L3S Strategy 2030

Balancing AI ambition and accountability: A path to Trustworthy AI


L3S
AI Research

INHALT

0. PROLOGUE

We are commencing from a position of strength, reflected in key performance indicators such as successfully secured ERC grants, the strong h-index of many L3S researchers, and the growth in third-party funding and flagship projects. The recent establishment of the ambitious new Center for AI and Causal Methods in Medicine (CAIMed), along with the enhancement of our compute infrastructure through the AI service centre (KISSKI) and the expansion of AI transfer services for industry (DAISEC), reaffirms our commitment to advancing cutting-edge research in highly relevant domains.

The rise of Generative AI in recent years did not catch us off guard. However, the speed and magnitude with which Large Language Models for text, images, code, videos, and more have permeated public consciousness surprised us as well. This evolution has reinforced our belief that an L3S strategy should aim for medium-term goals, dividing them into concrete measures and milestones while monitoring their achievement.

Through a collaborative process, we discovered ways to guide L3S toward the future as an internationally recognised AI research centre, aiming to rank among the top-five nationally. It is an ambitious goal, but without ambition, strategy becomes nothing more than a continuation of the past.

Wolfgang Nejdl, Bodo Rosenhahn, Wolf-Tilo Balke, Jörn Ostermann, Sören Auer, Marcus Magnor, Marius Lindauer, Markus Fidler, Claudia Niederée, and Johannes Winter

May 2025

L3S 2030 — Balancing AI ambition and accountability: A path to Trustworthy AI

1. INTRODUCTION: L3S IN 2025

The contemporary AI landscape is defined by unprecedented advancements in model architectures and their integration into scientific and industrial applications. However, significant challenges remain in addressing complex cognitive tasks involving reasoning, strategic planning, and uncovering cause-and-effect relationships. Hybrid AI systems incorporate causal reasoning to discern the underlying mechanisms driving observed phenomena, thereby enhancing trust, reducing hallucinations, and fostering more reliable and responsible decision-making. Such hybrid approaches are exemplified by the L3S's interdisciplinary research portfolio encompassing knowledge graphs, large language models, retrieval-augmented generation, and trustworthy AI solutions.

According to the Stanford AI Index 2025, private investment in generative AI reached $33.9 billion in 2024, an 18.7% increase from 2023 and over 8.5 times the 2022 level. This significant funding fuels ongoing research and development in academia and industry. As an ambitious AI research centre with a dynamic research community comprising more than 200 researchers, over 25 professors, and an annual funding volume exceeding 15 million euros (Fig. 1), L3S addresses critical challenges and explores novel opportunities in personalised medicine, intelligent production, the future of education, and cybersecurity across regional, national, and global innovation ecosystems.

L3S maintains strategic partnerships with prestigious organisations, including the pan-European Confede-

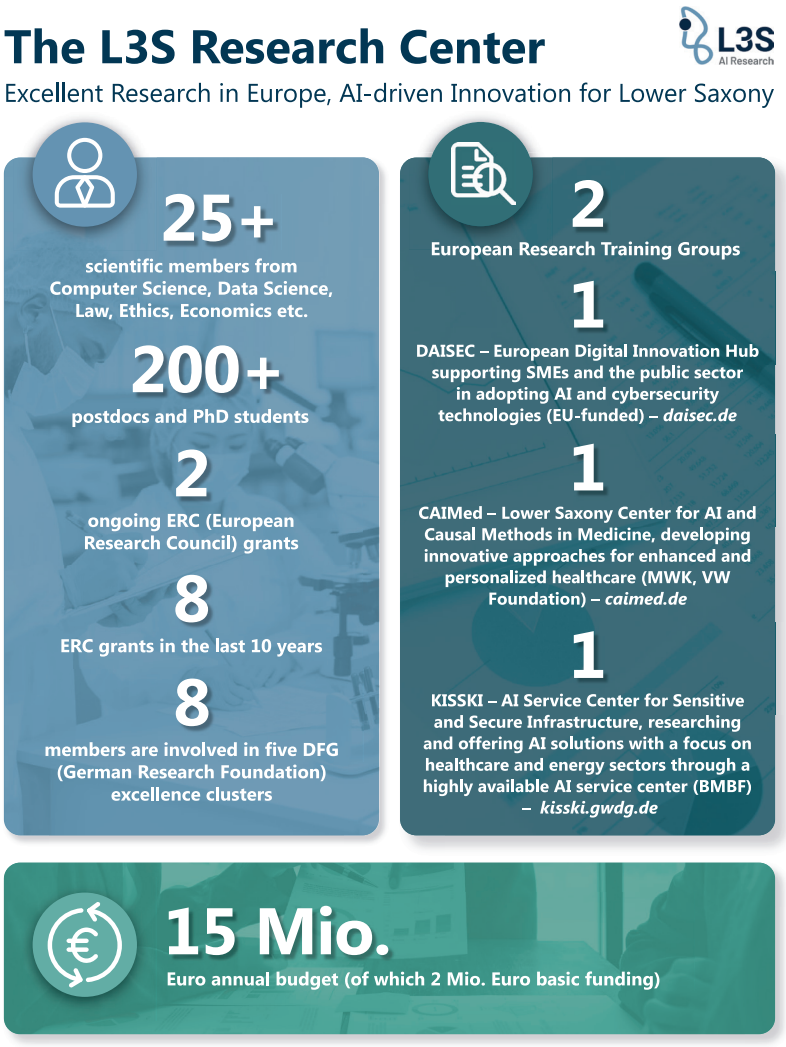

FIG. 1: L3S 2025 – Driving Trustworthy AI Research and Innovation.

ration of Laboratories for Artificial Intelligence Research (CAIRNE), the European Laboratory for Learning and Intelligent Systems (ELLIS), the Big Data Value Association, and the International Data Spaces Association.

The global regulatory landscape surrounding AI is marked by significant divergence. While the European Union has taken a proactive stance with the EU AI Act, emphasising a risk-based approach and stringent requirements for high-risk AI applications, other regions pose a critical challenge: balancing innovation with accountability. L3S adheres to the principles of reliable and responsible intelligent AI-systems and focuses on human-centred AI solutions, underlining L3S's commitment to trustworthy AI research.

Looking ahead, L3S envisions a future characterised by seamless human-AI interactions, fostering synergies that amplify both technological capabilities and human potential. By adhering to principles of explainability, trustworthiness, and reliability, L3S continues to develop cutting-edge, customised solutions that bridge the gap between fundamental research and practical applications. These efforts contribute significantly to scientific progress and societal benefit. At the intersection of global AI trends lies a pressing need to establish robust frameworks for human-AI interaction while maintaining rigorous scientific standards in development and deployment. L3S's efforts position it at the forefront of this endeavour, advancing both technological innovation and societal impact.

