance gains (numbers), societal or environmental relevance, industry or scientific benefit, etc. What were measurable benefits after the solution was implemented (e.g., cost savings, time reduction, improved decision-making, efficiency in operations, energy savings)?
6. Optional: Next Steps (approx. 350–450 characters, incl. spaces)
Optional: What are ongoing works, future applications, or plans for follow-ups? Are there next steps to scale or optimize the solution?
7. Organizations Involved
List the key participating institutions (optional logos will be added in the layout phase).
8. Contact
Name, Institution, Email.

4 Conclusion

Together, the user and success stories demonstrate the practical value and potential impact of HammerHAI's offers and the ecosystem in different applications.

The success stories demonstrate that HammerHAI can deliver tangible benefits to industry, especially in the engineering and manufacturing sectors. Particularly through improved efficiency, scalability of AI solutions, and improved access to specialized engineering knowledge. These include decreased training time from six to eight hours down to one hour, reduced inference latency from 9.03 seconds to 3.98 seconds, enabling accelerated development and training in CAD, as well as enabling federated learning without exchanging sensitive data to support secure industrial AI development through access to infrastructure. Additionally, HammerHAI's roadmapping workshop enabled the development of a RAG prototype. Together, they illustrate the real-world benefits that HammerHAI's expertise, infrastructure, and support can bring to technical and industrial challenges.

The user stories highlight the intended value and potential impact of solutions that are being developed with support from HammerHAI towards future industrial AI applications. They point to opportunities how AI can address practical challenges across autonomous driving, connecting fragmented construction data, and help robotics and manufacturing systems recognize production situations and respond reliably to change. These user stories highlight how the AI solutions could create practical benefits for industry as they are further developed and applied.

As engagement with HammerHAI users continues to grow, further potential contributors are being identified, which supports the ongoing development of additional user and success stories. The stories provide a foundation for continued dissemination activities within the project and will be further built upon in following communication outputs and the upcoming Deliverable 6.2 of WP6.

References

- [1] European High Performance Joint Undertaking. Grant Agreement — HammerHAI, Grant Agreement No. 101234027. 2025.
- [2] HammerHAI. Project website. <https://www.hammerhai.eu/>. Last accessed: 8 September 2026.
- [3] HammerHAI. LinkedIn. <https://de.linkedin.com/company/hammerhai>. Last accessed: 8 September 2026.
- [4] HammerHAI. BlueSky. <https://bsky.app/profile/hammerhai-eu.bsky.social>. Last accessed: 8 September 2026.
- [5] ARENA2036. <https://arena2036.de/>. Last accessed: 8 September 2026. Last accessed: 8 September 2026.

[6] Bosch BGS. <https://www.bosch-softwaretechnologies.com/en/locations/germany/>. Last accessed: 8 September 2026.

[7] MPA Universität Stuttgart. <https://www.mpa.uni-stuttgart.de/>. Last accessed: 8 September 2026.

[8] cogniBIT GmbH. <https://www.cognibit.ai/>. Last accessed: 8 September 2026. Last accessed: 8 September 2026.

[9] Ramblr GmbH. <https://www.ramblr.ai/>. Last accessed: 8 September 2026. Last accessed: 8 September 2026.

[10] flinq (TTI GmbH, Universität Stuttgart). <https://flinq.ai/>. Last accessed: 8 September 2026. Last accessed: 8 September 2026.

[11] HammerHAI. YouTube Channel. https://www.youtube.com/@HammerHAI_eu. Last accessed: 17 September 2026.

Appendix

The full story booklet PDF can be found in the Annex.

**1st Digital Edition
2026**



HammerHAI

Germany's AI Factory for Manufacturing,
Engineering, and Research

First Edition of User and Success Stories



HammerHAI Success Stories
1st Digital Edition

Publisher
HLRS on behalf of the HammerHAI Consortium
Nobelstraße 19
70569 Stuttgart, Germany
info@hammerhai.eu

Published
September 2026

Copyright
© 2026 Members of the HammerHAI Consortium. All rights reserved.

