

GDS AUTUMN CONFERENCE 2026

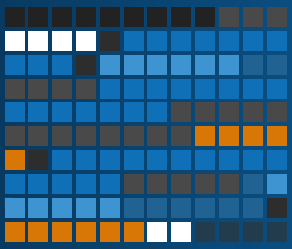
AI Sovereignty in the Age of Agentic AI

Research, Industry and Society

Thu, 19 November 2026
9:00 to 17:40 · doors open 8:45

Google Berlin · AI Innovation Hub
Tucholskystraße 2, 10117 Berlin

Book of Abstracts



The day in data pixels:
one square, five minutes,
one row per hour from 8:00 to 17:55.

BEFORE YOU START

How to Use This Book



Tabs on the right edge. Every page shows all sessions. The blue tab marks where you are. Tap any tab to jump to that session.



The day in data pixels. On page 6, one square stands for five minutes. Tap a square to open the talk running at that time.



One page per talk. Title, speakers, abstract and a short biography, all on one page.

◂ ▸ **Arrows at the bottom.** They lead to the previous and the next talk, in the order of the day.

12 **Page numbers.** The number at the bottom of a talk page takes you back to the programme.

S3.2 **Codes.** S3.2 is the second contribution in session 3. The same code appears in the programme and in the index.

IMPRINT

Published by the German Data Science Society (GDS) e.V., Berlin Regional Group, for the GDS Autumn Conference on 19 November 2026 at Google Berlin, AI Innovation Hub. Editor: Prof. Dr. Daniel Ambach.

Abstracts and biographies appear as the speakers provided them. Where we translated a text, a note says so. Titles, affiliations and times as of 18.11.2026; the programme may still change.

Download: gds-society.de/news/herbsttagung-2026.

Design based on the GDS design system. Set in Raleway and Roboto Mono.

CONTENTS

What Is Inside

Greeting from the Board p. 4

Welcome from the Organisers p. 5

The Day at a Glance p. 6

Programme p. 7

Good to Know p. 10

SESSIONS

S1 Opening 9:00–10:15 p. 11

S2 Data First 10:20–11:00 p. 14

S3 Sovereign and Trustworthy AI 11:20–12:35 p. 18

S4 Fireside Chat 13:40–14:05 p. 22

S5 Agents in Production 14:10–15:25 p. 24

S6 Agents at Work 15:50–16:55 p. 28

S7 Closing 17:00–17:40 p. 33

REFERENCE

Index of Speakers and Chairs p. 36

GREETING FROM THE BOARD

Welcome to the GDS Autumn Conference in Berlin!

A very warm welcome to all participants of the GDS Autumn Conference in Berlin. We are delighted to bring together experts and practitioners to share experiences, exchange ideas and explore the latest developments in Data Science and Artificial Intelligence.

At the GDS, we believe that events like this are about more than just discussing new technological developments. They are valuable opportunities to connect people and expertise, encourage dialogue across different perspectives, foster partnerships and translate innovative ideas into practical solutions. Our aim is to provide a platform for knowledge sharing, new impulses and meaningful collaboration.

The question of AI sovereignty is particularly relevant today. As AI agents increasingly shape our ways of working, decision-making and interacting, we need to ensure that we retain the ability to understand, shape and govern these technologies. Developing our own expertise, maintaining the freedom to make strategic choices and ensuring trust in these technologies will be essential.

We wish all participants an inspiring conference, interesting discussions and many new perspectives!

Göran Kauermann
Chairman GDS

Everything Around the Model

Models are getting better and easier to use. Whether AI works in practice, and whether we stay in control of it, is decided by everything we build around the model.

"The model was never the hard part", says the title of one talk this morning. Read the abstracts of this day side by side and the same idea comes back from every direction. Powerful models are arriving fast, from tabular foundation models to open European LLMs such as Soofi. Yet the speakers spend little time on the model itself. They talk about the work around it: which data to collect and how clean it is, which context an agent sees, which rules it must follow, how each step can be traced, and who decides what "better" means.

This is where trust comes from. Across the programme, trust is not a promise of the model but a property of the system: deterministic guards before an agent acts, certificates before an answer reaches a customer, test cases with cent-exact prices, and people who catch what the pipeline missed. It is also where sovereignty begins. Whoever owns the data, the rules and the judgement keeps the ability to act, whichever model runs underneath.

The programme follows this idea through the day. We start with data, turn to sovereign and trustworthy AI, hear the view of healthcare in a fireside chat, and close the afternoon with agents in production and agents at work, with honest numbers on cost, control and limits.

Many contributions came through our call for papers. They appear as ten-minute talks, followed by a joint discussion with all speakers of the session. Please join in: the discussions are where questions from the room meet experience from the field.

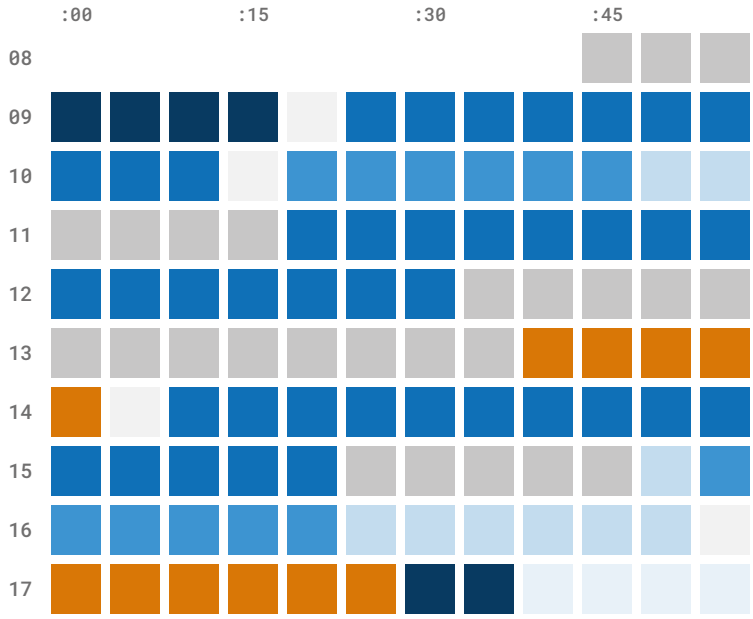
We thank all speakers for their time and their openness, our host Google Berlin for the venue, and everyone in the GDS who made this day possible.

