




BEHIND THE CODE

Computer Vision

Smart Vision for a Smarter World

Dr. Natália S. Teixeira Silva





Computer Vision Insights

Smart Vision for a Smarter World

WHAT IS IN THIS MONTH'S ISSUE:


- An Introduction to the **Computer Vision world** and the Software Campus **projects and partners** contributing to the field
- **Deep Dive:** navigate our Computer Vision projects and their journeys into innovation and applied research
- Computer Vision Innovation: our **Industry Partners'** insights in the era of smart environments and tools
- The Computer Vision **Community of Practice:** a founder story full of purpose
- Conclusions and Remarks

In recent years, Computer Vision has evolved from a specialized research field into a tech area with extensive implications across multiple industries. At its core, computer vision empowers machines to interpret and understand visual information, enabling advancements in areas as diverse as

manufacturing, urban management, biometric security, healthcare, and autonomous driving. As these applications continue to grow, so does the necessity to deepen the scientific knowledge around them.

When it comes to industrial development, automation plays a crucial role in optimizing production processes and improving labor-force safety and efficiency. Integrated computer vision-based systems are capable of detecting manufacturing defects on production lines, predicting maintenance needs, or ensuring quality control with inspection sensors. This class of technology is also significantly improving safety by monitoring hazardous environments and identifying anomalies that could lead to accidents. Manufacturing leaders have expressed their growing belief that smart factories will transform the way goods are produced over the next five years. Implementing smart solutions is also expected to raise key performance indicators such as quality and labor productivity by 12%, according to Deloitte's 2024 manufacturing industry outlook.

As the benefits of computer vision extend beyond the factory floors, cities worldwide are experimenting with pilot projects to solve complex urban management challenges. Computer vision systems have the potential to optimize tasks such as traffic flow adjustments to reduce congestion at peak hours, improve public transport systems, or enhance public safety.



Biometric security, particularly facial recognition systems, has gained considerable attention and is making significant improvements in terms of research and innovation. Despite all the controversy surrounding its use, including concerns with privacy issues and algorithm bias, facial recognition is rapidly advancing to make it more accurate and reliable. This progress, though, comes with the growing need to address the ethical implications of implementing such technologies, particularly in public spaces.

In the healthcare sector computer vision is revolutionizing medical diagnosis and patient care. Especially with AI-powered imaging tools the accuracy and speed of disease detection have improved in areas like radiology and pathology. Beyond diagnostics, computer vision techniques have been used to monitor patients remotely, analyze surgical procedures, and assist in personalized treatment strategies.

As these examples illustrate, the application power of computer vision is vast and varied, touching nearly every aspect of modern life. With such expansion comes a set of challenges to be meticulously addressed by the scientific community, industrial bodies, and governments. Issues like privacy, misuse potential, and the need for transparent and unbiased engines are critical, now that computer vision-based systems become more and more integrated into our daily lives.

In this edition of the Behind the Code series, we will cover the projects of four Software Campus participants who are dedicated to driving progress in the field. We will address the complexities and current trends, offering insights from those at the forefront of these rapidly evolving technologies. These projects are being developed at **TU Darmstadt** and two **Fraunhofer Institutes** in collaboration

with **ZEISS**, **Software AG**, and **SHS - Stahl-Holding-Saar (SHS)**. The Industry Partners (IP), along with **Volkswagen**, will also contribute with their valuable insights into this prominent field. To conclude, we will present the contribution from a **Software Campus founder** that illustrates how medical computer vision and AI are improving cancer prognostics in the healthcare sector.


Deep Dive:

The computer vision community projects

SMART FACTORIES: BALANCING WORKSPACE SAFETY AND PRIVACY

The integration of Computer Vision in industrial settings brings challenges to the intersection of technological innovations and safety. In the steel industry, occupational safety is of utmost importance, and a vast range of measures are in place to prevent accidents. The latest technologies, including AI, have been incorporated to improve safety standards and reduce the likelihood of accidents. This is achieved by modernizing existing methods and safety precautions.

Mickael Cormier, a PhD candidate at the **Fraunhofer Institute of Optonics, System Technologies and Image Exploitation (IOSB)** is developing his Software Campus project under the supervision of Prof. Dr.-Ing. Jürgen Beyerer and in collaboration with **SHS**. His project aims to employ thermal imaging, computer vision techniques and strong privacy shields to detect potential hazards in real time. The practical application deploys an integrated alarm system to prevent work-related accidents in industrial settings.