Editor:
HLRS on behalf of the HammerHAI Consortium

Content contributions
Success stories, images, and logos provided by members of the HammerHAI Consortium and participating users.

The acquisition and operation of the EuroHPC AI-optimised supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Digital Europe Programme, as well as by the German Federal Ministry of Research, Technology and Space (BMFTR) and the Baden-Württemberg Ministry of Science, Research and the Arts.

Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the EuroHPC JU. This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101234027. It 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.

HammerHAI
Germany's AI Factory for
Manufacturing, Engineering, and Research



Index

About us 6

The Consortium 7

Success Stories

Accelerating Real-Time Shopfloor Intelligence with HammerHAI..... 8

Accelerating AI Robot Planning with Large Language Models 10

GenCAD – Foundation Model for Engineering Design Generation 12

Securing AI across the Machine Learning Operations chain: Piloting cybersecurity and data sovereignty of AI in supercomputing..... 14

Testing a RAG Prototype with Real Engineering Questions 16

User Stories

Bridging Simulation and Reality for Safer Autonomous Driving..... 18

JAWS: Anticipating Industrial Operating Conditions from Point-of-View Video..... 20

Training Europe’s Construction Embedding AI with HammerHAI..... 22

List of Abbreviations 24

Further Information 25

About us

Established by the EuroHPC Joint Undertaking under its AI Factories initiative, HammerHAI is helping to accelerate the artificial intelligence revolution in Europe. Supporting start-ups, SMEs, industrial companies, and academic research – particularly in manufacturing and engineering fields – HammerHAI provides a comprehensive portfolio of AI computing infrastructure, services, and tools. These resources are empowering organizations to more quickly seize the opportunities that emerging AI technologies now offer.

Going into service in early 2027, the HammerHAI supercomputer is a cost-efficient option for European organizations whose AI computing needs exceed their in-house capabilities. With up to 15 exaflops of AI inference performance and cloud-like usability, it can support medium- to large-scale workloads for model training, inference, and data analytics. Located at the High-Performance Computing Center Stuttgart, it is operated securely and in compliance with European data protection regulations.

HammerHAI also offers professional training courses, consulting and strategic planning services, software tools, and data resources, as well as personalized support from our team of AI experts. Regardless of the maturity of an individual application, HammerHAI can provide scalable solutions, addressing challenges at all stages of the AI lifecycle.

The following User and Success Stories highlight innovative applications of artificial intelligence, machine learning, and data analytics being developed with HammerHAI's support. The stories give insights into our users' needs and experiences, including their goals, needs, and progress achieved. Together, they illustrate the value and impact of HammerHAI's resources and services.



The Consortium



Leads HammerHAI and will host the AI-optimized supercomputer, including computing and storage resources, software, and secure access.

One of Europe's leading research institutions, contributing expertise in trustworthy AI, data-driven methods, and innovation for SMEs.

Supports companies, especially SMEs, in adopting simulation, HPC, data analytics, and AI through consulting, training, and technology transfer.

Provides HPC services and AI infrastructure, with a strong focus on secure AI solutions for research, SMEs, and startups.

A leading German supercomputing centre providing supercomputing resources, AI expertise, user support, and training across disciplines for research and science.

Accelerating Real-Time Shopfloor Intelligence with HammerHAI

Highlights

- Reduced model training time from 6–8 hours to less than one hour.
- Supported AI applications in manufacturing, robotics, and digital twins.
- Enabled industrial testing of AI-based shopfloor intelligence.

Challenges

ARENA2036 stands for „Active Research Environment for the Next generation of Automobiles“ and is one of nine research campuses of the „Research Campus - Public-Private Partnership for Innovation“ funding initiative in Germany. ARENA2036 is supported by the Federal Ministry of Research, Technology and Space and is run as a registered association with members from science and industry. These members are active in various disciplines – from the automotive industry to aerospace technology, textile and materials research to ergonomics.