2. L3S 2030 — OUR VISION

L3S aims to be one of the top-five AI research centres in Germany by 2030, 30 years after its founding. This ambitious goal will be achieved through a combination of cutting-edge research, successful knowledge transfer, and the nurturing of exceptional AI talent, as well as the strong connection to CAIMed as a strategic extension of our activities in the context of Artificial Intelligence and Medicine. Benchmarks are the federal and state-funded AI centres BIFOLD (Berlin Institute for Foundations of Learning and Data), DFKI – German Research Center for Artificial Intelligence, MCML – Munich Center for Machine Learning, LAMARR – Institute for Machine Learning and Artificial Intelligence, ScaDS – Competence Center for Scalable Data Services and Solutions Dresden / Leipzig, and Tübingen AI Center – Competence Center for Machine Learning, as well as strong regionally driven initiatives such as hessian.ai and RWTH Aachen – Center for Artificial Intelligence (see Fig. 2).

By partnering regionally, nationally, and internationally, we aim to develop AI that benefits both research and society. We are committed to making significant contributions toward global development goals, including climate neutrality, prosperity, and social cohesion.

Diversity is already at the core of our research, including many international students and researchers from diverse social and cultural backgrounds. Therefore, we will encourage all researchers and students to take advantage of the extraordinary opportunities for international and diverse research, study and working conditions that Hannover offers, and make the L3S one of the best AI research centres in Germany and Europe. Within this context, we aim for at least 30% female researchers across all academic levels, ranging from students, PhD students, group leaders, up to professors.

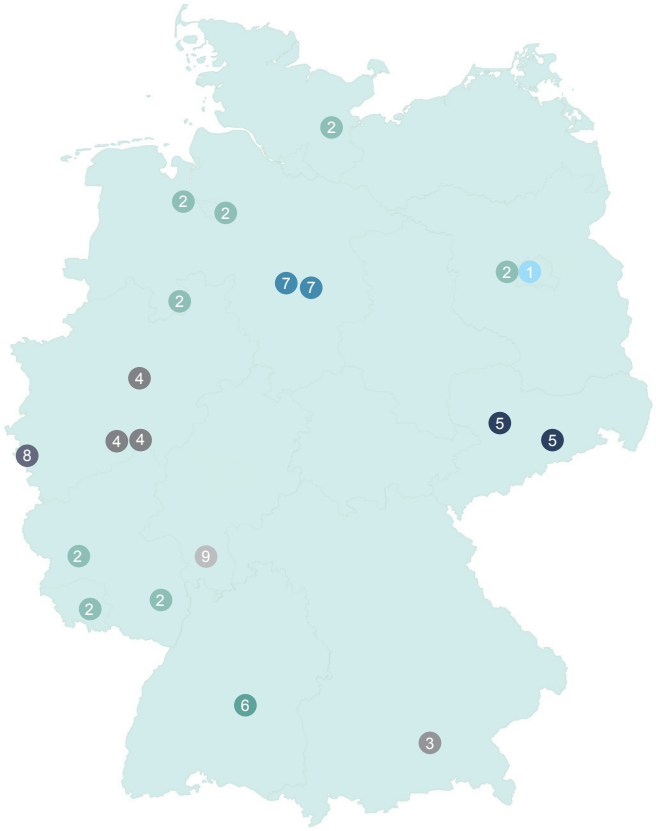

FIG. 2: Striving to be on a par with leading AI centres in Germany.

1. *BIFOLD – Berlin Institute for Foundations of Learning and Data*
2. *DFKI – German Research Center for Artificial Intelligence*
3. *MCML – Munich Center for Machine Learning*
4. *LAMARR – Institute for Machine Learning and Artificial Intelligence*
5. *ScaDS.AI – Competence Center for Scalable Data Services and Solutions (Dresden / Leipzig)*
6. *Tübingen AI Center – Competence Center for Machine Learning*
7. *L3S Research Center*
8. *RWTH Aachen – Center for Artificial Intelligence*
9. *hessian.AI – The Hessian Center for Artificial Intelligence*

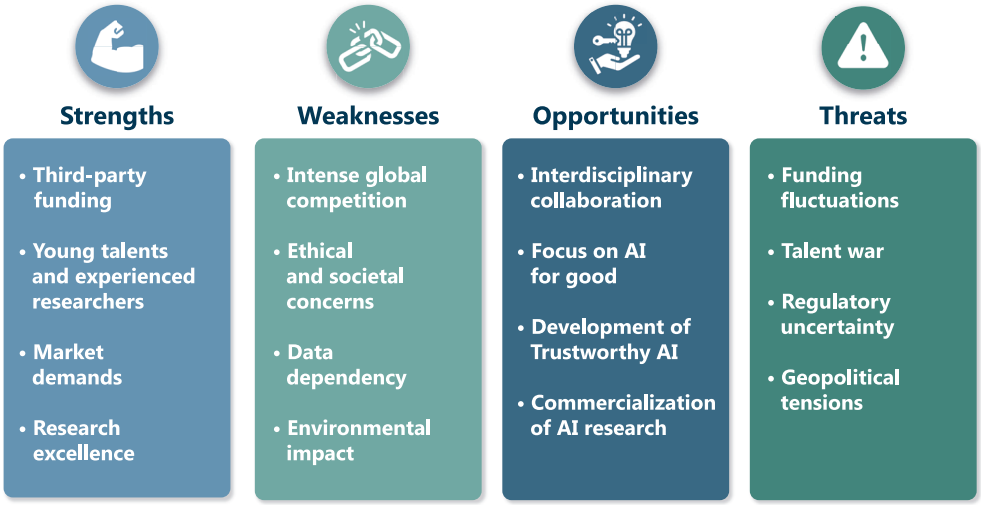

FIG. 3: Results of the L3S's SWOT analysis, 2023 / 2024.

The L3S will be recognised as an exceptionally well-managed research institution, attracting top international researchers and the best young scientists. As a globally connected institution, L3S will build on its reputation in AI research and scholarship both nationally and internationally. Our path forward will be characterised by innovation, collaboration and interdisciplinary synergy. Examples for interdisciplinary synergies and collaborative efforts are the participation of L3S members in Hannover's Clusters of Excellence, PhoenixD, Quantum Frontiers and Hearing4All.

In pursuit of these goals, our research centre will leverage and further develop cutting-edge AI technologies across active projects. To optimise the benefits of a culture rooted in intellectual inquiry and challenge, we will seamlessly integrate AI teaching and research, combining our efforts to address societal challenges through cross-disciplinary projects. Our ethos strongly values diversity, underpinning our commitment to inclusivity and providing opportunities for those who can benefit. Guided by robust ethical and technical principles for research and development of reliable and responsible intelligent AI-based systems and applications, we ensure that our values are embedded in every facet of our plans and actions.

