WHITEPAPER 2023

The rise of industrial explainable artificial intelligence (XAI) – **Insights across the AI life cycle**

This whitepaper is intended to be beneficial for decision makers, product and project managers as well as AI experts and developers.

SIEMENS

Contents

Brief Summary

In this whitepaper we shed light on several pressing questions surrounding the increasing use of explainable artificial intelligence (XAI) in manufacturing applications.

Explainability is quickly becoming an essential component of industrial-grade AI. It is necessary throughout the AI life cycle, from business considerations to monitoring and maintenance of productive AI systems. The benefits and difficulties of explainability are interwined with AI and domain expertise. Additionally, explainability might contribute to meet the transparency and human oversight requirements of the upcoming AI Act.

1. Key findings

As the impact of artificial intelligence in manufacturing environments continues to increase, Industrial-grade AI and the related field of explainable AI (XAI) will play an increasingly important role. Furthermore, explainability is expected to play a role in meeting the transparency and human oversight requirements of the upcoming European Union AI Act. Our key findings are summarized below.

Figure 1: Key findings

2. Introduction

The emerging role of Industrial XAI

Artificial intelligence (AI) is taking the world by storm. It is also fast becoming a reality in many industries, including those operating in manufacturing environments.

For effective implementation in industrial production, it is essential that AI models fulfill certain requirements specific to the industry. Key characteristics of these so-called **Industrial-grade AI** systems include reliability, robustness, testability, traceability, certifiability and explainability. The ability to efficiently scale available resources, processes and technologies is equally important. Despite growing AI pervasiveness in manufacturing, moving from an initial concept to a widely deployed AI model of industrial scale still poses significant challenges. On top of that the upcoming EU AI Act will also require explainability for high-risk use cases which we expect to become relevant for industrial use cases that relate to safety as stated in article 13: "High-risk AI systems shall be designed and developed in such a way to ensure that their operation is sufficiently transparent to enable users to interpret the system's output and use it appropriately." (European Commission, 2021, p. 50).

Special focus of this white paper was directed at explainable AI (XAI), which serves as a key attribute both for Industrial-grade AI as well as the development process itself. Explainable AI is a relatively new and emerging concept in manufacturing and **refers to key factors that influence the results of AI systems in a way that allows human users to understand and trust them.**

XAI has already proven its efficacy in academic disciplines, delivering impressive results with image, text, tabular and time series data. The potential advantages of XAI for manufacturing environments and machine learning operations (MLOps), by contrast, have found only minimal consideration as of yet.