Modern manufacturing environments generate vast amounts of spatial data that must be processed in real time to support robotics, autonomous systems, and digital twins. At ARENA2036, we developed a deep-learning-based Shopfloor Model capable of mapping, segmenting, and classifying industrial environments. However, training these large neural networks on local GPU workstations required overnight computation (6–8 hours), significantly slowing research iterations and limiting rapid model optimization.

Solution

Through the HammerHAI project, ARENA2036 gained **access to the HLRS Hunter supercomputer, enabling the migration of our training pipeline** from sequential local execution to a highly optimized parallel multi-node supercomputing environment. Hunter's scalable


computing resources accelerated deep learning workflows while maintaining reproducible training across large datasets. HammerHAI provided the computational foundation that made this acceleration possible.

Benefits and Impact

The accelerated Shopfloor Model now serves as a key building block for intelligent manufacturing, robotic navigation, and digital twin applications. Through AI-MATTERS, these technologies are being transferred into industrial testing and experimentation services, allowing companies to evaluate AI-based shopfloor intelligence in realistic production environments.

one hour, enabling significantly faster experimentation and validation. Through HammerHAI, companies can access the required HPC capabilities without investing upfront in dedicated computing infrastructure. Combined with AI-MATTERS, the solution can be tested in realistic manufacturing environments, reducing the technical and financial risks of industrial AI adoption.

Using HLRS Hunter, model training time was reduced from 6–8 hours to less than



Organizations Involved



Contact

Muhammad Saeed, ARENA2036, muhammad.saeed@arena2036.de

Accelerating AI Robot Planning with Large Language Models

Highlights

- Reduced average inference latency from 9.03 seconds to 3.98 seconds.
- Achieved more than 2× higher throughput.
- Enabled benchmarking and scalable testing of open-source LLMs for industrial robotics.

Challenges

ARENA2036 stands for „Active Research Environment for the Next generation of Automobiles“ and is one of nine research campuses of the „Research Campus - Public-Private Partnership for Innovation“ funding initiative in Germany. ARENA2036 is supported by the Federal Ministry of Research, Technology and Space and is run as a registered association with members from science and industry. These members are active in various disciplines – from the automotive industry to aero-

space technology, textile and materials research to ergonomics.

Programming industrial robots traditionally relies on manually designed rule-based systems that are difficult to scale. Large Language Models offer a new paradigm by translating natural-language instructions into executable robot actions, but evaluating hundreds of planning tasks requires substantial computational resources.

Solution



Through HammerHAI, ARENA2036 leveraged the HLRS Hunter supercomputer to benchmark an LLM-based robot task planning framework using DeepSeek-Coder. Hunter enabled large-scale evaluation across 90 manipulation tasks comprising 450 robot actions, providing the computational performance required for systematic benchmarking and optimization. HammerHAI provided the computational infrastructure that transformed experimental AI concepts into scalable industrial solutions.



Benefits and Impact

Using HLRS Hunter, average inference latency was reduced from 9.03 s to 3.98 s, with more than 2× higher throughput across a benchmark of 90 tasks and 450 robot actions. HammerHAI enables scalable evaluation of LLM-based robotic

intelligence without requiring companies to build dedicated AI computing infrastructure. Through AI-MATTERS, these approaches can be tested in realistic industrial robotics environments before deployment.



Organizations Involved



Contact

Muhammad Saeed, ARENA2036, muhammad.saeed@arena2036.de

GenCAD – Foundation Model for Engineering Design Generation

Highlights

- Accelerated development and training of a generative AI framework for Computer-Aided Design (CAD).
- Enabled training on large engineering datasets.
- Supported model optimization and robustness evaluation.