As we tackle future challenges with confidence and ambition, we support the personal and professional development of our researchers and employees in a dynamic but also safe environment.

3. EXTERNAL FACTORS INFLUENCING L3S'S DEVELOPMENT

The SWOT analysis presented in Figure 3 provides a comprehensive overview of the external landscape within which L3S operates. This analysis reveals a complex interplay of opportunities and threats that will significantly shape L3S's strategic direction and capabilities in the coming years. By understanding these external factors, we can better position L3S to leverage its strengths, address weaknesses, and navigate the evolving AI research ecosystem as it strives to become one of Germany's premier AI research centres by 2030.

The global AI landscape presents significant opportunities for L3S through increased demand for intelligent solutions across healthcare, production, energy, mobility, and cybersecurity sectors. This expansion in market demand coincides with greater funding potential from both government initiatives and private sector investments seeking AI advancements that drive economic growth, sustainability, and social innovation. The rising popularity of AI and data science programs worldwide ensures a continuous influx of skilled researchers, particularly as Germany remains an attractive research destination amid global uncertainties.

Interdisciplinary collaboration represents another key external opportunity, as AI's integration with neuro-

science, economics, engineering, law, and ethics opens pathways to address complex societal challenges. The growing emphasis on "AI for good" in domains like healthcare, environmental sustainability, and disaster relief attracts both funding and talent to research with positive societal impact. These application areas align with L3S's focus on developing trustworthy AI systems that build public confidence and comply with emerging regulatory frameworks.

The competitive landscape presents substantial challenges, with fierce competition from leading global AI research institutions and companies necessitating clear differentiation through unique research agendas and impactful collaborations. This competition extends to human resources, with an intense "talent war" for top AI researchers in a highly competitive job market requiring research centres to offer attractive environments and competitive compensation.

The evolving regulatory landscape, particularly the EU AI Act with its risk-based approach and stringent requirements, creates compliance demands that competitors in less regulated regions may not face.

External economic and geopolitical factors further shape L3S's operating environment, with potential fluctuations in funding sources requiring diversified financial models and sustainable approaches. Environmental concerns regarding the high energy consumption of large AI models drive the need for energy-efficient algorithms and hardware solutions. Meanwhile, geopolitical tensions increasingly impact international research collaborations, challenging the principles of open science and knowledge sharing that have traditionally underpinned academic research networks. These external pressures collectively influence L3S's strategic pathway toward positioning itself among Germany's top AI research centres by 2030.

FIG. 5: The European Digital Innovation Hub on AI and Cybersecurity (DAISEC) effectively facilitates regional transfer of AI research into practical applications.

4. L3S — ADVANCING TO 2030

4.1. Our Mission

At L3S, we drive breakthroughs in trustworthy AI to create a positive global impact. We envision a future where artificial intelligence powers innovation and economic growth while safeguarding our planet's resources and fostering social equity. We partner with industry leaders, policymakers, and societal stakeholders to transform cutting-edge research into practical solutions. Through innovation strategies and evidence-based policy recommendations, we empower organisations and the public to thrive in the evolving AI landscape while enhancing competitiveness, sustainability, and social responsibility. Together, we are building AI systems that do not just push technological boundaries – they reshape them responsibly for a better tomorrow.

4.2. Our Approach to Trustworthy AI

We aim to explore and develop trustworthy AI systems for the benefit of humans, which requires clear values, competencies, and guiding principles for research and development. The L3S brings together the necessary core competencies in computer science, ethics, and law to develop reliable and responsible intelligent systems for digital transformation. In doing so, we operate on a firm foundation of values reflected in the following duality of research and development principles (see Fig. 4).

4.2.1. Accountability and responsibility

As artificial intelligence continues to play an increasingly significant role in decision-making across critical domains, it is imperative to integrate accountability and responsibility into AI systems. This encompasses the elimination of bias, the assurance of explainability, and the application of fairness-aware federated learning principles. Research at L3S prioritises the creation of AI governance frameworks that comply with legal standards, foster ethical decision-making, and ensure robustness against errors and exploitation. The European AI Act plays a crucial role in setting these legal standards and promoting responsible usage of AI technologies. Key initiatives include addressing bias in decentralised systems, developing AI security mechanisms to prevent malicious attacks, and establishing certification methods for AI models to ensure interpretability, regulatory compliance, and resilience to tampering. Implementing hybrid methodologies, such as Retrieval Augmented Generation, which enable large language models to access and refer to local document collections, databases, and semi-structured knowledge, is also crucial for promoting accountability. L3S is dedicated to ensuring that AI systems meet societal, business, and political trust requirements by embedding accountability and ethical standards throughout every stage of development and deployment.

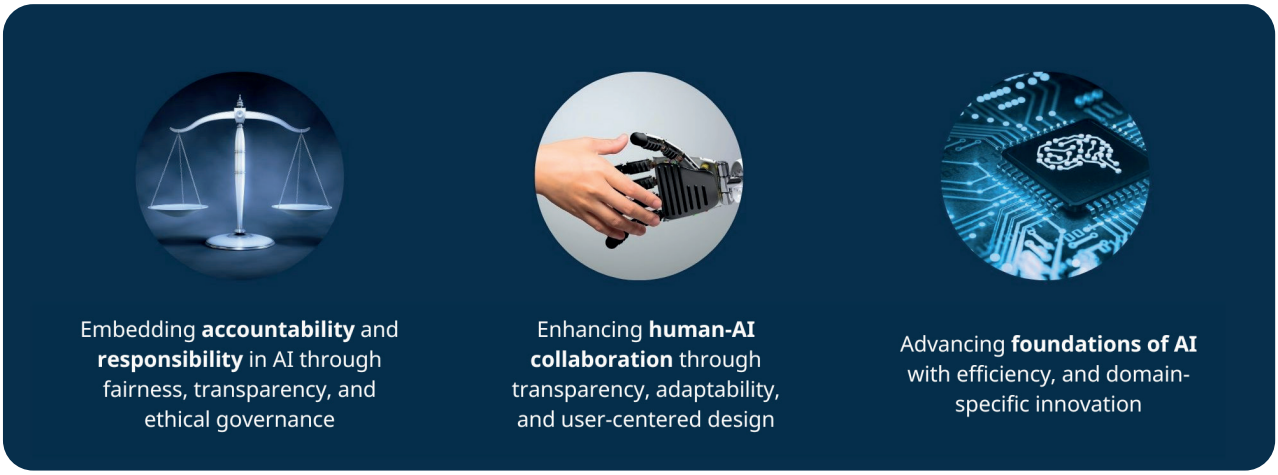

FIG. 4: L3S research – Developing Trustworthy AI.