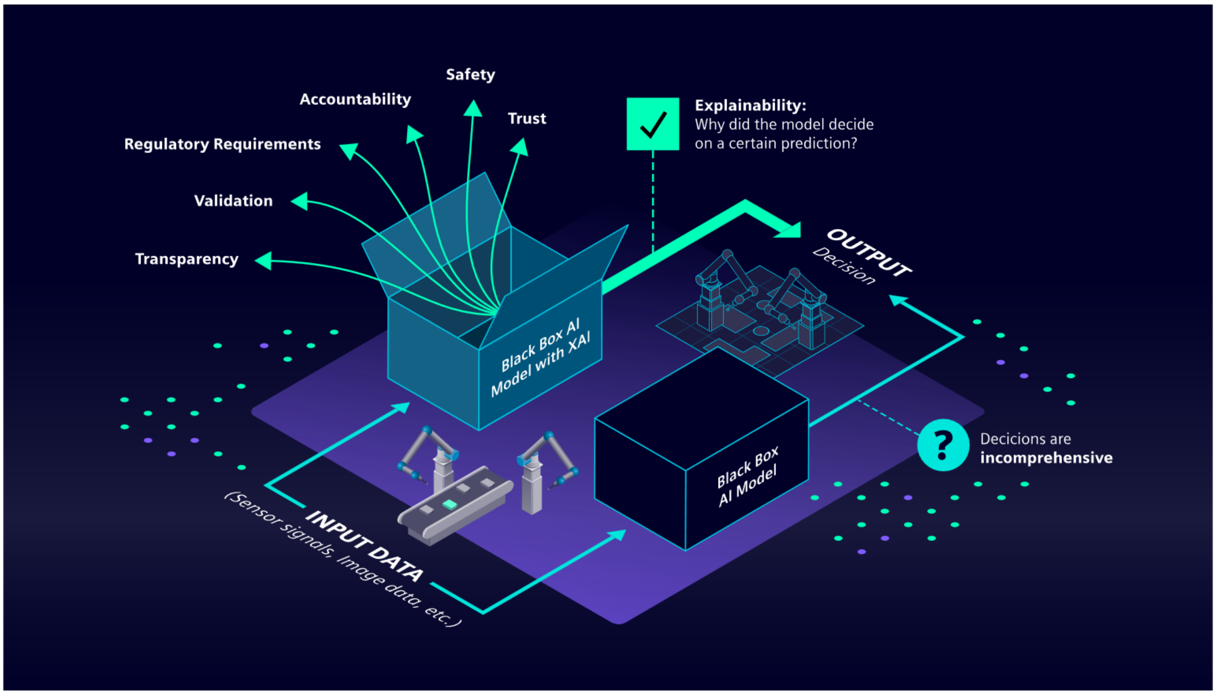

Figure 2: Conceptual image that compares Industrial AI with and without explainability at the shopfloor

3. The Whitepaper research approach

Listening to what the experts are saying

To explore the relevance of XAI across the entire AI life cycle of production industries, we conducted a series of in-depth interviews with 36 key experts and decision-makers. The group included both Siemens participants (21) as well as external authorities from various domains not affiliated with Siemens (15).

Each interview lasted between 45 and 60 minutes. The discussions were recorded and followed a predefined questionnaire, but also left room for responses that were not related to the original guidelines and questions.

Role / Domain	Industrial Automatization	Technical Associations	Auto-nomous Vehicles	Startups	Research Institutes	SUM
Data Scientist	7	–	1	3	–	**11**
Researcher		–	–	–	3	**3**
Sales	3	–	–	–	–	**3**
Safety Engineer	2	–	1	–	–	**3**
Team Lead	1	–	–	2	–	**3**
Service Technician	2	–	–	–	–	**2**
Certification Engineer	–	2	–	–	–	**2**
Manager	3	–	3	2	–	**8**
Domain Expert	1	–	–	–	–	**1**
SUM	**19**	**2**	**5**	**7**	**3**	**36**

Figure 3: Overview of the interview participants and their roles

4. XAI insights across the AI life cycle

Explainability benefits all tasks

Analysis of the interviews brought forth several key insights. Above all, it was found that explainability is not only beneficial for the technical properties of AI models, but for the entire development process that leads to a specific model and to the tasks that follow. XAI has the potential to positively impact every stage of the AI life cycle, from the initial **business considerations** to the **data preparation, modeling, evaluation** as well as **deployment, monitoring** and **maintenance**.