Challenges

ARENA2036 stands for „Active Research Environment for the Next generation of Automobiles“ and is one of nine research campuses of the „Research Campus - Public-Private Partnership for Innovation“ funding initiative in Germany. ARENA2036 is supported by the Federal Ministry of Research, Technology and Space and is run as a registered association with members from science and industry. These members are active in various disciplines – from the automotive industry to


aerospace technology, textile and materials research to ergonomics.

Generative AI is rapidly transforming engineering design by enabling AI models to automatically generate, understand, and optimize CAD models from limited inputs such as sketches, point clouds, or natural-language descriptions. Training these large foundation models requires processing millions of geometric primitives and complex 3D representations, resulting in computational demands that exceed con-

Solution

Through the HammerHAI project, ARENA2036 utilized the High-Performance Computing Center Stuttgart (HLRS) **Hunter supercomputer to accelerate the development and training of GenCAD**, a generative AI framework for computer-aided design. Hunter's scalable computing resources enabled efficient training of deep neural networks on large engineering datasets while significantly reducing model development time. The supercomputing environment also allowed extensive experimentation with


model architectures, hyperparameter optimization, and large-scale benchmarking that would have been impractical using local GPU resources. HammerHAI transformed what had previously been resource-constrained experimentation into a scalable AI development workflow and provided the computational backbone that enabled this next generation of engineering AI research.



Benefits and Impact

GenCAD represents an important step toward AI-assisted engineering design and digital product development. The technology has strong potential for automated CAD generation, design optimization, digital engineering, and future industrial AI assistants. Through AI-MATTERS, these capabilities can be demonstrated and validated with industrial partners, accelerating the adoption of generative AI in manufacturing and product development. HammerHAI provided the computational backbone that enabled this next generation of engineering AI research.

Using HLRS Hunter, GenCAD training became more scalable, supporting larger engineering datasets and systematic evaluation of different model and design-generation strategies. HammerHAI enables companies to explore compute-intensive Generative AI for engineering without making major upfront investments in dedicated HPC infrastructure. Through AI-MATTERS, resulting Proofs of Concept can be tested with industrial partners before larger deployment decisions.



Organizations Involved



Contact

Muhammad Saeed, ARENA2036, muhammad.saeed@arena2036.de

Securing AI across the Machine Learning Operations chain: Piloting cybersecurity and data sovereignty of AI in supercomputing


Highlights

- Federated learning applied to steering angle prediction reduced the need for data transfers.
- The federated steering-angle model reduced estimated data network transfer approximately by 60%.
- Continuous model evaluation integrated as a “verifier-as-a-client” architecture.

Challenges

During industrial partnerships, cross-organizational AI development often depends on large, distributed and sensitive datasets that cannot be freely transferred to a central supercomputing environment because of privacy, security, regulatory and bandwidth constraints. Centralized training typically requires

moving substantial volumes of raw data, increasing network load, storage demand, computing time and operational cost, while also exposing AI workloads and models to additional cybersecurity risks. This can limit collaboration between organizations and slow the validation of AI models on supercomputing infrastructure.



Solution

Bosch Global Software Technologies GmbH (BGSG), located in Stuttgart, Germany, develops software, digital systems, and IT solutions for enterprises. Working with HammerHAI, BGSG tested an approach that combines federated learning and secure model validation to train automotive AI models without the need to exchange raw vehicle datasets.

Distributed clients trained local steering-angle models using the Flower open-source federated learning framework. No data were exchanged among clients except for model parameters, which were

aggregated into a global model. HLRS GPU infrastructure provided through HammerHAI was integrated into a customized validation workflow, which included a “verifier-as-a-client” concept for evaluating global models during training. The AI security platform “AI Shield” was deployed in the HPC environment to strengthen protection of AI workloads. The approach enabled collaborative model development among data owners and sources and created a reusable edge-to-HPC workflow for secure industrial AI development and validation.