4.2.2. Human-AI interaction

Enhancing collaboration between humans and AI is essential for developing reliable and user-centreed AI systems. L3S prioritises promoting usable safety, human-centreed AutoML, as well as immersive experiences such as analysing video recordings with action cameras, to improve practical applications in sectors such as healthcare, education, and industrial automation. The institute's research emphasises explainable AI to ensure AI-driven tools are transparent and adaptable. These tools will support knowledge transfer and enable intuitive user interactions with AI systems. This includes the development of multimodal interfaces and reinforcement learning methods that tailor responses to the user's intent and context. The participatory design approach allows individuals from various domains to influence AI behaviour, ensuring its adaptation to real-world needs. Through these initiatives, L3S aims to create AI technologies that augment, rather than replace, human decision-making.

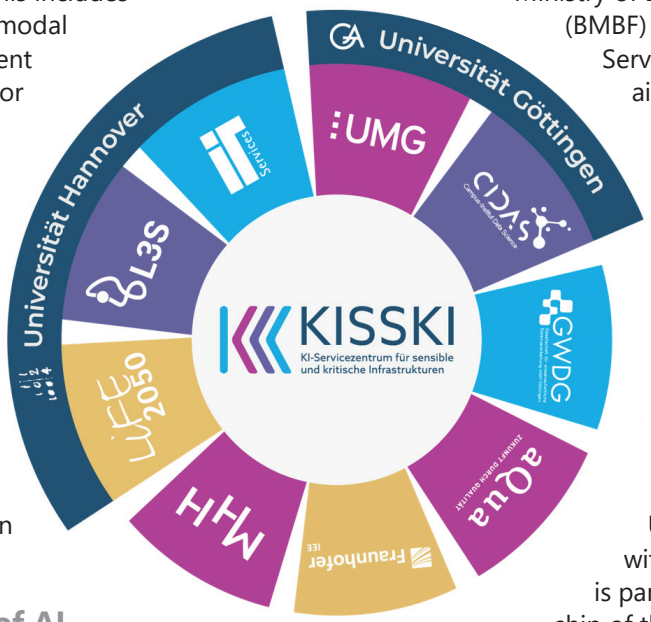

FIG. 6: Areas of expertise in the KISSKI consortium: L3S stands for cutting-edge AI methods.

4.2.3. Foundations of AI

Foundation models, encompassing large language models and multimodal systems, are transforming AI research and its applications. The L3S prioritises the development of these models with a focus on domain-specific applications, including medicine, cybersecurity, and energy-efficient computing. Fundamental research in machine learning, natural language processing, information retrieval, video analysis, and semantic technologies forms the basis for innovative approaches and solutions. L3S also promotes usable safety, interactive AutoML, and explainable AI that ensures transparency and adaptability.

As L3S, we pursue an interdisciplinary approach through collaborations across society, science, and business domains. The centre strengthens connections between computer science and medicine, mechanical engineering, law, and social sciences. Impact extends beyond academia through startup incubation (Digital Health City Hannover, Life Science Factory) and advisory services via the European Digital Innovation Hub on AI and Cybersecurity (DAISEC, see Fig. 5) and CISPA Helmholtz Center. In healthcare, initiatives by L3S, such as CAIMed – Lower Saxony Center for Artificial Intelligence and Causal Methods in Medicine, integrate clinical records, imaging, and genetic data into AI-driven systems aimed at optimising diagnostics and treatments.

Since the fall of 2022, the German Federal Ministry of Education and Research (BMBF) has been funding four AI Service Centers in Germany, aiming to facilitate access to computing infrastructure and AI expertise while promoting the transfer of AI to a broader audience. One of these centres is the joint project "AI Service Center for Sensitive and Critical Infrastructures" (KISSKI; led by L3S Director Prof. Rosenhahn), in which Leibniz University Hannover, along with the L3S Research Center, is participating. Under the leadership of the University of Göttingen, this project is funded by the BMBF for a period of three years, to be continued in the future. Other project partners include the Society for Scientific Data Processing (GWDG) Göttingen, aQua – Institute for Applied Quality Promotion and Research in Healthcare GmbH, Göttingen, and the Fraunhofer Institute for Energy Economics and Energy System Technology, Kassel (see Fig. 6).

Despite having less funding than national AI competence centres, L3S ranks sixth among German AI institutes based on publications at top-tier venues (3.77 Rank-1 publications per member between 2020 and 2024). To reach its goal of becoming a top-five AI research centre by 2030, L3S will expand AI professorships to strengthen its research potential while continuing to leverage its expertise in data science, fundamental AI research, and interdisciplinary applications.

5. L3S 2030 – BALANCING AI AMBITION AND ACCOUNTABILITY: A PATH TO TRUSTWORTHY AI

For years, L3S members have been organising top-tier international AI conferences (rated A/A* by Computer Science Conference Rankings / CORE) and publishing high-impact research in leading journals (Q1 by SCImago Journal Rankings). This enhances Germany and Europe's reputation as a key location for cutting-edge research and innovation in AI. The L3S is a nationally and internationally recognised centre for research and education in data science and AI, bridging science and industry across various regions, as highlighted by its national and international collaborations.

The "h-Index list (according to Google Scholar)" of L3S members currently ranks Thomas Illig (162), Wolfgang Nejdl (79), Sören Auer (72), Marcus Magnor (62), Bodo Rosenhahn (58), Sándor Fekete (56) and Jörn Ostermann (53), to name the most highly ranked members. In addition, Prof. Nejdl is listed among the top fifty scholars worldwide in the AMiner portal's "Most Influential Scholar Annual List" in the field of Web and Information Retrieval, and he is ranked 807th worldwide (national: 29th), while Sören Auer is ranked 2661st (national: 124th) among the best scientists in computer science for the year 2023, according to research.com, a leading academic platform for researchers.

Building on this strength, we will take the following strategic actions to fulfil our mission of being a leading AI research institution and think tank of international reputation (see Fig. 7):

5.1. Research and Compute Infrastructure

L3S provides cutting-edge AI research infrastructure, including high-performance computing clusters, GPUs, and specialised hardware. Investments in national and international projects have built a competitive infrastructure, primarily used within individual research groups. Future plans include partial resource sharing to foster collaboration. An initial resource inventory will be compiled with technical administrators, followed by regulations to ensure fair access and prioritisation.

The joint AI Service Center KISSKI (with the University of Göttingen) operates an HPC system with NVIDIA A100 and H100 GPUs for research, development, and technology. Resources are geo-redundantly installed at Leibniz University IT Services (LUIS) in Hannover and the GWDG Göttingen. Models and data can be integrated via self-service or KISSKI catalogues.