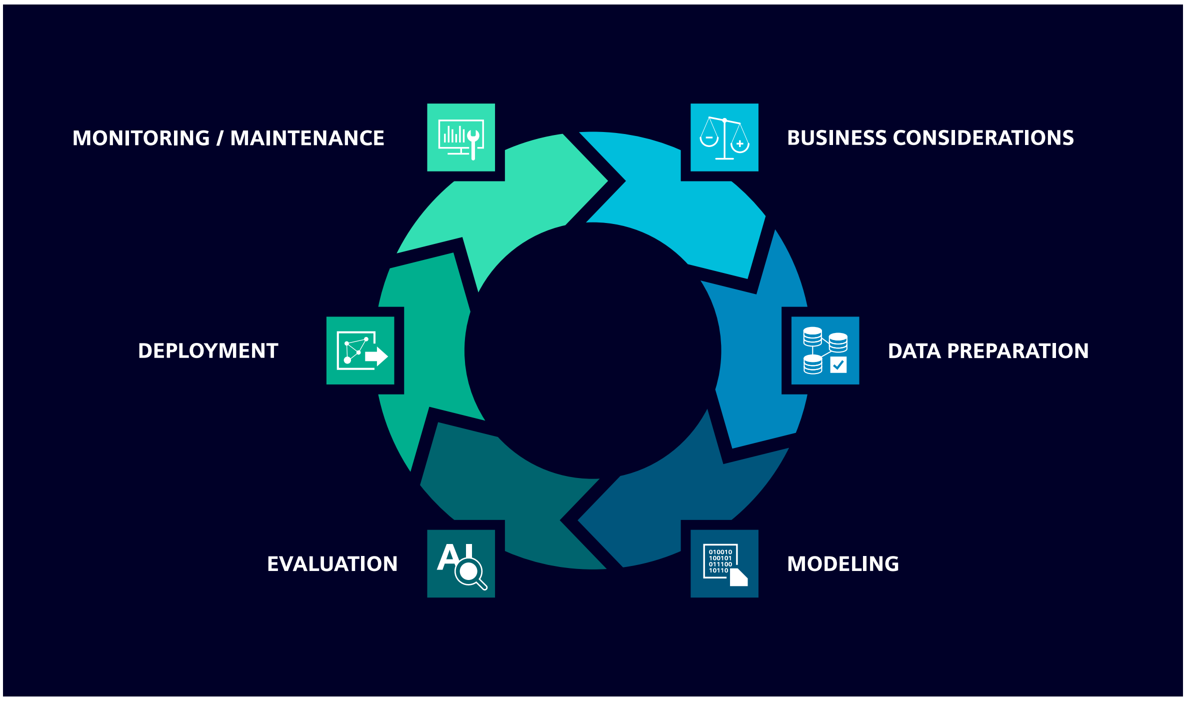

Figure 4: Representation of the AI life cycle

Business considerations

AI as a business enabler

As AI maturity continues to evolve, expectations of how to evaluate its success are gaining in importance. The days of simply riding the AI hype wave are over. Today's AI products, services and solutions strongly focus on financial measures. XAI is also quickly gaining relevance as a means for enabling business:

- To enhance differentiation in competitive environments
- To gain access in strongly regulated markets (e.g., food, pharma, aerospace)
- To address requirements of the upcoming EU AI Act
- To ensure meaningful interaction between humans and AI along the whole life cycle

Model building starts with building trust

Explainability is also vital to increase trust and Transparency in AI – not only in technology, but also in terms of business habits, customer relations and brand perception. "Fear of the unknown," for instance, was found to be present at all organizational levels, although for different reasons. End users and domain experts tend to fear job loss and replacement by AI. Management level fear of the unknown is often accompanied by anxiety regarding the reliability and robustness of the AI system.

Brand reliability helps to increase customer and user confidence in the technical reliability of AI systems. The need for change management and iterative approaches to help overcome a lack of trust when introducing AI systems was frequently highlighted. Honesty, in general, was found to be very important for sustaining trust in AI models, as were traceability, data transparency as well as information about how AI systems are built and what their limitations are.

Data preparation

Creating effective AI models is a collaborative effort that begins with data gathering and preparation.

Collecting the right data

Proper data gathering as needed for solving a particular problem during a typical AI project requires the combined expertise of many individuals and roles, in particular, that of data scientists and domain experts. The latter provide support with data selection and with defining the "correct" data to be collected. They also help AI developers to understand data quality evaluation. Both tasks are essential. Measurement and domain expertise must be combined to acquire data suitable for AI model building.

Ensuring data quality

Equally important, especially for achieving AI model objectives such as performance, trustworthiness, auditability, and maintainability, is the quality and quantity of available data. Data coverage and variety also play a vital role for developing robust AI functionality, safety arguments, and for proper model building in general. Defining the required data coverage is a task for entire teams: domain, machine learning and safety experts.

Although XAI typically only plays a major role after modeling, several experts stated that it helps with the iterated cycles between model building and data collection; for example, with determining influential data points that are then checked for data quality, data bias and potentially incorrect label issues.

Modeling

Developers of AI systems, as the interview findings show, have a strong desire to understand how models work at multiple levels, including the underlying math and algorithms, but also, more concretely, how models reach predictions that directly touch on the topic of explainability.