Benefits and Impact

The project demonstrated that federated learning can combine knowledge from distributed automotive datasets without requiring the transfer of the underlying raw data.

Compared to a centralized approach, the federated steering-angle model prediction demonstrated significantly better generalization across operating environments, without a huge drop in accuracy. Estimated network traffic was reduced by approximately 60%, lowering bandwidth, storage and data-management requirements. Integration with the HLRS supercomputing environment established a reusable workflow for GPU-based model

validation by internal teams, while AI Shield extended the approach toward securing AI workloads on HPC infrastructure.

This solution offers a scalable foundation that will make it easier for industrial partners to collaborate on the development and validation of secure AI models without exposing proprietary or sensitive data. It holds the additional benefit that reduced data movement can also decrease infrastructure load and associated energy consumption. Direct energy or monetary savings were, however, not measured in this project.

Contact

Rane Abhijeet, Bosch Global Software Technologies GmbH,
Abhijeet.Rane@in.bosch.com

Testing a RAG Prototype with Real Engineering Questions

Highlights

- Developed a retrieval-augmented generation (RAG) prototype
- RAG prototype uses 7,282 MPA publications on welding and joining
- MPA engineers write benchmark questions and evaluate answers and sources

Challenges


The Material Testing Institute (MPA) at the University of Stuttgart combines materials testing, joining research, and industrial transfer. Its knowledge is spread across publications, material tables, formulas, standards, and expert experience. A retrieval-augmented generation (RAG) system — a large language model (LLM) that retrieves and uses information from external knowledge sources — could improve access to this information. However, fluent answers alone do not

guarantee engineering correctness. Relevant sources must be retrieved and cited traceably, calculations must be reliable, and uncertainties must remain transparent. Earlier work by the High-Performance Computing Center Stuttgart (HLRS) and the University of Stuttgart's Institute for Metal Forming Technology (IFU) showed that language models can partially predict temperature-dependent yield curves.

Solution

The collaboration with HammerHAI began with a **roadmapping workshop** that identified a benchmark test for a RAG system and initiated the development of a RAG prototype. **MPA and HLRS ran benchmark tests** to see where the RAG prototype helps, where it fails, and how much expert review is required before industrial use could be considered. A living roadmap and weekly meetings now guide prototype reviews and planned tests. The RAG testbed works with 7,282 MPA Stuttgart publications on welding and joining. Its interface allows experiments with source counts,

multi-step retrieval, source, page-level references, PDF links, calculations, plots, and a second verification pass. These are prototype options, not yet integrated or validated production features. MPA engineers develop representative questions and evaluate the answers, cited sources, and calculations using the benchmark tests developed by HammerHAI and MPA.



Benefits and Impact

The main result is an evaluation approach and early RAG prototype. The benchmark records assessments conducted by MPA engineers on technical correctness, source traceability, and reproducibility. An exploratory comparison of 11 openly available models illustrates the need for systemic evaluation. This solution has the potential to support future AI-assisted

engineering workflows by improving access to specialized engineering knowledge, reducing information retrieval efforts, and supporting the selection of welding parameters to accelerate the calibration of welding machines.

Organizations Involved

Contact

Gerrit Koch, Material Testing Institute (MPA), University of Stuttgart, Gerrit.Koch@mpa.uni-stuttgart.de

Bridging Simulation and Reality for Safer Autonomous Driving

User Goals

- Convert driving simulations into photorealistic video.
- Scale synthetic data generation to support verification and validation of autonomous vehicles.

Current State and Key Challenges

cogniBIT is a Munich-based startup that develops cognitive traffic agents for the validation of automated driving and driver-assistance systems. Our software, cogniBOT, creates controllable and repeatable scenarios involving vehicles and other road users. The simulated agents behave based on insights from neuroscience, which enables them to model human perception, decision-making, and behavior more realistically than conventional traffic agents while also ensuring that cogniBOT remains computationally tractable for large-scale simulations. The scenarios are independent of any visualization platform.