Prof. Dr. Daniel Ambach
for the organising team of the GDS Berlin Regional Group

THE DAY AT A GLANCE

The Day in Data Pixels

Each square stands for five minutes, each row for one hour. Tap a square to open the talk that runs at that time.



- Talk, 20 + 5 min
- Short talk from the call for papers
- Joint discussion, introduction
- Welcome and closing
- Registration, coffee, lunch
- Changeover
- Networking
- Highlights: fireside chat, closing talk and book launch

PROGRAMME

Thursday, 19 November 2026

Tap a title to open its abstract. The codes S1 to S7 mark the sessions.

8:45 ***Doors open, registration and coffee***

S1 Opening 9:00–10:15

Chair: Prof. Dr. Anne Schwerk

9:00 *Welcome by Prof. Dr. Göran Kauermann (GDS chairman) and our host Google*

9:25 How to deliver a successful Data Science project in business — S1.1 — p. 12
Zlata Tanovic

9:50 Soofi: A European LLM. The Next Step: Modular, Self-Improving AI — S1.2 — p. 13
Prof. Dr. Alexander Löser

S2 Data First 10:20–11:00

Chair: Martin Habedank

10:20 Adaptive Sampling with a Gravity Prior: Can Tabular Foundation Models Better Predict Which Data Should Be Collected? — S2.1 — p. 15
Christian Hoops

10:30 Automated Data Quality Control in the Age of Generative AI – An Empirical Evaluation of Leading Large Language Models — S2.2 — p. 16
Niklas Schüschke; Johannes Rieper

10:40 The Model Was Never the Hard Part: Agents for Predictive Data Science — S2.3 — p. 17
Mattia Dimauro

10:50 *Joint discussion with all speakers of the session*

11:00 ***Coffee break***

S3 Sovereign and Trustworthy AI

11:20–12:35

Chair: Prof. Dr. Daniel Ambach

11:20 Trustworthy, Reliable & Sovereign AI – When can we Count on AI? S3.1
Prof. Dr. Gerard de Melo p. 19

11:45 Trust, but Trace: Observability for Agents that act S3.2
Sasa Fajkovic p. 20

12:10 Shaping Sovereign AI: Europe's Path from Dependence to Strength S3.3
Jörg Bienert p. 21

12:35 *Lunch break*

S4 Fireside Chat

13:40–14:05

Host: Prof. Dr. Anne Schwerk

13:40 Who Controls AI When AI Takes Control? A Fireside Chat on Healthcare Data, Agentic AI & Digital Sovereignty S4.1
Bettina Goerner p. 23

S5 Agents in Production

14:10–15:25

Chair: Prof. Dr. Anne Schwerk

14:10 When Performance alone is not Enough - a Multidisciplinary View on Clinical Decision Support S5.1
Dr. Roland Roller p. 25

14:35 Building AI Systems with Event-driven Operational Agents and Real-time Context Serving S5.2
Steffen Hoellinger p. 26

15:00 MLOps for Multi-Agent architectures S5.3
Kai Sahling p. 27

15:25 *Coffee break*

S6 Agents at Work

15:50–16:55

Chair: Martin Habedank

15:50 *Introduction to the session*

15:55 Can You Trust What Your Agent Just Told a Customer? A Certificate-Based Approach — S6.1 p. 29
Dr. Lukas Schack; Ulrich Arndt; Michele Zilli

16:05 When agents run the experiments: AI-first model development at Sixt, and what stayed human — S6.2 p. 30
Dr. Vasyl Hafych

16:15 Agents, Cost, and the Limits of Automation: A Production Case Study — S6.3 p. 31
Zino Lang

16:25 A Treasure's Impact: AI Agents Turning Hospital Data into Better Staffing and Care (cancelled) — S6.4 p. 32
Dr. Martin Weiss

16:25 *Joint discussion with all speakers of the session*

S7 Closing

17:00–17:40

Chair: Prof. Dr. Daniel Ambach

17:00 Avoiding the AIceberg: How blindly trusting tech hurts us all — S7.1 p. 34
Chris Heilmann

17:25 Launch of the 2nd edition of "Künstliche Intelligenz und Data Science in Theorie und Praxis" (Springer Spektrum) — S7.2 p. 35
Prof. Dr. Göran Kauermann

17:30 *Closing remarks and outlook*

17:40 ***Networking with drinks***

GOOD TO KNOW

Venue, Format and Team

VENUE	Google Berlin, AI Innovation Hub Tucholskystraße 2, 10117 Berlin
TIME	Doors open at 8:45. Programme from 9:00 to 17:40, networking with drinks afterwards.
LANGUAGE	Most talks are in English.
FORMATS	Talks take 20 minutes plus 5 minutes for questions. Short talks from the call for papers take 10 minutes; all speakers of a session then join a common discussion.
CATERING	Coffee breaks and lunch are included.
SLIDES	With the speakers' consent, we will share the slides with all participants after the conference.
WI-FI	Access on site: please ask the organising team at the registration desk.
QUESTIONS	Ask the speakers via Slido, also anonymously: open **our Slido room** or go to slido.com and enter the code **#4038 262**.