Although we instinctively consider such technology of undeniable importance, monitoring systems raise important ethical and privacy concerns. Mickael highlights the delicate balance between safety and worker privacy when using surveillance technologies in a workspace. The project's main challenge, and primary novelty, is the mixture of thermal imaging and strong firewalls with multi-step verification to enforce data protection. Unlike traditional optical cameras, thermal imaging can operate effectively in multiple conditions, typical in dynamic industrial environments. The technology enables individuals' surveillance, but not workers' recognition. One important technical aspect faced by Mickael's team relates to the extreme temperatures found during metal manufacturing, and the use of thermal imaging allows for the differentiation between a human presence and the high-temperature surroundings. However, all technologies have their limitations. The project is still exploring the full capabilities of thermal imaging, particularly in situations where a person is partially obscured by larger objects. The development of new algorithms capable of detecting and interpreting such scenarios is ongoing, and continuous learning may eventually play a role in refining the models.

The collaboration between Mickael and the SHS specialists has been dynamic and close, involving data collection and scenario testing. These are crucial steps for developing a robust and reliable model. Although the data processing and analysis is still labor-intensive for the team, the researcher emphasizes that this is an essential step that will ensure the new system will be capable of performing effectively in real-world conditions in the future.

Mickael's experiences demonstrate the challenges and rewards of the academic-industry exchange throughout the project's development. Working closely with an Industry Partner, his team had to navigate the complexities of project management, data privacy concerns, and shifting project requirements. These challenges are not uncommon in tech-driven projects, where the rapid pace of innovation often outpaces the ability of organizations to adapt. As the project moves forward, the focus remains on improving detection accuracy and overcoming the technical limitations of current thermal imaging technology. The interviewee expressed optimism about the future, anticipating that continuous learning and further refinements in Computer Vision algorithms will enhance the system's capabilities. The project represents a significant step toward integrating AI-driven safety mechanisms in industrial environments. For computer scientists and industry professionals, it's a highlight of the importance of interdisciplinary collaboration and the careful consideration of ethical issues in the deployment of new technologies.

SENSORS AND ACTIVITY RECOGNITION, AN INNOVATIVE WAY TO MAKE OUR CITIES SMARTER

As urban environments become increasingly interconnected, the role of innovation in ensuring efficient public safety and dynamics has rapidly grown. **Thomas Kreutz** provided us with valuable insights into the challenges and solutions shaping the field. He has been a Software Campus participant since 2022, partnered with **Software AG**, and has developed his PhD at **TU Darmstadt**, under Professor Dr. Max Mühlhäuser's supervision.



Thomas's project combines LiDAR and IMU sensors to enhance public safety in urban environments. While IMU measures acceleration, orientation, and rotation changes and the acquired data is used for activity recognition, the LiDAR data is primarily used for object recognition (through motion segmentation) and semantic scene information (semantic segmentation). These sensors have enormous application power and have been extensively used in smart environments and autonomous driving, combined with optical cameras and other sensors.

In smart city research, multiple strategies are necessary to analyze the vast amount of data generated while identifying behavioral patterns and detecting anomalies. Thomas used LiDAR sensors to acquire detailed 3D point cloud data of the surroundings and combined with sophisticated machine learning architectures, to enable systems to understand and interpret the scene's complexity. With the recently published cross-modal learning approach LiOn-XA (Kreutz et al., 2024), LiDAR-based models can accurately perform semantic segmentation, for instance by clearly differentiating between the background (buildings, trees, drivable surfaces, etc.) and foreground (pedestrians, vehicles, cyclists, etc.). Integrating two data streams derived from different LiDAR representations helps compensate for information loss, ensuring high-resolution spatial awareness and the precise semantic segmentation of an urban scene.


To ensure the architecture generalizes well across different data domains (e.g., from simulation to real-world data), he combined cross-modal learning and adversarial training. The technique minimizes the discrepancy between the source domain (training data)

and the target domain (real-world or different dataset) representations. By aligning feature representations in both domains, the model becomes more robust and adaptable to different environments.

In addition to semantic scene structure, the occupancy changes in LiDAR data over time enable systems to detect and track dynamic objects, such as pedestrians or vehicles. LSTM and CNN architectures are being employed to process this sequential data, with models learning patterns over time to identify and segment motion. This unsupervised approach (Kreutz et al., 2023) can automatically discover moving objects, improving urban monitoring systems without requiring manual labeling.