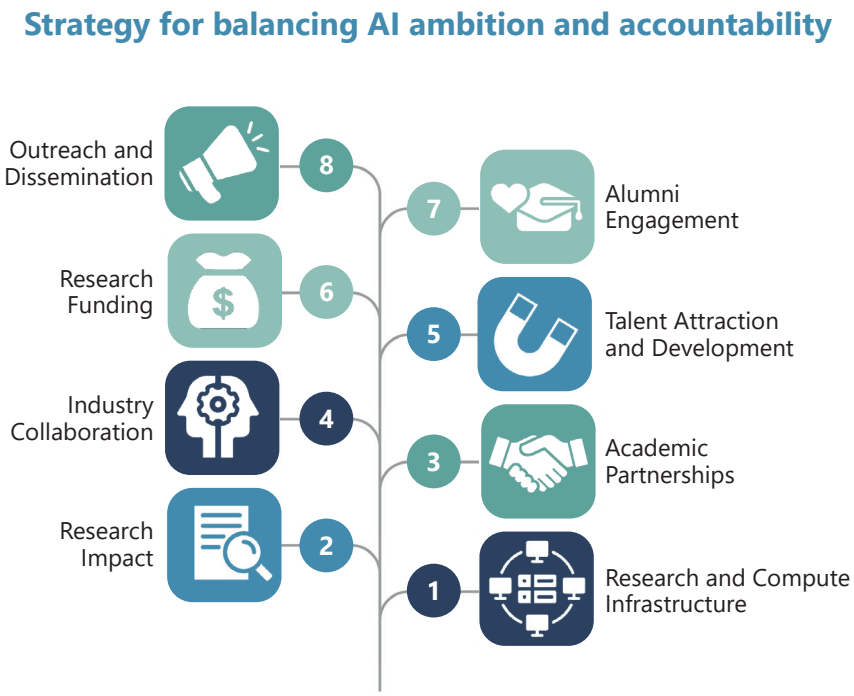

FIG. 7: Advancing 2030: L3S's strategic fields of action.

The SoBigData Research Infrastructure (RI) enhances national research systems within a federated European platform. It enables large-scale data science and social mining experiments in the cloud (aligned with EOSC guidelines) and leverages supercomputing. The platform makes social data analysis more efficient, repeatable, and accessible to users from various fields. Its approach follows FAIR and FACT principles.

5.2. Research Impact: How we Measure our Scientific Excellence

L3S publishes in high-ranking journals (Q1, SCImago) and conferences (A*, CORE) and fosters collaborations with top institutions. To increase visibility and citation rates, research findings are also shared via open-access journals and preprint archives.

L3S identifies emerging AI research areas such as Trustworthy AI, Generative AI, and AI for Medicine, integrating them into initiatives like CAIMed (Lower Saxony Center for AI and Causal Methods in Medicine). The centre is funded by the Lower Saxony Ministry for Research and Culture and the Volkswagen Foundation and is driven by scientists from the L3S (lead), Hannover Medical School, the Helmholtz Centre for Infection Research Braunschweig / Centre for Individualised Infection Medicine (CiiM) as well as the Georg-August-Universität Göttingen / Campus Institute Data Science (CIDAS) and the University Medical Center Göttingen. The centre develops innovative AI methods for personalised healthcare and combating cancer, cardiovascular diseases, and infections by using multimodal AI (Symbolic & Causal Representations, Neural & Deep Learning Based Representations, Connected to the Scientific Body of Knowledge). CAIMed links research, clinical, and patient data to improve prevention, diagnostics, and therapy.
The combination of excellent locations in Lower Saxony for methodological AI research, data-intensive medicine, medical informatics and basic medical research creates a unique collaborative space for basic research, knowledge transfer and research-driven innovations in the healthcare system (see Fig. 8).

5.3. Academic Partnerships

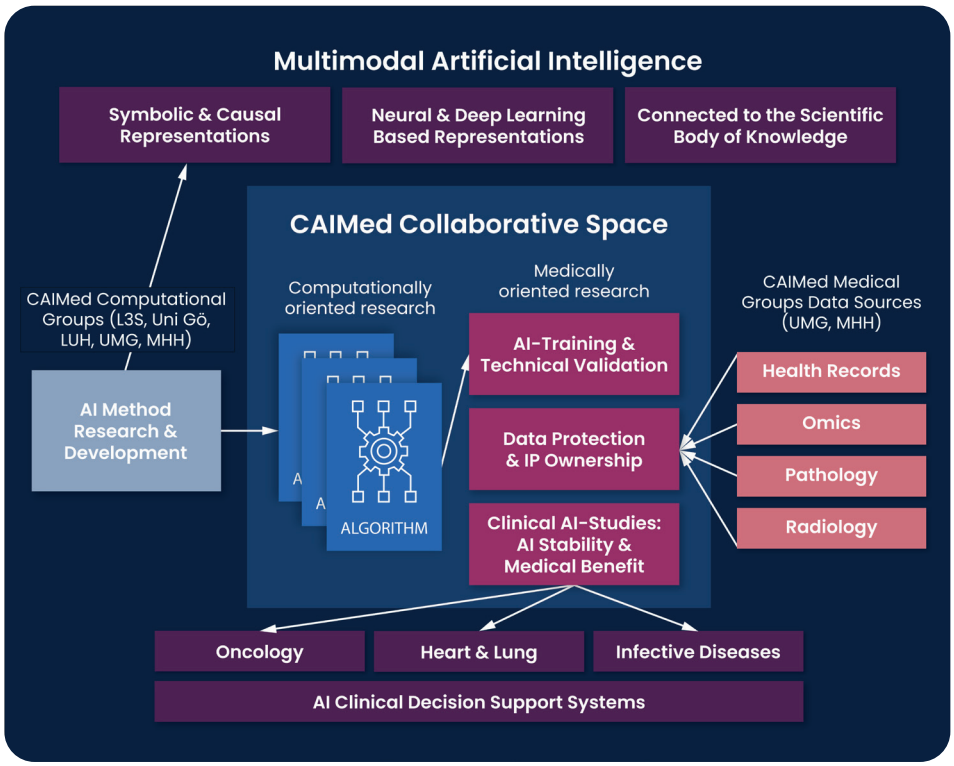

FIG. 8: The CAIMed – Lower Saxony Center for AI and Causal Methods in Medicine Collaborative Space.