ORGANISING TEAM

PROGRAMME

The GDS programme team

ORGANISATION

GDS Berlin Regional Group:
Prof. Dr. Daniel Ambach (lead) and **Martin Habedank**

with the GDS board, the GDS office and the GDS marketing team

Meet the GDS team: `https://gds-society.de/ueber-uns`

HOST

Google Berlin, AI Innovation Hub



SESSION 1 · 9:00–10:15

Opening

Welcome by Göran Kauermann, chairman of the GDS, and our host Google, followed by two opening talks from practice and research

Chair: Prof. Dr. Anne Schwerk, IU International University of Applied Sciences, GDS board

---

9:00 *Welcome by Prof. Dr. Göran Kauermann (GDS chairman) and our host Google*

9:25 **How to deliver a successful Data Science project in business**
Zlata Tanovic

9:50 **Soofi: A European LLM. The Next Step: Modular, Self-Improving AI**
Prof. Dr. Alexander Löser

S1.1 · 9:25–9:50 OPENING TALK · 20 + 5 MIN

How to deliver a successful Data Science project in business

Zlata Tanovic
Associate Director of Data Science, HelloFresh SE

ABSTRACT

Drawing on real experience leading Data Science work in a business setting, this talk walks through the practical considerations that separate a project that works in a test setting from one that is successful in production. It covers how thinking about customer experience—internal or external—needs to happen at design time and why aligning with the people who produce your data is as much a part of the job as modeling it. It also addresses a less technical but equally critical ingredient: earning trust from senior leadership and stakeholders, and why that trust gives a Data Science team the independence to do its best work. Finally, it touches upon whether GenAI can replace the Data Scientist in business and what junior Data Scientists can do today to get their foot in the door in a field that's shifting under everyone's feet. Attendees will leave with a high level strategy and experience based learnings for delivering a real world Data Science project.

About the speaker

Zlata is Associate Director of Data Science at HelloFresh SE in Berlin. She has worked in the Data Science field for almost ten years, leading teams since 2021. Originally from Bosnia and Herzegovina, she grew up in the Netherlands. She holds a PhD in Economics (VU Amsterdam), a master's in Econometrics (University of Amsterdam), and a bachelor's in Math & Physics (Leiden University).

S1.2 · 9:50–10:15 TALK · 20 + 5 MIN

Soofi: A European LLM. The Next Step: Modular, Self-Improving AI

Prof. Dr. Alexander Löser
Berliner Hochschule für Technik (BHT), DATEXIS

ABSTRACT

Europe is already at the grown-ups' table in AI. But it doesn't know It.
Europe needs a different path to powerful AI. The US and China can pool data, capital, and computing power. They focus on large, centrally trained models. Europe, however, must take a technically more complex, modular approach.

The reason lies in Europe's industrial structure. Valuable data and expertise are distributed across many companies. These data remain behind corporate firewalls. At the same time, companies lack the resources to train their own large foundation models. Electricity, GPUs, talent, investment courage, and more.

The first part of the solution is an open, good-enough foundation model, like Soofi. Specialized expert models build on this foundation. They recursively self- improve through post-training. Intelligent routers combine their capabilities. Learning harness systems optimize their performance at runtime. All behind the firewall.

The keynote explores how this modular architecture can become a competitive advantage for Europe. The largest model will not determine success. What matters is the ability to recombine proprietary data in new ways and continuously self-improve.

About the speaker

Alexander Löser is Professor of Data Science at the Berlin University of Applied Sciences and Technology since 2013. He founded DATEXIS and the Data Science+X center, co-founded the MSc in Data Science (2017), and joined the DFG "Berlin Initiative for Research on Foundation Models" (2023). He is also Expert at the BMFTR Platform Learning Systems since 2018 for LLMs and also played a pivotal role in securing doctoral degree–awarding rights for Berlin's Universities of Applied Sciences (HAWs).
His work focuses on deep learning, language models, and explainability. He is also PI in Soofi for Agentic Reasoning , a 30B MOE European LLM, made in Germany. Since 2012, he has advised DAX and NASDAQ firms and co-designed 60+ data products.



SESSION 2 · 10:20–11:00

Data First

Quality, collection and tabular foundation models

Chair: Martin Habedank, GDS board

10:20 **Adaptive Sampling with a Gravity Prior: Can Tabular Foundation Models Better Predict Which Data Should Be Collected?**
Christian Hoops

10:30 **Automated Data Quality Control in the Age of Generative AI – An Empirical Evaluation of Leading Large Language Models**
Niklas Schüschke; Johannes Rieper

10:40 **The Model Was Never the Hard Part: Agents for Predictive Data Science**
Mattia Dimauro

10:50 *Joint discussion with all speakers of the session*

S2.1 · 10:20–10:30 SHORT TALK · 10 MIN · CFP

Adaptive Sampling with a Gravity Prior: Can Tabular Foundation Models Better Predict Which Data Should Be Collected?

Christian Hoops
PredictiveStatistics.eu / TU Dortmund

ABSTRACT

Machine learning typically requires relevant data to be available. In many real-world applications, however, this is not the case: The population of potential observations is very large, while data collection is possible only for a small portion. This applies, for example, to companies in Industry 4.0, where data from machines, equipment, or production processes may first need additional sensors. Similar challenges exist in medical screening, environmental monitoring, or disaster management, where selecting observed units can determine which patterns and structures are detected. The decision of which data to collect largely determines what information is available to the model.

To make this selection adaptive, a sampling framework is developed and evaluated in which a spatial gravity prior guides initial data collection and is continuously adjusted based on newly acquired observations. The prior uses population size, spatial proximity, and regional characteristics to estimate the relevance of unobserved units. After each survey wave, the model updates its estimates and adaptively determines the next sample.

The focus is on a systematic comparison of Gradient Boosting and Tabular Foundation Models under identical conditions. The study investigates whether TFMs, because they can learn effectively from small samples, benefit disproportionately from adaptive collection of targeted observations and can thereby correct misspecified prior knowledge more efficiently.