Despite the potential of building smart cities and leveraging machine learning methods to support this type of solution, several technical challenges need to be addressed and kept in mind. Different from what we saw for smart factories, where many circumstances are often repetitive and to a certain point predictable, urban settings are characterized by a wide range of human behaviors. Developing models that can accurately distinguish between normal and suspicious activities in such a dynamic setup is a complex task. While training the models, context becomes an important step in activity recognition, and the model should be capable of understanding, for instance, how running in a park can be different from running in a restricted area and its relationship with security breaches.

Under this scenario we all have a persistent thought: is it safe – and necessary – to be tracked? As already discussed in the smart industry project developed by Mickael Cormier, analyzing activity in public spaces raises significant privacy and ethical concerns, an issue Thomas addressed with



careful consideration. While the primary goal of the models and later systems to be implemented is to enhance safety and speed response time whenever needed, personal data is inevitably collected and analyzed. Although the LiDAR technology collects data differently from regular cameras, for instance, where faces and physical traces are not accessible, patterns of behaviors and routines can be easily assessed. To achieve the real goal, developers and policymakers must work together to ensure these systems are used responsibly and strong privacy shields are enforced to ensure data is protected from misuse.


An important aspect raised by Thomas is that, although we are still realistically far from having true smart cities implemented, the scientific community must build high-quality data sets that allow researchers to tackle the best practices and strategies to train the models and support the development of useful and safe systems. For computer scientists and professionals in the tech industry, this interview serves as a reminder of the importance of interdisciplinary collaboration. As cities continue to grow and become more complex, the work being done today in AI, semantic segmentation and activity recognition will play a crucial role in shaping tomorrow's safe, efficient, and equitable urban environments.

THE GROWING ROLE OF BIOMETRICS AND FACIAL RECOGNITION

In the rapidly evolving landscape of machine learning, integrating biometric security systems presents a fascinating intersection of technology, privacy and ethics. **Marco Huber** is a Software Campus participant developing his PhD at the **Fraunhofer-Institut für Graphische Datenverarbeitung**

(IGD) under the supervision of Prof. Dr. Arjan Kuijper and Dr.-Ing. Naser Damer. In collaboration with **Software AG**, Marco's project involves developing machine learning models for biometric security, with a special focus on facial recognition systems. The fundamental idea is to create a secure relevant model capable of reliably identifying and verifying individuals in critical contexts, such as border control and secure facility access. As these systems process sensitive information, there is a pressing need to make them more transparent and understandable to both users and developers.

Marco emphasizes the concept of explainable AI in biometrics to enhance the transparency and interpretability of facial recognition models. Besides providing an understanding of the model's reasoning and outcome, his team explored how these explanations can be utilized beyond human interpretations. For example, the team investigated the possibility of using explanation maps – tools that visualize the decision-making process of the models – to identify and address issues such as ethnicity or gender bias. This approach reflects an important trend in biometrics research, where the focus is shifting from purely performance-based metrics to a more holistic understanding of how the systems function and how they can be improved. On the technical side, Marco employed advanced machine learning architectures, primarily based on variations of the ResNet model, which are specially adapted for facial recognition tasks. Unlike standard classification models, these tasks require distinguishing between subtle differences in facial features, needing specialized loss functions that can handle all the nuances effectively and deal with the fact that the identity classes used during training are not the same the system is later confronted with



in practice. One of the project's innovations is the use of ElasticFace ([Boutros et al., 2022](#)), a loss function tailored to improve the accuracy and reliability of face recognition. This new strategy was published in 2022 by Fadi Boutros, who is now a Software Campus alumnus. Alongside other state-of-the-art methods like ArcFace and CurricularFace, it forms the backbone of Marco's approach, where the different model behavior is analyzed to derive implications on the final model.

At this point, it is clear that computer vision and ethical issues walk together. For Marco's project, the challenges are inherent as people's faces are used to train the models. Especially in countries like Germany, where privacy and data protections are highly considered, the potential for misuse of these technologies raises important debates on the thin balance between privacy and increased security. Facing this challenge, a hot topic in Computer Vision is the use of synthetic data, where generative models produce a face dataset for the model to be trained. This technique allows computer scientists an alternative source without relying on individual permissions. The use of synthetic data for face recognition also plays an important role in training models to detect deepfakes. Although the performance of models trained on synthetic databases continues to improve, and they still exhibit significant bias, Marco believes it is only a matter of time before they reach the same level of performance as natural data.