L3S actively collaborates with leading AI research institutions worldwide, including the AI centres at RWTH Aachen and TU Darmstadt (hessian.AI). These partnerships foster high-impact research and joint publications. L3S members organise and present at top AI conferences, including the 18th ACM International WSDM Conference (Hanover, 2025), enhancing visibility and fostering collaborations. We are deeply integrated into international AI networks such as ELLIS, CAIRNE, and AI Grid. Efforts to expand ELLIS membership have strengthened ties within the European AI research community.

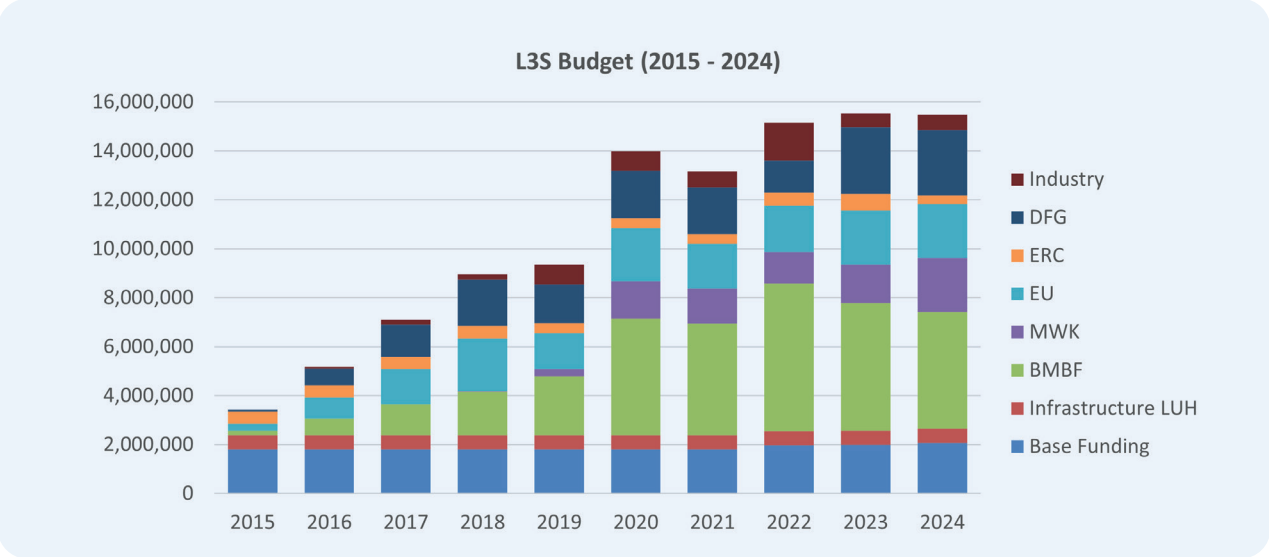

FIG. 9: Development of Third-Party Funding (2015 – 2024).

As a Gold Member of CAIRNE, L3S reinforces its European AI leadership. L3S actively engages with acatech (Germany's National Academy of Science and Engineering), contributing to AI policy, agenda-setting, and national research strategies. Collaboration with open-source AI organisations is a strategic priority, enhancing research capabilities and fostering innovation. Further engagement with key figures will deepen these partnerships. Interdisciplinary cooperation between L3S computer scientists and legal scholars has intensified through joint proposals and initiatives.

5.4. Industry Collaboration

L3S collaborates with leading AI-driven industries, including Volkswagen, Sartorius, Bosch, Intel Labs, and Infineon, fostering technology transfer in projects like IIP Ecosphere. Through the European Digital Innovation Hub (DAISEC), L3S transfers AI and cybersecurity research into SME applications via training, advisory services, and joint projects.

L3S supports high-tech incubators (Life Science Factory, Venture Villa, Starting Business, SMINT) to commercialise AI research through startups and licensing agreements.

To bridge academia and industry, L3S is establishing an Industrial Advisory Board (IAB), complementing the Scientific Advisory Board (SAB). The IAB will align research with industry needs, ensuring AI innovations translate into real-world applications. By actively engaging with industry leaders, L3S strengthens knowledge transfer, facilitates funding opportunities, and accelerates the adoption of cutting-edge AI solutions across sectors. Collaboration with industry partners also provides researchers with valuable insights into real-world challenges, refining AI methodologies for broader impact.

5.5. Talent Attraction and Development

People are the key value of an organisation; their skills and commitment are crucial to the success of research efforts. That's why the L3S is dedicated to nurturing and leveraging its talented team, ensuring that their expertise and dedication drive both academic and industrial advancements. Efforts include attracting renowned and emerging researchers, applying for an Alexander von Humboldt Professorship, and providing mentorship for early-career researchers.

L3S supports the national AI Grid program, enhancing networking between young and experienced AI researchers across Europe. As an official AI Grid partner, L3S collaborates with top German, European and Global AI institutions, with key members serving as mentors and mentees. To further strengthen talent development, L3S considers launching its own fellowship program, providing funding and resources to distinguished AI researchers, modelled after the Lamarr Fellows program, and potentially supported by the Ministry for Science and Culture (MWK).

5.6. Research Funding

L3S actively secures competitive research grants from national and international agencies, maintaining its position among Europe's most successful AI centres with eight ERC grants over the past decade. Institutional funding (2 million € per year, see Fig. 9) is supplemented by third-party funding, as illustrated in funding growth trends. We attract top talent and resources through structured study programs, MWK-supported graduate schools (Hybrint), DFG clusters of excellence (PhoenixD, Hearing4All, QuantumFrontiers), and initiatives like AI Grid and the DAAD Postdoctoral Networking Tour in AI.

Strengthening research consortia remains a priority, with DFG proposals such as the SciKnow Collaborative Research Center and an International Research Training Group in partnership with the Indian Institute of Technology Kharagpur. Additionally, L3S actively engages with the Ministry of Science and Culture (MWK) to align strategic objectives, collaborating with leading AI institutions in Germany and federally funded AI centres. Discussions at state and federal levels aim to establish a long-term vision for L3S as a leading AI research centre with robust resources and partnerships.

5.7. Alumni Engagement

We actively maintain connections with alumni who have advanced in AI research, leveraging their contributions to our research network and potential funding opportunities. We continue to nurture these relationships through seminars, conferences, research collaborations, and joint publications. To strengthen alumni ties, we will launch an annual homecoming event starting in 2025, providing a platform for alumni to reconnect, share experiences, and explore collaboration opportunities. Additionally, we will enhance support for networking, workshops, and research seminars to foster ongoing engagement throughout the year.