About the speaker

Christian Hoops is a statistician with senior experience at Ipsos, OTTO, and Haspa, where he led data-driven projects. He played a key role in developing a Data Science & AI platform for the German Savings Banks. His work combines statistical methods, machine learning, and AI. He has spoken at Google Analytics Summit & Predictive Analytics World and received the EHI Science Award from GS1 Germany.

S2.2 · 10:30–10:40 SHORT TALK · 10 MIN · CFP

Automated Data Quality Control in the Age of Generative AI – An Empirical Evaluation of Leading Large Language Models

Niklas Schüschke
AI Project Engineer, Mittelstand-Digital Zentrum Rostock

Johannes Rieper
Protiviti

ABSTRACT

Manual data cleaning is slow and error-prone, so many teams now upload spreadsheets to chat assistants and ask them to find the mistakes. We tested whether that actually works. For two years we have evaluated LLMs on error detection in tabular data. Our latest round covers 11 current models from Google, Anthropic, OpenAI and Mistral on four public datasets ranging from 100,000 to 1 million cells. Each dataset contained 50 manually injected syntactic, semantic and structural errors, and every model received the same short prompt.

The results are mixed. The two most consistent models stayed above 83% F1 on all three smaller datasets, but rankings shifted from dataset to dataset: one model scored 100% on one and 51% on another, and the newest model did not lead everywhere. Weak runs were dominated by false alarms, up to 853 for 50 real errors, and one model found no errors at all. One model family struggled even to read the supplied Excel files. On the one-million-row dataset, no model finished within 15 minutes, as loading failures, processing time and token limits made the workflow impractical.

Attendees will learn why a single benchmark win is a weak basis for choosing a model, why false-positive rates matter as much as recall, where scale breaks the upload-and-ask approach, and how a hybrid human-AI review process can use LLMs today without trusting them unattended. We are also open about the limits of our controlled setup.

About the speakers

Niklas Schüschke works at Mittelstand-Digital Zentrum Rostock as AI Project Engineer, supporting small and medium-sized companies with digitalisation and AI. Together with Johannes Rieper (Protiviti), he has spent two years testing how well LLMs handle practical data tasks and on bridging the gap between theoretical AI capabilities and operational enterprise requirements.
Johannes Rieper: I am an SAP consultant at Protiviti and hold a bachelor's degree in business informatics. My interests include process automation and artificial intelligence, which I also explore through an independent research project on data quality control.

S2.3 · 10:40–10:50 SHORT TALK · 10 MIN · CFP

The Model Was Never the Hard Part: Agents for Predictive Data Science

Mattia Dimauro
Founder, Predictabl

ABSTRACT

Predictive modelling is becoming more accessible. Tabular foundation models, time-series foundation models and a new wave of decision models are making powerful predictive capabilities available with less model-specific work.

Yet turning a business question into a useful predictive system still requires substantial work. It means framing the problem, understanding how enterprise tables fit together, assembling the right data and defining the target. It also means preventing leakage, designing realistic validation, choosing meaningful baselines and metrics, and connecting predictions to business decisions and operational workflows.

Much of applied data science lives in these steps. Agentic workflows offer a way to automate more of them, with agents coordinating data preparation, classical machine learning and specialised models.

Drawing on my experience building Predictabl, I'll share practical lessons from automating these workflows: where agents help, where they fail, how to recover and where human judgment remains essential. I'll argue that automating the work around the model will reshape the data scientist's role more profoundly than improvements in predictive models alone.

About the speaker

Mattia Dimauro is the founder of Predictabl, building an agentic data scientist for predictive modeling on tabular data. Previously Staff AI Product Manager at Aleph Alpha, he led product for Delivery Hero's Logistics Data Science and SAP AI. Previously a data scientist shipping production ML at SAP, Blinkist and Pro7Sat1. He builds at Tech:Europe's Sandbox and writes on honest model evaluation.



SESSION 3 · 11:20–12:35

Sovereign and Trustworthy AI

Reliability, observability and European sovereignty

Chair: Prof. Dr. Daniel Ambach, GDS Berlin Regional Group

11:20 **Trustworthy, Reliable & Sovereign AI – When can we Count on AI?**
Prof. Dr. Gerard de Melo

11:45 **Trust, but Trace: Observability for Agents that act**
Sasa Fajkovic

12:10 **Shaping Sovereign AI: Europe's Path from Dependence to Strength**
Jörg Bienert

S3.1 · 11:20–11:45 TALK · 20 + 5 MIN

Trustworthy, Reliable & Sovereign AI – When can we Count on AI?

Prof. Dr. Gerard de Melo
Hasso Plattner Institute / University of Potsdam

ABSTRACT

AI has become ubiquitous in everyday life, but how reliable are the models really, and how reliable is our access to them? In this talk, I will first outline why sovereign AI is of crucial importance despite the substantial cost of training foundation models. I will then discuss the potential of smaller LLMs to power agentic AI, enabled by more efficient model architectures (e.g., in our SLlama model) as well as by new paradigms such as Knowledgeless Language Models, which are trained to exclude factual knowledge from their parameters. Finally, I will discuss techniques to make LLMs more reliable by improving their meta-cognitive ability to distinguish what they know from what they do not know (e.g., by adding an I-Don't-Know token to the model's vocabulary or adopting our CONCORD training paradigm to avoid overconfident responses). Together, these ideas point towards a future of highly capable agentic systems with on-premises deployment and enhanced trustworthiness.

About the speaker

Gerard de Melo is Professor of Artificial Intelligence and Intelligent Systems at the Hasso Plattner Institute and the University of Potsdam. His group works on large language models and multimodal models. He received his doctorate from the Max Planck Institute for Informatics and held posts at Rutgers, Tsinghua and UC Berkeley. He has published over 200 papers and chairs the AI@HPI Conference.