XAI is not yet put fully into practical use in manufacturing

Some developers are already using available XAI tools, for example, to link modeling and evaluation with computer vision methods. The tool stack of AI practitioners already includes standard techniques from the XAI literature. Many developers would welcome additional XAI tools, provided that they were convenient and offered additional benefits for their domain.

By contrast, a number of the data scientists that were interviewed still preferred to focus exclusively on performance metrics. They view XAI tools as being of limited relevance, raising doubts about their technical capabilities to deal with complex models and about general assumptions typically associated with XAI models. Their statements also provide an indication of the gap that still exists between research and practical application.

Collaboration and communication are key

AI model development for the real world is critically dependent on close collaboration between data scientists and domain experts. Domain-specific input is essential for selecting features, understanding model limitations and for proper model development in general. Integrating domain knowledge into AI models is a key responsibility of data scientists. The latter face tremendous challenges if they have no access to domain knowledge.

XAI helps to facilitate collaboration between data scientists and domain experts. However, in the case of domain experts, it is important to present XAI in their "language," using semantics familiar to them.

Evaluation

Testing AI models in industrial settings is challenging and frequently not adequately addressed through conventional machine learning techniques such as proper train-validation-test splits.

No access to real environments

For many use cases, AI developers lack access to data that accurately reflects realistic operating conditions. We strongly recommend testing AI models at the customer's site and under realistic circumstances. Extensive on-site testing also helps to increase trust.

Insufficient testing options

There is also a lack of established testing approaches. According to developers, there are three different ones. First, there is:

- Black box testing, which only deals with input-output relations

And second, two pre-specified testing approaches:

- One that defines edge cases and generates synthetic test data
- The other involves cross-functional teams, including domain experts, that cooperate to define relevant test cases/scenarios and carefully select datapoints

Particularly the last approach has shown to be effective in addressing general data quality and data insufficiency challenges. It is crucial that domain experts thereby assist data scientists in evaluating model decision correctness and in defining relevant test cases and scenarios. Data scientists reportedly spend significant time and resources on communicating with domain experts. XAI can help to bring about relevant communication efficiency improvement.

The question of responsibility and liability

Testing and validation are also crucial for the responsibility and liability that developers, product managers, product owners and even entire companies assume. The need for risk control, ownership and responsibility has many facets:

- End users of AI systems overweight mistakes, accidents and wrong decisions
- Authorities demand safety measures and certification
- Companies have product liabilities that also apply to AI components
- AI models generally decay over time, but someone must assume responsibility for ensuring long-term functionality

As the interview findings show, data scientists tend to be averse to assuming responsibility for their AI developments, citing struggles to accurately test models and difficulties in understanding the risks associated with AI. Company-wide recommendations included the introduction of processes to ensure AI traceability and defining complete chains of responsibility for AI products and solutions.

Regarding technical measures aimed at mitigating AI risks, participants proposed building “safety net processes” and programmed “safety fallbacks.” This is not the same as “safe AI” where the AI component itself is safe. The topic of proving safe AI was also discussed, given the fact that statistical measures such as accuracy are generally not accepted by authorities, prompting some companies to even stop planned market launches. In this case, AI testing difficulties become a barrier for innovation.

Rethinking AI safety

Addressing standards, regulations, and safety-related concerns with currently available XAI methods is a complex undertaking. Various factors need to be considered.

Future standards and audits

While there was general consensus that current standards were inadequate, opinions regarding the role that XAI should play in defining future regulations differed greatly.

Certification in safety-critical fields

AI systems certification for safety-critical use cases such as healthcare, public transport, industrial automation, and shopfloor control will likely be mandatory in the future. However, views on implementation varied greatly. Some called for the certification of the entire system, rather than just the subsystem. Others proposed that the development process itself be certified. It was noted that end-to-end AI systems certification may prove too difficult. Some argued that certification should be made contingent on the implementation of monitoring processes as part of the certification process.