5.8. Outreach and Dissemination

L3S maintains a strong commitment to outreach and dissemination, integrating research, transfer, and education to foster public understanding and responsible engagement with AI. Through participation in internationally and nationally recognised initiatives – such as ELLIS, CAIRNE, and Plattform Lernende Systeme – and through transfer-oriented projects – including the European Digital Innovation Hub (DAISEC), the AI Service Center (KISSKI), and the demonstration, research, and transfer networks for production (ProKI and ReGap) – L3S demonstrates the real-world applicability of AI. These efforts are supported by demonstrators developed in collaboration with academic and industrial partners, which are presented at major industry events like Hannover Fair and EMO. The Mittelstand-Digital Zentrum Hannover (MDZH) initiative further exemplifies L3S's approach by offering tailored formats such as AI company dialogues and pilot projects, fostering innovation across regional industries. These activities underline the centre's pivotal role in connecting AI research with practical applications and industrial needs.

To ensure broad impact, L3S leverages a range of media and event formats, including the widely read BINAIRE magazine and high-engagement digital platforms, to reach varied target audiences. The centre contributes to international discourse by organising specialist symposia and participating in global events, such as the NII Shonan Meeting in Japan. Initiatives like "Bringing AI Knowledge to Schools" highlight the centre's long-term vision of democratising AI knowledge and preparing future generations. Internally, L3S has strengthened its visibility through the successful launch of the AI Day at Leibniz University of Hannover, reinforcing the call to establish AI as a strategic focus within the university and in broader contexts. The redesigned website and active GitHub repository support transparency and collaboration, while platforms and services like Open Research Knowledge Graph and ORKG ASK illustrate L3S's commitment to structured knowledge dissemination. Collectively, these efforts consolidate L3S's reputation as a nationally and internationally recognised think tank and innovation driver in the AI ecosystem.

6. FUTURE PERSPECTIVES: SHAPING TOMORROW'S AI

As we approach 2030, L3S will persistently pursue key strategic areas with adaptability, staying at the forefront of the fast-evolving AI landscape. We will continue refining our research strategy while exploring new and emerging areas where we can make impactful contributions. Our research focuses on both expanding existing fields and addressing new challenges. Promising areas of interest include:

- **Trustworthy AI:** This area continues to grow in importance as ethical considerations and regulatory frameworks, like the EU AI Act, shape the future of AI. Our work spans transparency, fairness, explainability, and security in AI systems, with a focus on mitigating biases and ensuring accountability.

- **Agentic and Embodied AI:** We are advancing autonomous AI agents capable of decision-making and problem-solving in real-world environments. Our research aims to enhance productivity in sectors like healthcare and customer service while addressing safety, transparency, and ethical concerns.

- **Hybrid AI:** Combining neural and symbolic approaches, hybrid AI integrates pattern recognition with logical reasoning, offering enhanced transparency and adaptability in AI applications. We are developing tools for multimodal fusion and knowledge graph integration, with applications in personalised medicine, autonomous systems, and more.

- **Biological and Medical Foundation Models:** These large-scale AI systems aim to revolutionise medical diagnostics, drug discovery, and personalised treatments. Despite challenges like data accessibility and regulatory concerns, we believe these models hold immense potential for advancing healthcare and biomedical research.

- **Energy-Efficient AI:** As AI systems grow in complexity, we are focusing on reducing their energy consumption through techniques like network pruning, knowledge distillation, and the use of alternative hardware. Energy-efficient AI is critical for applications in autonomous driving, industry 4.0, and medicine.

- **AI for Science:** L3S envisions AI as a collaborative partner in scientific discovery, not a replacement for human researchers. Our focus is on AI that augments human creativity, facilitates new insights from large datasets, and accelerates the scientific process through enhanced reasoning and transparency.

By staying true to our core values of curiosity, research excellence, and ambition, L3S is poised to shape the future of AI, ensuring it is both groundbreaking and aligned with societal needs.

7. EPILOGUE: UNVEILING A DYNAMIC L3S STRATEGY FOR A COMPLEX WORLD

L3S is positioned at the forefront of AI research, contributing to innovations that drive industrial transformation and improve global quality of life. Our interdisciplinary team of around 200 researchers from diverse fields is committed to creating AI systems that are trustworthy, ethical, and beneficial for society. By fostering research that aligns with legal, sustainability, and societal standards, we aim to push the boundaries of data science while ensuring our work adheres to the highest ethical principles. Our strategy focuses on continued innovation, dynamic adaptation to emerging trends, and partnerships that advance AI for the public good.

Our goal is to rank among Germany's top-five AI research centres by 2030, building on existing strengths and exploring promising new frontiers. Flagship initiatives such as KISSKI, DAISEC, and CAIMed will be key to our progress, supporting groundbreaking research in critical fields like cybersecurity, healthcare, and infrastructure. By continuing to prioritise education, research, and knowledge transfer, we are committed to pushing the limits of AI to benefit humanity.

BOARD OF DIRECTORS

Prof. Dr. Wolfgang Nejdl
Executive Director

Prof. Dr.-Ing. Bodo Rosenhahn
Director

Prof. Dr. Wolf-Tilo Balke
Director

Dr. Claudia Niederée
Managing Director, Group Leader

Dr. Johannes Winter
Chief Strategy Officer, Deputy Managing Director, Member of Extended Board of Directors

Prof. Dr. Sören Auer
Member of Extended Board of Directors

Prof. Dr.-Ing. Markus Fidler
Member of Extended Board of Directors

Prof. Dr. Marcus Magnor
Member of Extended Board of Directors

Prof. Dr. Marius Lindauer
Member of Extended Board of Directors

Prof. Dr.-Ing. Jörn Ostermann
Member of Extended Board of Directors

APPENDIX

Future Perspectives:
As we move towards 2030, L3S is committed to pursuing our defined strategic areas with determination. We will proceed in stages to not only implement our research and transfer strategy consistently but also to continually evaluate and adjust it as needed. A dynamic approach is crucial because the last 70 years of AI research have shown that significant progress and breakthroughs are accompanied by disappointments and the need for adjustments. Therefore, L3S will be persistent where we anticipate breakthroughs and want to advance our research with high intensity, while remaining adaptable where new insights from research and practice require corrections. This blend of rigour and flexibility has made us successful over the past 20 years, and we are eager to take the next steps. Beyond existing research and projects, we have identified key areas of interest in AI. Either because we are already conducting extensive research in these areas, such as Trustworthy AI, Explainable AI, Causal AI, etc., which we want to intensify, or because weak signals suggest that new fields in AI research are emerging that we want to actively shape. These promising areas include:

Trustworthy AI.
Trustworthy AI has established itself as an important area of research and innovation in the field of artificial intelligence, addressing the need for reliable, ethical AI systems that are aligned with societal values. With the enactment of EU regulation, in particular the EU's AI Act, trustworthiness has become a legal prerequisite for AI generation and development. This area focuses on the development of technologies that are fair, transparent, robust, safe and respectful of privacy and human rights, and also represents the core principle of L3S. Hybrid AI methods, such as Retrieval Augmented Generation (i.e. combining LLMs with knowledge graphs), combine knowledge-based and data-driven approaches to create more powerful and flexible AI systems. As AI systems are increasingly used in decision-making processes in various sectors, ensuring their trustworthiness is of paramount importance for their legality and their broad acceptance in society. Research in fairness-aware federated learning and bias mitigation strengthens AI's ethical foundations, particularly in decentralised environments. Our research in this area now and in the future covers several dimensions, including explainability, accountability, bias mitigation and security against adversarial attacks.