S3.2 · 11:45–12:10 TALK · 20 + 5 MIN

Trust, but Trace: Observability for Agents that act

Sasa Fajkovic
Lead Applied AI, Trade Republic

ABSTRACT

I built an agent orchestration platform running a fleet of AI colleagues handling over 20,000 tasks a week. They review pull requests, answer help-desk channels, create work and triage alerts. In a regulated environment, every action one of them takes has to be reconstructable afterwards. This talk is a field report on how we made that true.

Trust: the agent holds no permissions of its own. Every command it tries to run passes a chain of deterministic guards. No model is consulted, the same input always gets the same verdict, and the first denial wins. The same allowlist is checked against every prompt and playbook at build time, so a document that teaches a forbidden command fails the build instead of producing a refusal somebody has to debug live, while a colleague waits.

Trace: a run id is attached the moment a message arrives and travels with every log line, guard decision, approval and cost record. Sixteen dashboards grew out of that, none of them planned. Each one exists because somebody asked a question nobody in the room could answer.

I will show three cases where our telemetry lied: a p95 pinned at its top histogram bucket, looking healthy, cost counters that reset on every process restart, and a run that exited zero while posting a half-finished answer.

No vendor lock-in and no agent-observability product anywhere in it. Prometheus and Loki, where label choices mattered more than the model choice. I close with what the system still cannot tell us.

About the speaker

I lead Applied AI at Trade Republic, where I am also a Staff Engineer. Based in Berlin, I set technical direction, grow engineering teams and stay hands-on with architecture and production AI systems. I share practical lessons about leadership, hybrid human and AI teams, infrastructure ownership and responsible adoption.

S3.3 · 12:10–12:35 TALK · 20 + 5 MIN

Shaping Sovereign AI: Europe's Path from Dependence to Strength

Jörg Bienert
Managing Director, Center for Sovereign AI (LEAM GmbH)

ABSTRACT

Artificial intelligence is transforming business, public administration and society. Whoever controls models, data and computing power also shapes the ability of companies and states to act. Digital sovereignty is therefore becoming a strategic question.

The talk examines Europe's challenges: dependence on international providers, access to computing power and data, and building up its own expertise. At the same time, it shows the opportunities offered by open models, strong partnerships and a sovereign AI infrastructure.

Using Soofi and the Center for Sovereign AI as examples, it shows how joint initiatives can strengthen a European AI ecosystem. The focus is on how companies and public institutions can use AI while securing their technological capacity to act.

Translated from the German original by the organisers.

About the speaker

Jörg Bienert is Managing Director of the Center for Sovereign AI (LEAM GmbH) and co-founder of the German AI Association (KI Bundesverband e.V.). He currently also leads the Soofi project, in which a consortium of leading research institutions is training a family of sovereign foundation models. After studying computer engineering and holding several positions in the IT industry, he founded several companies. ParStream, a Silicon Valley big data startup, was acquired by Cisco in 2015.

Translated from the German original by the organisers.



SESSION 4 · 13:40–14:05

Fireside Chat

Healthcare data, agentic AI and digital sovereignty

Host: Prof. Dr. Anne Schwerk, IU International University of Applied Sciences, GDS board

---

13:40 **Who Controls AI When AI Takes Control? A Fireside Chat on Healthcare Data, Agentic AI & Digital Sovereignty**
Bettina Goerner

S4.1 · 13:40–14:05 FIRESIDE CHAT · 25 MIN

Who Controls AI When AI Takes Control? A Fireside Chat on Healthcare Data, Agentic AI & Digital Sovereignty

Bettina Goerner
Chief Data Officer, Eterno Health

Prof. Dr. Anne Schwerk
Host, IU International University of Applied Sciences, GDS board

ABOUT THE CONVERSATION

Description of the conversation to follow.

About the guest

Bettina Goerner specializes in healthcare data, AI and compliance. She is a Senior Advisor to investors and their European healthcare companies at every stage, from due diligence to exit.
She was founding Chief Data Officer at ETERNO, leading innovation and compliance for AI-powered clinical software in the outpatient sector. Previously, she was CDO at Centogene, where she built the data function for multi-omics patient data. Before that, she ran a 450-person global data business at Springer Nature with full P&L ownership. Bettina graduated in Molecular Biology at the Max Planck Research School and trained at McKinsey and INSEAD. She has been recognized in multiple "40 under 40" lists.

About the host

Anne Schwerk is Professor of Artificial Intelligence at IU International University of Applied Sciences and Senior Manager for Responsible AI at Accenture, advising on AI governance and EU AI Act compliance. With a PhD in Medical Neurosciences from Charité Berlin, she has worked at the interface of medicine and AI, from rare disease diagnostics at CENTOGENE to her managerial work at Charité. She sits on the board of the German Data Science Society and co-authored a book on Human-Centered AI in Medicine. Her research centers on explainable AI for medical decision-making, so clinicians and patients can trust AI-supported decisions.



SESSION 5 · 14:10–15:25

Agents in Production

Governance, safety and running agent systems

Chair: Prof. Dr. Anne Schwerk, IU International University of Applied Sciences, GDS board

---

14:10 When Performance alone is not Enough - a Multidisciplinary View on Clinical Decision Support
Dr. Roland Roller

14:35 Building AI Systems with Event-driven Operational Agents and Real-time Context Serving
Steffen Hoellinger

15:00 MLOps for Multi-Agent architectures
Kai Sahling

S5.1 · 14:10–14:35 TALK · 20 + 5 MIN

When Performance alone is not Enough - a Multidisciplinary View on Clinical Decision Support

Dr. Roland Roller
DFKI (German Research Center for Artificial Intelligence)

ABSTRACT

Abstract to follow.