MEDICAL COMPUTER VISION AND DIFFUSION MODELS

Medical computer vision has achieved significant technological advancements, particularly in image generation and

segmentation, offering great potential in clinical training and diagnostics. The global medical imaging market, worth approximately \$38.9 billion in 2020, is expected to continue growing, with Europe at the forefront of the technological development of AI-based solutions.

Dr. Moritz Fuchs, who graduated with his PhD last June 2024, explored aspects of applied computer vision in synthetic dataset generation for cataract surgery, an area still largely unexplored. Moritz pursued his PhD at **TU Darmstadt** under the guidance of Dr. Anirban Mukhopadhyay and in close collaboration with the **ZEISS Group** during the Software Campus program. The project's approach aimed to create highly realistic simulations that can be used to train surgeons. The long-term goal is to provide an interactive simulation tool allowing medical professionals to practice surgeries in a controlled, virtual environment. Although achieving real-time performance in the next few years is unlikely due to resource limitations, the potential impact on surgical training and medical computer vision research is enormous.

A significant part of the research utilizes CycleGANs and some adaptations to improve temporal consistency, which is critical for medical imaging where sequential images (like we see in videos) need to maintain a consistent motion. CycleGANs operate on a "game-theory-like" principle, where two neural networks – the generator and discriminator – compete to produce and verify synthetic images. This architecture proved effective in domain transfer, although its application in videos remains a challenge for the future. The current challenge here was to translate motion between different datasets with varying styles and conditions – such as lighting and magnitudes. The use of



optical flows to measure and transfer motion between frames made the synthetic images more realistic and cohesive over time.


The initial approach began with traditional scan data, but the deeper achievement came when diffusion models were implemented in the image generation, connecting real and simulated data. The model brought a great increment to the image's realism, improving data interpretability. A key advancement was integrating uncertainty guidance into the models, enabling a form of autonomous training. By iterating between generating images and training on uncertain or rare cases, the model continuously improved. While the current technology does not yet support real-time video generation, the foundations laid by Moritz's project open new perspectives for future research on surgical model training.

Interestingly, Moritz's project was supported by Yannik Frisch for the past years and now this project will continue with Yannik as a Software Campus participant. This second project, also supported and partnered with ZEISS, makes the next move by integrating SceneCraft into the model. The tool enables more controlled simulations and is expected to further enhance the realism and utility of the datasets. By making a virtual tool and by creating a more controllable environment, the project seeks more immersive professional training experiences for the future.

Despite the promising results, real-time implementation remains a significant challenge. As the researcher noted, while the technology may not require real-time capabilities for creating artificial datasets, it is essential for surgical training, where real-time feedback is crucial. The transition from simulation to real-world application is a complex process regulated by stringent

regulatory frameworks, as it should be, especially when considering the use of AI in sensitive medical procedures. The EU AI Act and MDR (Medical Device Regulation) impose strict guidelines that make it challenging to bring such innovations to market, particularly when the safety and reliability of AI-driven robotics still pose a risk.

The project also benefitted from collaboration with the industry leader in optics and photonics, ZEISS, a company with extensive experience in medical devices. ZEISS supported Moritz by providing access to data and their microscopes for the required analytical procedures. This partnership highlights the importance of industry-academic collaborations in advancing medical technology. The researcher emphasizes that working with private companies brings outstanding achievements yet poses challenges concerning intellectual property (IP) and data accessibility. Obtaining the necessary approvals to publish data can be a lengthy process, sometimes hindering the ability to share research findings promptly. For researchers working with proprietary data, academic freedom, and IP restrictions remains a topic to be carefully balanced. Looking forward, this project represents an important step in the development of AI-driven surgical simulations. The synthetic data generation techniques explored here can lead to enormous progress: interactive simulation tools that could accurately mimic real surgeries. As collaborative work between academia and industry continues to grow with programs like the Software Campus, projects like this will be critical in pushing the boundaries of what is possible in the future of medical computer vision.