In the project for which we use HammerHAI resources, we use the open-source autonomous driving simulation software CARLA to visualize our scenarios. We export synchronized RGB, depth, edge, and semantic-segmentation data. We then evaluate generative video models that transform CGI outputs into photorealistic dashcam-style footage while preserving the logic of the scenario being investigated. This allows us to obtain photorealistic footage of any scenario we wish, including

the ones including safety-critical events. This addresses a key industry challenge, as safety-critical events are rare, costly, and difficult or impossible to reproduce in real-world testing. Photorealistic outputs could help autonomous vehicle (AV) developers test their perception and decision-making software against scenarios that might not occur during normal operational testing.

We require HammerHAI's computing infrastructure and AI expertise to overcome limitations of other available computing systems, including their limited computational power and insufficient GPU-memory capacity. Being able to use HammerHAI resources is essential for cogniBOT's ability to fulfil the objectives of this project, which is to obtain photorealistic videos from CGI inputs that maintain temporal consistency and visual fidelity.

Support from HammerHAI

Through HammerHAI, we seek **access to high-performance NVIDIA GPU infrastructure and expert support** in selecting, configuring, and evaluating generative video models. This work requires intensive computing resources that would be

prohibitively expensive for us at standard commercial rates. We currently use NVIDIA Omniverse and Cosmos, making NVIDIA-based systems particularly suitable.

Intended Value and Impact

Our goal is to design a reliable pipeline that transforms controllable simulated scenarios into realistic sensor videos while preserving the underlying ground truth. This would let AV developers test perception and downstream driving functions on diverse, repeatable, safety-critical situations without depending solely on

costly road recordings. Faster video generation and systematic quality metrics could shorten autonomous vehicle development cycles and improve confidence in simulation-based validation. This also addresses a very serious shortcoming of the EU-based AV companies.


Organizations Involved



Contact

Dr. Vahid S. Bokharaie, Chief Scientist at cogniBIT, vb@cognibit.ai

From Controllable Simulation to Photorealistic AV Test Data



Key requirement: improve visual realism without changing geometry, actors, or behavior

JAWS: Anticipating Industrial Operating Conditions from Point-of-View Video

User Goals

- Anticipate future scene graphs and process states from egocentric industrial video
- Fine-tune compact multimodal models that output validated, machine-readable results
- Enable real-time planning and decision-making for robotics and manufacturing

Current State and Key Challenges

Ramblr GmbH is a German AI startup that builds models that understand the physical world to support robotics and industrial production.

Our technology turns egocentric video (i.e., video shot from an individual subject's perspective) into structured representations of objects, attributes, spatial relations, activities, and process states. A new technology currently under development at Ramblr—called JAWS—goes further, anticipating plausible short-term futures as scene graphs in validated, machine-readable form. These graphs can then be used for downstream planning.

We train with a production-tested stack based on PyTorch FSDP, Ray, and containerized workflows for distributed GPUs. However, we currently face the challenge

of scaling. Fine-tuning multimodal models on industrial video data requires substantial computing power, fast connections between computing resources, reliable storage, and repeatable execution across multiple nodes. In addition, industrial data are not consistent across sites, processes, camera positions, and conditions. Models must learn robust representations while producing outputs in a format upon which machines can act.

For a startup, limited GPU access makes systematic experimentation and validation difficult. HammerHAI provides the infrastructure to move from smaller experiments to production-relevant training.

Support from HammerHAI

Through HammerHAI, we have been using **scalable NVIDIA GPU infrastructure, high-performance networking, and storage for distributed multimodal training at KISSKI (AI service center for critical and sensitive infrastructures)** in Göttingen. Technical onboarding and support

for efficient multi-node execution have helped us to validate our FSDP, Ray, and container-based pipeline at production-relevant scale.