About the speaker

Roland Roller is a Senior Researcher at the Speech and Language Technology Lab of the German Research Center for Artificial Intelligence (DFKI) in Berlin. He studied Computer Science at the University of Trier and Saarland University in Germany and received his PhD from the University of Sheffield, UK. He has also conducted research stays at the University of the Basque Country in Spain, Université Paris-Saclay in France, and the Nara Institute of Science and Technology in Japan. His research focuses on natural language processing and machine learning, with a particular emphasis on applications in medicine and healthcare. His interests include trustworthy clinical decision support, multimodal machine learning, LLM-based agents, and text anonymization. Roland is involved in a range of interdisciplinary research projects spanning several medical domains, including nephrology, emergency medicine, tumor boards, and sexual medicine.

S5.2 · 14:35–15:00 TALK · 20 + 5 MIN

Building AI Systems with Event-driven Operational Agents and Real-time Context Serving

Steffen Hoellinger
Field CTO, Confluent, an IBM Company

ABSTRACT

AI agents are only as good as the context they act on. In many organisations, that context sits in batch pipelines and is already outdated when an agent uses it. This talk shows how event-driven architectures close that gap.

First, streaming agents: Apache Flink jobs that continuously monitor real-time event streams, detect anomalies and patterns, and trigger actions while a customer interaction is still under way. Second, reusable data products and a real-time context engine that serve continuously aggregated, consistent context to AI agents, independent of the agent framework and shared with analytical systems. Third, the step towards an agentic data mesh: how data governance and lineage keep agent development manageable when many teams build agents at scale.

Attendees take away architectural patterns for running operational agents on real-time data.

About the speaker

Steffen is a Field CTO with Confluent, an IBM Company, helping global enterprises harness the full potential of real-time data and AI by bringing Apache Flink and Apache Kafka to the core of their architectures. He partners with customers to modernize their data infrastructure, integrating AI models, metadata, data governance and data lineage to unlock new capabilities for agentic AI powered by continuous intelligence on streaming data.

S5.3 · 15:00–15:25 TALK · 20 + 5 MIN

MLOps for Multi-Agent architectures

Kai Sahling
Founder and CEO, &act

ABSTRACT

Classic MLOps is about training, versioning and deploying models. With LLM-based multi-agent systems there's no model to train, but the operational questions stay. How do you know a change made things better? Who controls what the system decides? What does a single run cost, and what happens when it fails? Answering them shifts the work to prompts, configuration, rules and evaluation data.

This talk uses a production system as the example: a team of agents that turns real customer emails, PDFs and technical drawings into priced, binding quotes for industrial manufacturers. It covers three lessons:

Control and trust. Agents read requests and make decisions. Prices, discounts and quantities come from rule sets that the customer's own staff maintain. Rule changes go through an impact check before release, and every position on a quote shows where it came from and why.

Evals that grow with the product. Each customer has a set of real test requests with the expected business result: products, quantities, cent-exact prices, which rules fired. We check facts, not generated text. Every bug becomes a new case. We'll show how the sets grew to over 100 cases per customer and how we're moving them into CI.

Cost and reliability. We measure what each quote costs, match the model to the task, and design for failures that are visible and recoverable: rate limits, broken tool calls, cut-off output, and agents that loop.

You'll leave with practical patterns for operating multi-agent systems, and for deciding where an LLM belongs and where plain rules do the job better.

About the speaker

Kai Sahling is the founder of &act and responsible for AI that runs reliably in day-to-day business operations. At &act, Kai builds multi-agent systems that turn customer emails, PDFs and technical drawings into priced quotes for industrial manufacturers. Most recently, Kai was Team Lead and Principal MLOps Engineer at Alexander Thamm, bringing AI systems into production for large enterprise clients. Before that, Kai worked as a Machine Learning Engineer at Deutsche Bank and SAP Fioneer, and as a Data Science Consultant at Senacor. Across these roles, the focus was on building ML and data platforms.



SESSION 6 · 15:50–16:55

Agents at Work

Lessons from practice, with value and limits

Chair: Martin Habedank, GDS board

15:50 *Introduction to the session*

15:55 **Can You Trust What Your Agent Just Told a Customer? A Certificate-Based Approach**
Dr. Lukas Schack; Ulrich Arndt; Michele Zilli

16:05 **When agents run the experiments: AI-first model development at Sixt, and what stayed human**
Dr. Vasyl Hafych

16:15 **Agents, Cost, and the Limits of Automation: A Production Case Study**
Zino Lang

16:25 **A Treasure's Impact: AI Agents Turning Hospital Data into Better Staffing and Care (cancelled)**
Dr. Martin Weiss

16:25 *Joint discussion with all speakers of the session*

S6.1 · 15:55–16:05 SHORT TALK · 10 MIN · CFP

Can You Trust What Your Agent Just Told a Customer? A Certificate-Based Approach

Dr. Lukas Schack
Principal Machine Learning Engineer, TUI

Ulrich Arndt
Enterprise Architect, TUI

Michele Zilli
Principal of Immersive Technology, TUI

ABSTRACT

How do you trust what an AI agent tells a customer across millions of interactions a day?

TUI's customer-facing travel assistant operates where a wrong answer is not a typo but a mispriced fare or a rebooking rule that does not hold. Two risks converge: the model can be confidently wrong, and the system can be attacked through prompt injection, tool misuse, or answers that read well but are false. At thousands of conversations a day and scaling toward millions, trust cannot be established by sampling, and prompt-level guardrails enforce nothing. Trust must be structural.

We present the Trust Kernel, a reference monitor between agent and customer. Every answer is checked against policy and released only with a signed, verifiable certificate: it is grounded in permitted, current facts, stays within the customer's rights, and meets its specification. Nothing reaches the customer without a valid certificate, every decision is written to an append-only log, and uncertified output, including from shadow AI, is never delivered.

The kernel is neuro-symbolic. Language models operate only at the seams, turning requests into typed facts and decisions back into language, while a deterministic solver decides over a rule base written as code, cited to source policy, and owned by people.