How to do certification for AI is still unclear

With respect to the methodology, some suggested black-box testing, as the details are not relevant for certification. One participant felt that understanding the AI was not necessary for certifying units. Many, however, did recognize the value of XAI for certification, citing various reasons:

- XAI could help certification authorities understand AI models
- Explainability could become a mandatory feature to satisfy regulators
- It would require developers to understand model characteristics and decisions
- XAI would likely be a separate component of the certification process

Demand for up-to-date best practices to accurately evaluate AI safety is tremendous and will continue to grow. XAI can assume a key role in future safety methodologies – for example, in providing tools for identifying errors, safety verification and to facilitate safety audits. It can help strengthen AI safety arguments beyond qualitative reasoning, even as a requisite component for functional safety validation. This will require a selection of independent and redundant XAI tools used in combination with traditional approaches.

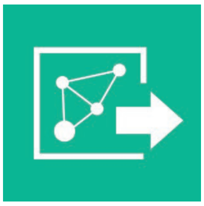

Deployment

AI models should be designed to integrate seamlessly with existing systems. There is also evidence that acceptance of new systems is increased when guidance for users, drill-down capabilities and root cause reasoning are provided.

The tricky task of scaling

Scaling XAI poses multiple challenges that are not only of a technical nature:

- XAI is computationally complex
- Scaling requires experts capable of interpreting XAI method outcomes
- XAI itself does not scale, because it needs to be domain and role specific

To make a specific example, most XAI techniques so far are developed with AI experts in mind. Hence all information from explanation systems comes in the semantics and language of AI experts and are, therefore, no real explanations for experts in other domains. Hence, further translation of the explanations is needed.

There is evidence, however, that XAI that are able to perform this translation tasks can support AI scaling by enabling non–AI experts to build, understand and use AI models.

Monitoring and Maintenance

Effectively monitoring and maintaining AI systems involves equally critical and also difficult tasks requiring various profiles and skill sets. In addition to highlighting helpful strategies and best practices currently in use, the interviews also underscored the potential XAI holds for facilitating the monitoring and maintenance of AI systems.

Training non-AI experts

In essence, monitoring and maintenance functions can be performed by non-AI experts with many different backgrounds – application engineers, automation engineers, service technicians, operators, even newer profiles such as DevOps engineers.

AI engineers were frequently mentioned as ideal candidates. After all, they designed the system. AI monitoring and maintenance involves many tasks that only they can perform. Sometimes junior data scientists are hired as a means of also providing them with the necessary foundation for their new roles. In other cases, the development team initially handles the monitoring and maintenance tasks before transferring the responsibilities to other personnel. What is important is that they be provided with adequate training, user-friendly interfaces, dashboards and other tools.

Potential strategies and best practices

Monitoring systems do not scale if human oversight is required. However, even simple alterations may need to be inspected manually. To minimize human intervention, it is essential that monitoring systems be self-sustaining, holistic and as automated as feasible.

Whereas some experts emphasized the need to monitor the functionality of a model, others felt that monitoring should focus on the data rather than the model. Still others noted that, with certain applications, neither the model nor the data should be monitored, but rather the business metrics.

The growing importance of XAI-based support

XAI is expected to play an increasingly important role in supporting operational and monitoring tasks of AI systems, especially for the purpose of enhancing AI maintainability, strengthening trust and mitigating safety concerns. One drawback: XAI is not able to identify causal relationships in the data, yet it provides tools for model-causal insights that allow users to identify the factors that lead AI models to certain decisions.

The following conditions and requirements were identified as essential for XAI systems used to support monitoring and maintenance:

- Proper handover and explanation of the model to person monitoring it
- Error/incident handling support for persons monitoring the model
- Support for ad hoc questions and requests
- Maintenance support

Also, it was noted that the interface for these XAI monitoring solutions must be intuitive and user-friendly and, ideally, allow persons monitoring the system to resolve issues independently.