Agentic and Embodied AI.
AI agents are developed to create autonomous systems capable of perceiving, reasoning, and acting in complex environments to achieve specific goals. The primary approaches to building these agents include leveraging foundation models and generative AI, employing multi-agent systems for collaborative problem-solving, integrating tools and APIs for real-time interaction, and utilising reinforcement learning for decision-making. However, challenges such as data quality and bias, ethical considerations, security and safety concerns, transparency, integration with existing systems, and scalability must be addressed to ensure the effective functioning of AI agents. Soon, AI agents are expected to enhance productivity by automating routine tasks and providing personalised experiences in sectors like customer service and healthcare. Long-term potential includes the development of fully autonomous systems managing complex operations, fostering human-AI collaboration in creative tasks, democratising access to specialised knowledge, and significantly transforming societal structures. As technology advances, it will be essential to tackle the ethical and societal implications of AI agents to maximise their benefits while minimising risks. With the broad activities of the L3S, including strong research on robotics, LLM and foundation models, computer vision, reinforcement learning and predictive control, we are in an excellent position to strive towards high impact in this area.

Hybrid AI.
Hybrid AI models integrating neural and symbolic approaches aim to combine the capabilities of neural networks' pattern recognition with symbolic AI's logical reasoning. This integration addresses limitations in current AI systems, particularly in transparency, adaptability, and domain-specific knowledge integration. The state-of-the-art demonstrates advanced capabilities across industries, including healthcare, autonomous vehicles, manufacturing, and education. Key technical innovations include the use of knowledge graphs to provide structured context for neural networks, multimodal fusion combining visual and textual data processing with logical rule extraction, and modular integration enabling communication between neural and symbolic components. However, hybrid AI models face challenges such as integration complexity, high data requirements, computational overhead, interpretability trade-offs, and ethical risks. These models also pre-

sent opportunities in areas like personalised medicine, quantum-hybrid systems, edge AI deployment, automated scientific discovery, and context-aware cybersecurity. As integration tools mature, particularly in neurosymbolic frameworks and cognitive architectures, hybrid AI models may significantly affect AI applications in critical domains by improving explainability and adaptive learning.

Biological and Medical Foundation Models.
Such models are driven by the need to tackle complex healthcare and biomedical research challenges using advanced AI techniques. These large-scale AI systems, trained on extensive datasets, aim to revolutionise medical diagnostics, drug discovery, and personalised treatments. The current state-of-the-art includes language models for biomedical text processing, vision models for medical imaging, multimodal models integrating various data types, and bioinformatics models for molecular-level analysis. However, these models face significant challenges, including data accessibility, computational costs, interpretability issues, ethical considerations, regulatory compliance, and integration with existing healthcare systems. Despite these challenges, foundation models offer immense potential for advancing medical therapies and research. They promise to enable more personalised medicine by integrating diverse patient data, accelerate drug discovery processes, provide advanced clinical decision support, enhance biomedical research capabilities, improve precision diagnostics, and offer a more holistic understanding of biological systems. As these technologies continue to evolve and address current limitations, they have the potential to usher in a new era of data-driven, personalised medicine that could significantly improve patient outcomes and accelerate scientific discoveries in the medical field.

Energy-Efficient AI.
Energy-efficient AI comprises techniques to reduce the energy demand for both the training procedures as well as the online inference. The increasing energy hunger of AI is well known, and the plans to use power plants next to data centres raise severe concerns. Several techniques for energy-efficient AI are highly promising, e.g. by applying network pruning techniques, knowledge distillation or by selecting appropriate and most useful data for training itself. Even further, it is possible to use alternative hardware for parts of deep neural networks, e.g. diffractive optical devices to model convolutions (photonic chips), or by using dedicated hardware (quantum chips, FPGAs, neuromorphic chips) for hybrid optimisation and efficient inference. Such techniques have a direct impact on many applications in industry 4.0, autonomous driving, medicine or everyday tools on a smartphone and will be of utterly importance in modern systems.

AI for Science.
AI for Science represents a paradigm shift, moving beyond AI as a mere tool to AI as a collaborative partner in the scientific discovery process. At the L3S, we envision a future where researchers across disciplines are augmented by sophisticated AI capabilities, not replaced by them. This human-centreed approach focuses on developing AI systems that enhance human intuition, creativity, and expertise. We will concentrate on AI that can manage and analyse vast, heterogeneous datasets (from simulations, experiments, literature, and real-world observations), identify non-obvious patterns and correlations, generate novel hypotheses, design and even autonomously execute experiments, and accelerate the iterative cycle of scientific inquiry. Crucially, these AI systems will be designed to explain their reasoning and provide researchers with transparent insights, fostering trust and allowing scientists to retain ultimate control and understanding. The research will focus, among others, on causal inference and causal representation learning, Bayesian deep learning, Bayesian optimisation, uncertainty quantification, scientific knowledge graphs, and hybrid methods that seamlessly integrate domain expertise with data-driven learning. The goal is not to automate science, but to empower scientists, enabling them
to tackle more complex problems, make breakthroughs faster, and ultimately, deepen our understanding of the world.

Forschungszentrum L3S
Appelstraße 9a
30167 Hannover

AUTHORS:
The Extended L3S Board of Directors
The L3S Members
The L3S Group Leaders

EDITORS-IN-CHIEF:
Prof. Dr. Wolfgang Nejdl
Prof. Dr.-Ing. Bodo Rosenhahn
Prof. Dr. Sören Auer
Prof. Dr. Marius Lindauer
Dr. Johannes Winter

ILLUSTRATIONS & EDITORIAL ASSISTANCE:
Alina Nekrasova
Annika Schumann

EDITORIAL DESIGN:
Priska Tosch
Falko Lohrenscheit
www.tosch-kommunikation.de

IMAGE CREDITS:
L3S, unless otherwise noted
Cover image created by Falko Lohrenscheit using Midjourney

PRINTED BY:
Ströher Druckerei und Verlag GmbH & Co. KG
www.stroeher-druck.de

EDITORIAL STATUS:
May 2025

L3S
AI Research

www.L3S.de