We report from the concept and its early implementation, using claim and refund processing as a worked example, from customer request to certified decision.

About the speakers

Dr. Lukas Schack (Principal Machine Learning Engineer), Ulrich Arndt (Enterprise Architect) and Michele Zilli (Principal of Immersive Technology) are a TUI team working together on innovative engineering solutions that bring AI agents to life to serve real-world, enterprise-level problems.

S6.2 · 16:05–16:15 SHORT TALK · 10 MIN · CFP

When agents run the experiments: AI-first model development at Sixt, and what stayed human

Dr. Vasyl Hafych
Senior Data Scientist, Sixt SE

ABSTRACT

What does a data scientist still do when AI agents write the code, run the experiments and draft the report?

At Sixt, every rental price rests on an opportunity-cost curve: how much one more car is worth at a given location and day. Over the past year we rebuilt the predictive model behind these curves and we did it AI-first. Coding agents wrote most of the implementation, executed the tuning campaigns we designed, and drafted the reports. Several thousand cross-validated training runs, for a model serving hundreds of locations, were completed in days rather than months, by a team of three.

The results are real: the new model is more accurate than its predecessor, more robust across locations, and trains orders of magnitude faster. But the more interesting outcome is what did not get automated. The key modelling idea came from the team, a large share of the "improvements" the agents found evaporated under proper statistical control, and the most dangerous errors were subtle enough to pass automated check.

This talk is an honest field report: how we organised AI-driven experimentation at scale, what it cost and what it delivered, where it misled us, and how the data scientist's job shifted from building models to deciding what "better" means and standing behind the result.

About the speaker

Vasyl Hafych is a Senior Data Scientist in the revenue management department at Sixt in Munich. He holds a PhD in physics from TU Munich and worked on Bayesian inference and machine learning at the Max Planck Institute for Physics and at CERN before moving into industry.

S6.3 · 16:15–16:25 SHORT TALK · 10 MIN · CFP

Agents, Cost, and the Limits of Automation: A Production Case Study

Zino Lang
Founder, Avolang

ABSTRACT

Does unattended agentic AI beat a human doing the same task while staying in the loop? I built an automated newsletter for LMK, a German subcontractor, tracking forced auctions, renewable-energy planning, and insolvencies across three states, using Haiku for scraping, Sonnet for orchestration, and a three-way Sonnet audit gate before delivery, and compared it against manual, self-hosted, and hybrid setups.

Automated runs cost about $59.60 across a month; the manual, subscription-based path cost roughly $0. But money isn't the real story: verification is mandatory either way, and cost is just what happens when that verification runs through metered API calls instead of a person's attention.

The pipeline's own deduplication missed about 20% of real duplicates, and a missing pagination parameter once produced a silently incomplete report the audit gate didn't catch until pointed at it; human review caught it first. A fully deterministic pipeline with no AI failed outright at recognizing when two sources described the same case. A graph-based retrieval hybrid did help here, cutting duplication by roughly 20% in a short test, real evidence that retrieval-side fixes work. But the serious, generalizable fix isn't a bigger model or better retrieval alone: it's human-in-the-loop oversight and real domain grounding, above all. Neither substitutes for the other.

About the speaker

Zino Lang is the founder of Avolang, helping mid-sized businesses turn AI and reliable data into better decisions. A Claude Certified Architect - Professional and certified by AWS and Microsoft, he combines technical expertise with a pragmatic approach: secure, practical AI solutions, if practical, with measurable value, making advanced technology accessible and useful in everyday business.

S6.4 · 16:25 – 16:35 SHORT TALK · CANCELLED

A Treasure's Impact: AI Agents Turning Hospital Data into Better Staffing and Care

Dr. Martin Weiss
Thought Leader and Strategic Trusted Advisor, Dataciders SE

This talk has been cancelled at short notice. We keep the abstract here for reference.

ABSTRACT

Hospitals possess a treasure whose value is still largely untapped: the data generated every day across workforce management, finance, operations and clinical care. In this thought-leadership session, Dr. Martin Weiss, Thought Leader and Strategic Trusted Advisor at Dataciders SE, demonstrates what AI agents can already deliver when this fragmented data is transformed into governed, decision-ready intelligence.

Using hospital workforce management as a concrete example, the talk shows how specialised agents can identify data inconsistencies, forecast demand, absences and capacity gaps, compare staffing scenarios and explain recommended actions. The goal is not automated cost cutting or the replacement of professional judgement. It is better leadership: less avoidable overtime, lower agency dependency, more stable rosters, improved capacity utilisation and safer, more reliable care.

The session also addresses where agentic ambition must stop. Data quality, bias, permissions, co-determination, privacy, traceability and patient-safety requirements define clear boundaries. Human leaders retain decision rights and accountability.

Participants will gain a vendor-neutral framework that connects use cases, data foundations, agent roles, governance and measurable KPIs. It offers a realistic path from isolated data and promising pilots to trusted agents at work—turning a hospital's data treasure into operational impact for people, performance and care. This is AI with real purpose.

About the speaker

Dr. Martin Weiss is a Thought Leader and Strategic Trusted Advisor at Dataciders GmbH. He advises C-level teams on data-driven transformation, AI governance and operating models, with a focus on Healthcare, Public Sector and Social Economy. His work connects strategic ambition with measurable value, responsible implementation and sustainable organisational change. He is a member of GDS e.V.