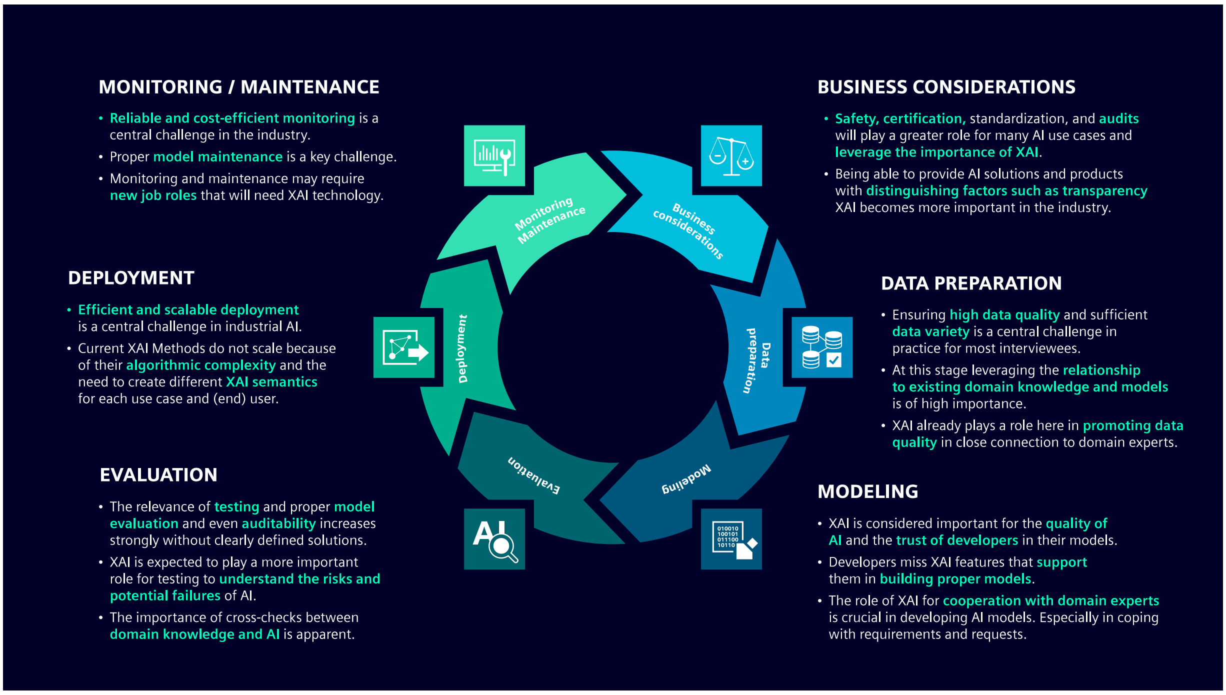

Figure 5: Representation of the AI life cycle with central findings from the interviews

5. Conclusion

Actionable XAI insights across the AI life cycle

Findings of the interviews showed that XAI can provide **valuable support across the entire life cycle of Industrial-grade AI systems** used in manufacturing. For one, XAI enables **easy integration of stakeholders** at various stages. It offers a variety of tools for **removing barriers between roles** and can be adapted for individual backgrounds and expertise. This is a key benefit. The development and operation of effective AI systems requires quality input from various stakeholders throughout the life cycle.

Through its ability to **improve AI project team communication efficiency,** XAI helps to streamline process stages such as data preparation and model evaluation. It can provide **assistance with certification,** for example, in **helping authorities better understand AI models**. XAI also shows tremendous promise for **facilitating machine learning operations (MLOps)** tasks such as commissioning, monitoring and maintenance.

Knowing that competition for highly trained Industrial AI personnel is fierce, **recruitment efforts for AI monitoring supported with XAI** should be widened to include non-AI experts from diverse backgrounds: application engineers, automation engineers, data scientists, service technicians and operators.

Interview findings confirm our assumption that **explainability not only has relevance for the technical properties of AI systems, but also for human-centric dimensions,** in particular, trustworthiness. To enhance confidence in AI models, it is essential that all stakeholders have access to the Industrial-grade AI development process. In addition to enabling easy adaptability to wide-ranging use cases, XAI seamlessly integrates different roles and stakeholders along with their individual needs:

Developers find XAI particularly useful for debugging, testing, and evaluating AI systems.

Domain experts are generally interested in the relevant components and the overall logic of AI systems.

Service technicians focus on AI (un) certainty, drill-down capabilities, and root causes of alarms / incidents.

Decision makers, managers and sales personnel, also customers and users with some degree of AI system accountability also need to be better informed about XAI.