Innovation in Computer Vision:

Our Industry Partners' perspectives

ZEISS has been a Software Campus industry partner since 2017 and, ever since, has developed numerous projects in close collaboration with its partnered participants. ZEISS's commercial success is founded on high-precision optics. While 30 years ago, providing the highest possible optical quality was still a strong differentiator in the market, today, every product of the group also depends on the use of advanced algorithms.

"Computer vision plays a major role for the ZEISS group and already has found its way into many products."

Max Riedel

Head of ZEISS Innovation Hub
at KIT


A few examples: ZEISS arivis Pro provides automated cell identification and counting and other advanced imaging analysis tools for microscopy users. ZEISS Automated Defect Detection (ZADD) can detect, localize, and classify defects or anomalies in industrial computer tomography images. The software is specially developed for applications such as castings, injection-molded parts, batteries, and printed components. The ZEISS Secacam trail camera is cloud-connected and automatically identifies and classifies animals in its images, providing decision

support to foresters and hunters.

While AI in the form of neural networks already arrived in those products, researchers at ZEISS are also exploring the latest developments in computer vision. As in the SWC project described above, generative AI is deployed to efficiently scale training data creation through synthesized and augmented data sets or to close the sim-to-real gap in simulations. Currently available foundation models are often not usable out-of-the-box for specific industrial applications, as they are trained on everyday scenes. However, they can serve as a basis for developing specialized solutions much quicker, as they can greatly reduce the amount of labeled data needed.

Computer Vision, in particular visual foundation models, is advancing at an incredible speed, which is both a chance and a challenge for research and development. A solution to keep up with this speed is to build software architectures that leverage externally available AI models and add use-case-specific functionalities. Every advancement in the external models then directly translates into performance improvement of the own software.

Stahl-Holding-Saar has been our partner from 2022 to 2023 and has supported our participants in developing their projects with special attention to the metal industry sector. Looking ahead, SHS continues to expand the application power of deep learning and computer vision, now integral parts of many industries.



“We, in the steel industry, use these technologies, e.g., to check the quality of the surfaces of our products, product tracking or ID checking. Another important application is the classification of materials. Due to the green transformation, the secondary raw material steel scrap is becoming increasingly important. With the help of computer vision, we have the opportunity to develop a fully automated scrap classification and confusion control system.”

Michael Schäfer

Stahl-Holding-Saar, Head of Digitalization & AI

The Computer Vision Community:

Creating impactful innovation out of purpose, a founder story


The Computer Vision Community of Practice also holds an inspiring journey when it comes to entrepreneurial spirit, motivation and focus. **Omar El Nahhas**, a young computer scientist, **founder and CEO of StratifAI**, a PhD candidate at **TU Dresden**, and a Software Campus participant, shared valuable insights into the challenges and

opportunities of founding a tech start-up in medical computer vision.

Omar's journey into entrepreneurship began in his early days as an engineering student, where he started by creating hardware solutions to address local issues. One of Omar's early start-ups, launched during his master's studies in Spain, involved optimizing factory processes with software development. Although this initiative did not develop as expected, as the potential clients were more interested in a service business than a product, the business was profitable for some time. The researcher emphasizes that his first company was a learning opportunity, shaping his understanding of the market, business strategies, and the adaptability and resilience required in entrepreneurship.

After relocating to Estonia during the 2019 pandemic, Omar was inspired by the agile start-up culture and widespread public support for innovation in the country. His exposure to different environments and problem-solving mindset proved essential to his development as a founder. This time, motivated to make an impact on the healthcare system, particularly in the area of cancer treatment, he sought out opportunities in medical computer vision, influenced by his personal experiences. The partnership with Prof. Dr. med. Jakob Kather (TU Dresden) marked a turning point, shifting his focus from purely technological solutions to science-based applications with a tangible societal impact.

The start of his PhD and being part of the Software Campus program further solidified Omar's interdisciplinary approach, blending academic rigor with industry-driven insights. He highlights that navigating both the academic and industry perspectives is crucial in a research-based start-up environment.



The dual exposure helps him to manage expectations, build valuable networks, and foster a collaborative mindset. Omar has **ZEISS** as an Industry Partner, where he develops part of his PhD dissertation in close collaboration with the company's team. StratifAI is a science-based and AI-driven

start-up in the field of cancer prognostics. The team, headed by four co-founders, has raised €1.5M in pre-seed funding. Omar will be pitching on StratifAI at the Software Campus Summit, happening on 18 and 19 of November 2024 in Berlin, during our "Founders Pitch" session.