SESSION 7 · 17:00–17:40

Closing

Closing talk, book launch and outlook

Chair: Prof. Dr. Daniel Ambach, GDS Berlin Regional Group

17:00 **Avoiding the AIceberg: How blindly trusting tech hurts us all**
Chris Heilmann

17:25 **Launch of the 2nd edition of "Künstliche Intelligenz und Data Science in Theorie und Praxis" (Springer Spektrum)**
Prof. Dr. Göran Kauermann

17:30 *Closing remarks and outlook*

S7.1 · 17:00–17:25 CLOSING TALK · 20 + 5 MIN

Avoiding the AIceberg: How blindly trusting tech hurts us all

Chris Heilmann
VP Developer Relations, WeAreDevelopers

ABSTRACT

As a veteran of Software development since 1996 and a survivor of the .com crash, it is safe to say that the inevitable crash feels different these days. We don't have to worry about a few e-Commerce sites being unavailable. Instead, we must deal with AI being deeply ingrained in people's day-to-day lives and increasingly in the delivery of software meant to make these lives easier. Are we bound to sink or swim with this wave of technology? Chris Heilmann takes a look at what's happening and what you can do to stay afloat.

About the speaker

Chris Heilmann is VP of Developer Relations at WeAreDevelopers, hosts the weekly WeAreDevelopers Live show and writes the Dev Digest newsletter. In software development since 1996, Chris led developer relations at Yahoo, Mozilla and Microsoft. Author of several JavaScript books and The Developer Advocacy Handbook, Chris works for an accessible, resilient web.

S7.2 · 17:25–17:30 BOOK LAUNCH · 5 MIN

Launch of the 2nd edition of “Künstliche Intelligenz und Data Science in Theorie und Praxis” (Springer Spektrum)

Prof. Dr. Göran Kauermann
Chairman, German Data Science Society (GDS)

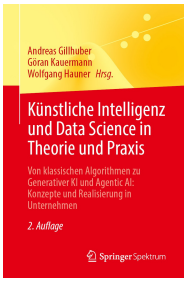
ABOUT THE BOOK

Edited by Andreas Gillhuber, Göran Kauermann and Wolfgang Hauner, all three members of the GDS board, this volume links the foundations of data science and AI with their use in companies. Part one covers methods and algorithms from statistics and computer science, including an introduction to modern AI models, and broader questions such as value creation, sustainability and data ethics. Part two turns to implementation: strategy, operations and the step from proof of concept to scaled production, also in light of the EU AI Act. Part three collects use cases from companies of all sizes, including projects with large language models and agentic systems.

The second edition adds foundations and examples of generative AI and agentic AI and covers reliable operation with MLOps and LLMOps. Many chapters are new, others were updated or fully revised.

Künstliche Intelligenz und Data Science in Theorie und Praxis. Von klassischen Algorithmen zu Generativer KI und Agentic AI: Konzepte und Realisierung in Unternehmen. 2nd edition, Springer Spektrum, October 2026. XX + 581 pages, in German. eBook and softcover.

Books can be ordered via `https://link.springer.com/book/10.1007/978-3-662-73474-2`. We also plan the publication of an English version.



Cover: Springer Spektrum

INDEX

Speakers and Chairs

Sorted by surname. Tap a name to open the talk or session.

A

Prof. Dr. Daniel Ambach — S3 p. 18, S7 p. 33
Chair of session S3; Chair of session S7; GDS Berlin Regional Group

Ulrich Arndt — S6.1 p. 29
Enterprise Architect, TUI

B

Jörg Bienert — S3.3 p. 21
Managing Director, Center for Sovereign AI (LEAM GmbH)

D

Mattia Dimauro — S2.3 p. 17
Founder, Predictabl

F

Sasa Fajkovic — S3.2 p. 20
Lead Applied AI, Trade Republic

G

Bettina Goerner — S4.1 p. 23
Chief Data Officer, Eterno Health

H

Martin Habedank — S2 p. 14, S6 p. 28
Chair of session S2; Chair of session S6; GDS board

Dr. Vasyl Hafych S6.2 p. 30
Senior Data Scientist, Sixt SE

Chris Heilmann S7.1 p. 34
VP Developer Relations, WeAreDevelopers

Steffen Hoellinger S5.2 p. 26
Field CTO, Confluent, an IBM Company

Christian Hoops S2.1 p. 15
PredictiveStatistics.eu / TU Dortmund

K

Prof. Dr. Göran Kauermann S7.2 p. 35
Chairman, German Data Science Society (GDS)

L

Zino Lang S6.3 p. 31
Founder, Avolang

Prof. Dr. Alexander Löser S1.2 p. 13
Berliner Hochschule für Technik (BHT), DATEXIS

M

Prof. Dr. Gerard de Melo S3.1 p. 19
Hasso Plattner Institute / University of Potsdam

R

Johannes Rieper S2.2 p. 16
Protiviti

Dr. Roland Roller S5.1 p. 25
DFKI (German Research Center for Artificial Intelligence)

S

Kai Sahling S5.3 p. 27
Founder and CEO, &act

Dr. Lukas Schack S6.1 p. 29
Principal Machine Learning Engineer, TUI

Niklas Schüschke S2.2 p. 16
AI Project Engineer, Mittelstand-Digital Zentrum Rostock

Prof. Dr. Anne Schwerk S1 p. 11 S5 p. 24 S4.1 p. 23
Host of the fireside chat; Chair of session S1; Chair of session S5; IU International University of Applied Sciences, GDS board

T

Zlata Tanovic S1.1 p. 12
Associate Director of Data Science, HelloFresh SE

W

Dr. Martin Weiss S6.4 p. 32
Thought Leader and Strategic Trusted Advisor, Dataciders SE

Z

Michele Zilli S6.1 p. 29
Principal of Immersive Technology, TUI

GERMAN DATA SCIENCE SOCIETY

Data Science Needs a Community

The German Data Science Society (GDS) e.V. connects people who work with data and AI in research, industry and the public sector. Its regional groups bring them together across Germany, at meetups and at conferences like this one.

Members join our events at reduced fees and with priority, and they help shape the work of the society.

BECOME A MEMBER

gds-society.de

Personal, student and corporate memberships

Become a Member

GDS GERMAN DATA SCIENCE SOCIETY