Intended Value and Impact

With JAWS we aim to deliver a compact multimodal world model that predicts scene graphs and process states from industrial video in near real time, targeting inference within 200 milliseconds per frame. Machine-readable outputs will help robots and assistance systems recognize production situations, anticipate what happens next, and react more reliably to change. Our goals are to

shorten development cycles, improve robustness across industrial environments, and build a scalable foundation for AI-enabled manufacturing.

Contact

Philipp Schubert, Ramblr GmbH, philipp@rambl.ai



Training Europe's Construction Embedding AI with HammerHAI

User Goals

- Scale our embedding model into a full cross-modal foundation model for construction data.
- Prove that domain-trained, EU-models can outperform large general-purpose ones.
- Give the European Architecture, Engineering, and Construction sector a sovereign semantic infrastructure layer.

Current State and Key Challenges

flinq is a Stuttgart-based deep tech startup and is currently a technology spin-off (TGU) of TTI GmbH at the University of Stuttgart.

We are training a domain-specific, cross-modal embedding model for the European construction industry — an industry that generates enormous amounts of data every year, including tenders and an increasing number of Building Information Modeling (BIM) models and product catalogs.

As a cross-modal model, our embedding model processes many data types simultaneously, converting BIM models (IFC), tenders (GAEB), manufacturer data, construction site images, and other construction-specific formats into a common vector space. We make this available via an OpenAI-compatible API.

Our finely tuned 600-million-parameter Proof of Concept (PoC) model shows promising benchmark results compared to

OpenAI's "text-embedding-3-large" and the "Qwen3" model, which is 13 times larger.

Initial tests also show that 3D geometry and text can be embedded in the same space. This approach could address the challenge of integrating diverse construction data.

To date, flinq has run all training on a single rented GPU. For scalability and robustness, we require more computing power than we currently have, highlighting the need for access to HammerHAI's computing infrastructure.

Support from HammerHAI

We are looking for continuous access to GPUs that can scale with us as we transition from today's fine-tuning experiments to the significantly larger training runs for new models that lie ahead. In addition to **raw computing power, HammerHAI's expertise in efficient large-scale training**

and its connections to the broader European AI ecosystem are also of great value.

Intended Value and Impact

Using larger, fully cross-modal models, we aim to make flinq the semantic interoperability layer that the construction industry lacks, connecting BIM, tenders, products, and cost data. While the field moves toward horizontal, general-purpose models, we want to establish flinq as a dedicated embedding model for the construction industry and a trusted, vertical, EU-hosted solution. This will provide a

common foundation that will reduce the need for time-consuming manual data reconciliation.

Contact

Matthias Hornung, flinq (TTI GmbH, Universität Stuttgart), info@flinq.ai

List of Abbreviations

AEC	Architecture, Engineering, and Construction
AI	Artificial Intelligence
AV	Autonomous Vehicle
BIM	Building Information Modeling
CAD	Computer-Aided Design
CGI	Computer-Generated Imagery
EU	European Union
GPU	Graphics Processing Unit
GWDG	Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen
HLRS	High-Performance Computing Center Stuttgart
HPC	High-Performance Computing
KISSKI	KI-Servicezentrum für sensible und kritische Infrastrukturen
KIT	Karlsruhe Institute of Technology
LLM	Large Language Model
LRZ	Leibniz Supercomputing Centre
PoC	Proof of Concept
RAG	Retrieval-Augmented Generation
RGB	Red, Green, Blue
SME	Small and Medium-sized Enterprises

Further Information

Learn more about HammerHAI, its activities, and latest updates:



Project website: www.hammerhai.eu
Contact: info@hammerhai.eu

Follow us:



Funded by:

The acquisition and operation of the EuroHPC AI-optimised supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Digital Europe Programme, as well as by the German Federal Ministry of Research, Technology and Space (BMFTR) and the Baden-Württemberg Ministry of Science, Research and the Arts.

Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the EuroHPC JU. This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101234027. It 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.