While there is still uncertainty as to how XAI is best deployed and what level of detail users require, interview recommendations stressed the need to present **XAI in easily accessible forms.**
There was general consensus that users should be involved in the XAI system process, for example, with the help of user studies, tests, feedback and carefully designed UX interfaces. To be a genuine benefit, XAI must assume the design and semantics that users are familiar with.

XAI has the potential to greatly improve the productive use of Industrial-grade AI. Although features are relatively new and still evolving, we are already supporting customers with a variety of AI model development, logging, monitoring and debugging products and services. Our **AI Inference Server** provides a host of advanced industrial applications for standardizing AI models on the **Siemens Industrial Edge**. The AI Inference Server also enables AI model deployment as content for a standard Industrial Edge app. These are just the beginnings of our commitment to building strong AI capabilities for manufacturing together with partners.

In the spirit of building an ecosystem, to address the upcoming AI Act let us work together on how to improve industrial AI systems with explainability. Hence, this whitepaper is also a call for discussion and collaboration. Only together we can unlock the power of industrial-grade AI.

6. References

Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION LEGISLATIVE ACTS, COM(2021) 206 final, *https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206*

Simonyan, K., Vedaldi, A., & Zisserman, A. (2014, April 14–16). Deep inside convolutional networks: Visualising image classification models and saliency maps [Workshop poster]. International Conference on Learning Representations, Banff, Canada. *https://arxiv.org/pdf/1312.6034.pdf*

Wang, J., Ma, Y., Zhang, L., Gao, R. X., & Wu, D. (2018). Deep learning for smart manufacturing: Methods and applications. Journal of Manufacturing Systems, 48, 144–156.

Tercan, H., & Meisen, T. (2022). Machine learning and deep learning based predictive quality in manufacturing: Systematic review. Journal of Intelligent Manufacturing, 33, 1879–1905.

Nehrkorn, J., & Ochsenfeld, H. (2020). Reliable operations of machine learning systems in industrial environments [White paper]. Siemens. *https://new.siemens.com/de/de/produkte/automatisierung/themenfelder/artificial-intelligence-in-industry/whitepaper-machine-learnings-systems.html#Registration*

Galassi, A., Lippi, M., & Torroni, P. (2021). Attention in natural language processing. IEEE Transactions on Neural Networks and Learning Systems, 32(10), 4291–4308. *https://arxiv.org/ftp/arxiv/papers/1902/1902.02181.pdf*

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?” Explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. Association for Computing Machinery. *https://dl.acm.org/doi/proceedings/10.1145/2939672*

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), Advances in Neural Information Processing Systems 30 (NIPS 2017). *https://papers.nips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html*

Rojat, T., Puget, R., Filliat, D., Del Ser, J., Gelin, R., & Díaz-Rodríguez, N. (2021). Explainable artificial intelligence (XAI) on time series data: A survey. arXiv:2104.00950 [cs.LG]. *https://arxiv.org/pdf/2104.00950.pdf*

Molnar, C. (2020). Interpretable machine learning. Lulu.com. *https://christophm.github.io/interpretable-ml-book/*

Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv:1702.08608 [stat.ML]. *https://arxiv.org/abs/1702.08608*

Independent High-Level Expert Group on Artificial Intelligence. (2019). A definition of AI: Main capabilities and disciplines. European Commission. *https://www.aepd.es/sites/default/files/2019-12/ai-definition.pdf*

Schölkopf, B. (2022). Causality for machine learning. In H. Geffner, R. Dechter, & J. Halpern (Eds.), Probabilistic and Causal Inference: The Works of Judea Pearl (pp. 765–804). Association for Computing Machinery. *https://dl.acm.org/doi/10.1145/3501714.3501755*

Contact us to learn more:

Michael Lebacher
Head of AI & Digital Business Germany

Ralf Gross
Head of Industrial AI Portfolio

Stefan Hagen Weber
Senior Key Expert Visual Analytics

Geirhos, R., Jacobsen, J. H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., & Wichmann, F. A. (2020). Shortcut learning in deep neural networks. Nature Machine Intelligence, 2(11), 665–673.