OMAR'S TOP PRACTICAL ADVICE FOR ASPIRING SCIENCE-BASED ENTREPRENEURS:

- **Clarify your Intellectual Property Rights early on:** it's a very different first tip, but it will save you time and support your company's future foundations.
- **Find the right co-founders, you don't have to do it alone:** similar motivation levels, complementary skills, and shared values are a must. Don't be afraid to be picky!
- **Be aware of a shift in roles and priorities:** for a while, you will be more of a business manager, an HR specialist, an accountant, a pitch expert and less of a researcher and computer scientist.
- **Embrace resilience and long-term vision:** your business will most probably not be ready and running in the blink of an eye. Be patient!
- **Develop a strong company culture from day one:** for science-based companies, fostering a balanced environment where researchers and industry professionals feel comfortable and valued is key. The more aligned you all are, the better your team will be.

Conclusions and Remarks

With the launch of the last four Communities of Practice (Data Science and Analysis, IoT and Distributed Systems, Interdisciplinary Machine Learning, and Agile and Software Development and Engineering), we are now prepared and enthusiastic about our

upcoming events and discussions. During the Software Campus Kickoff event for the new class of 2024, we welcomed 44 new participants to the program with a welcoming and engaging afternoon. The event took place on September 2, 2024, in



Berlin. Participants were tasked with creating a short video with GenAI tools that reflected the strengths – or the superpowers – of their community. This inaugural networking event marked the beginning of their Software Campus program. The participants will be introduced to the other community members in early 2025.

At the Software Campus Summit 2024, happening in November in Berlin, the Communities of Practice will come together again for a new activity. This time, the participants will also be supported by our Alumni Association and representatives of our Industry and Research Partners. In the next issue of the Behind the Code series, we will learn with the **Cybersecurity and Privacy** CoP. This Community is focused on topics related to network intrusion detection, API misuse detection, code-reuse attack mitigation, and much more. The Software Campus projects in Cybersecurity and Privacy are partnered with companies such as DATEV, TRUMPF, Huawei, and IAV. The participants develop their projects at many partner research institutions, such as: TU Darmstadt, TU Dresden, TU Munich, TU Berlin, University of Paderborn, Karlsruher Institut für Technologie and Fraunhofer IUK-Technologie.

Acknowledgements

The **Software Campus team** kindly acknowledges all the professionals who found time to contribute to this article. A special thank you to the **participants** who accepted being interviewed and shared their personal stories and perspectives beyond the programming. Your commitment to self-development within the Software Campus program is inspiring to witness.


To **ZEISS, Volkswagen, SHS - Stahl-Holding-Saar**, and **Software AG**, we appreciate your contribution and openness in supporting the Software Campus Communities of Practice and the production of the Behind the Code series. We also greatly acknowledge the **BMBF** funding (Förderkennzeichen 01IS22096). To the **DLR Projektträger**, thank you for closely supporting the Software Campus program and additionally, the Communities of Practice initiative.

References

Fadi Boutros, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Elastic face: Elastic margin loss for deep face recognition. In IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2022, pp. 1577–1586. 2022,

Thomas Kreutz, Max Mühlhäuser, Alejandro Sanchez Guinea. Unsupervised 4D LiDAR Moving Object Segmentation in Stationary Settings With Multivariate Occupancy Time Series. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pp. 1644–1653. 2023,

Thomas Kreutz, Jens Lemke, Max Mühlhäuser, Alejandro Sanchez Guinea. LiOn-XA: Unsupervised Domain Adaptation via LiDAR-Only Cross-Modal Adversarial Training. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2024.



LEGAL NOTICE

The content of this article is endorsed by all parties involved - the Software Campus participants, their Research Partners, and the project's Industry Partners. However, the views expressed are those of the authors and should not be attributed otherwise. We encourage the dissemination of this publication as it is, but reproduced copies may not be used for commercial purposes. Further use is permitted under the terms of the Creative Commons Licence (CC BY-SA). If you do so, please include a reference to EIT ICT Labs Germany GmbH. If you have any questions regarding the further use of this material, please write to info@softwarecampus.de.

Published by

EIT ICT Labs Germany GmbH, Berlin

Genthiner Straße 8
D-10785 Berlin

Stefan Jazdzejewski
Managing Director

Copyright

EIT ICT Labs Germany GmbH
Berlin, Germany
All rights reserved
2024