Nagarajan, V., Neyshabur, B., & Andreassen, A. (2021).
Understanding the failure modes of out-of-distribution generalization. arXiv:2010.15775v2 [cs.LG].
https://arxiv.org/pdf/2010.15775.pdf

Ntoutsi, E., Fafalios, P., Gadiraju, U., Iosifidis, V., Nejdl, W., Vidal, M.-E., Ruggieri, S., Turini, F., Papadopoulos, S., Krasanakis, E., Kompatsiaris, I., Kinder-Kurlanda, K., Wagner, C., Karimi, F., Fernandez, M., Alani, H., Berendt, B., Kruegel, T., Heinze, C., ... Staab, S. (2020).
Bias in data-driven artificial intelligence systems – An introductory survey.

WIREs Data Mining and Knowledge Discovery, 10:e1356.
https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/widm.1356

Jacovi, A., Marasovic, A., Miller, T., & Goldberg, Y. (2021). Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in AI. In FAccT 21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (pp. 624–635). Association for Computing Machinery.
https://dl.acm.org/doi/10.1145/3442188.3445923

Bhatt, U., Xiang, A., Sharma, S., Weller, A., Taly, A., Jia, Y., Ghosh, J., Puri, R., Moura, J. M. F., & Eckersley, P. (2020). Explainable machine learning in deployment. arXiv:1909.06342 [cs.LG].
https://arxiv.org/abs/1909.06342

Dillon, E., LaRiviere, J., Lundberg, S., Roth, J., & Syrgkanis, V. (2018). Be careful when interpreting predictive models in search of causal insights. SHAP documentation, revision 45b85c18.
https://shap.readthedocs.io/en/latest/example_notebooks/overviews/Be%20careful%20when%20interpreting%20predictive%20models%20in%20search%20of%20causal%C2%A0insights.html

Chen, H., Janizek, J. D., Lundberg, S., & Lee, S.-I. (2020). True to the model or true to the data? arXiv:2006.16234v1 [cs.LG].
https://arxiv.org/pdf/2006.16234.pdf

Selvaraju, R., Cogswell, M., Das, A., Vedantam, R., Parikh, D. & Batra, D. (2017). „Grad-cam: Visual explanations from deep networks via gradient-based localization." Proceedings of the IEEE international conference on computer vision.
https://openaccess.thecvf.com/content_ICCV_2017/papers/Selvaraju_Grad-CAM_Visual_Explanations_ICCV_2017_paper.pdf

Studer, S., Bui, T.B., Drescher, C., Hanuschkin, A., Winkler, L., Peters, S., Müller, K.R.: Towards crisp-ml (q): a machine learning process model with quality assurance methodology. Machine Learning and Knowledge Extraction 3(2),392–413 (2021)

Further readings

Topic-related Siemens white papers:

- How to combine AI and industrial automation: Artificial intelligence on the shopfloor: Envisioning the future of intelligent automation and its impact on manufacturing: 2018
  https://www.siemens.com/global/en/products/automation/topic-areas/tia/future-topics/whitepaper-shopfloor-ai.html
- The challenges of MLOps in industrial domains: Reliable operations of machine learning systems in industrial environments: 2020
  https://www.siemens.com/de/de/produkte/automatisierung/themenfelder/artificial-intelligence-in-industry/whitepaper-machine-learnings-systems.html#Registration

Acknowledgments

We wish to thank our internal and external interview partners for their open and honest contributions. We would also like to extend special thanks and recognition to the three publicly funded projects from which some of our interview participants were sourced:

- Cognitive Intelligence Platform for Economic Ecosystem Resilience – CoyPu: ***coypu.org/en/***
- Rethinking mobility – Safe AI using the example of a driverless regional train: ***safetrain-projekt.de/en/***
- Data-based manufacturing optimization of battery cells based on end-of-line data through massive use of machine learning algorithms and inline analytics – DAFODIL. („Datenbasierte Fertigungsoptimierung von Batteriezellen auf Basis von End-of-Line-Daten durch massiven Einsatz von Machine-Learning- Algorithmen und inline-Analytik – DAFODIL"):
  https://db.batterieforum-deutschland.de/verbundprojekte/dafodil/