Lufthansa
Industry Solutions

Artificial Intelligence

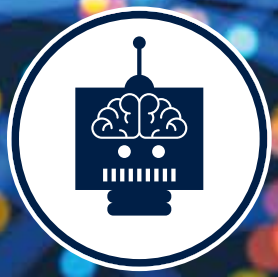

Definition of artificial intelligence

"Artificial intelligence (AI) concerns methods that enable a computer to perform tasks which, when performed by a human, require intelligence." (Gabler Wirtschaftslexikon, 2018 [Dictionary of Commercial and Business Terms])

Technological
Artificial intelligence denotes the use of parametrized models that can process new types of data or create new data in order to perform tasks. Selecting appropriate parameters creates an AI that can perform a task, including using new, previously unseen data.

Data analytics or artificial intelligence?

When does data analytics become artificial intelligence?
It is often not possible to draw a clear line between data analytics and AI: both disciplines are closely intermeshed and, in many cases, even build on one another. Both involve processing large volumes of data in an efficient manner. The insights gained through data analysis can help to create an AI.

Data analytics
The aim of data analytics is to glean insights from large volumes of data with the assistance of mathematical methods and, in so doing, to gain a better understanding of the process that produced this data. Data analytics combines in-depth knowledge of the respective field (known as domain knowledge) with an analytical approach specifically tailored to each data set. The insights gained are then evaluated by human experts.

Machine learning
The term machine learning refers to procedures or algorithms that automatically analyze context in historical data and universalize this to determine the parameters of an AI. Machine learning is a branch of AI. However, there are also other means by which to determine the parameters of an AI.

Artificial intelligence
An artificial intelligence gleans insights or creates new information on the basis of a parametrized model. This model and its parameters are designed in such a way that the AI can apply the insights it gathers in entirely new situations in a consistent manner. It reaches its result autonomously – that is to say, without human interaction.

What should a good artificial intelligence be able to do?

Creating an AI involves

- addressing and economically optimizing business cases,
- minimizing bottlenecks in existing business processes,
- taking scalability and platform independence into consideration, and
- calculating operating costs for various models (i. e. on-premises, cloud-based, edge devices).

A good AI

- provides people with seamless support and performs monotonous tasks,
- independently identifies and universalizes patterns in training data so it can also handle new data streams correctly,
- is robust and able to cope with data stream variations,
- can react appropriately to changes in the data stream and adjust results accordingly, and
- reaches reasonable and justifiable decisions.

Limits of artificial intelligence

Data focus
If the data source is flawed or does not contain all required data, an AI can only provide approximated answers. Conducting a differentiated analysis of the information content of the data used is imperative, as it makes it possible to evaluate the quality of the AI's results in advance.

Resources and costs
Due to the wide availability of computational power and the use of special hardware, an AI can also provide answers to highly complex questions. Nevertheless, the available computing and storage capacities and the time available for processing remain limiting factors. Yet, technological improvements and more sophisticated algorithms continue to push the boundaries of what is possible while simultaneously reducing costs.

Systematic restrictions
An AI is usually deployed to answer specific questions or problems in one or more specialist areas. An AI therefore only has an indirect regard for other domains. This means that AIs often produce an oversimplified image of the world because they lack general knowledge. As a result, an AI can in some situations be inferior to a human who can independently combine knowledge from different sources to draw new conclusions.

Ethical considerations
As AI primarily serves to support humans through (partially) automated decisions, it is also crucial to reflect on ethical issues. If the training data contains or suggests ethically dubious or undesirable decisions, the AI is also likely to reproduce such decisions. Potential ethical consequences therefore need to be considered and evaluated before creation and use of an AI.

How complex is an artificial intelligence?

Contrary to common misconception, AIs need not always be highly complex and sophisticated. A small-scale, easy-to-use AI can be created to handle clearly delineated tasks and issues and provide valuable insights. Analyzing business processes in detail makes it possible to determine the degree of AI complexity required for the desired purpose.

What degrees of complexity need to be considered?

Every artificial intelligence should be evaluated from various points of view. However, the mutually dependent nature of all degrees of complexity often necessitates compromise.

- **Complexity of problem domains:** The AI selected must be suitable for the problem or task at hand. As a general rule, the more complex the domains, the more parameters an AI requires. This has implications for the following degrees of complexity.
- **Complexity of learning:** All AIs require some degree of computation time and storage during its learning phase, dependent upon how many parameters are to be set. There is often no linear correlation between the two, and the time and outlay required can quickly and unexpectedly exceed the available resources.
- **Complexity of execution:** Another key aspect is the time and storage usage required for the AI to produce its results – especially in the case of real-time requirements. This complexity measure is particularly important in the context of live video processing. An AI can often produce results quickly even if it required a long learning period.

How does an artificial intelligence learn?

- AIs often learn laws and regularities and identify patterns from historical data. This data should ideally be representative; the raw data should contain all information required to produce the required results. To simplify an AI's learning process, easily computable mathematical features of data points are used as a substitute for raw data. While these features are currently determined by data scientists and experts in most cases, new approaches now allow AIs to calculate them.
- However, it is also possible to create an AI without historical data: by applying their specialist knowledge, experts can develop and configure an AI directly.

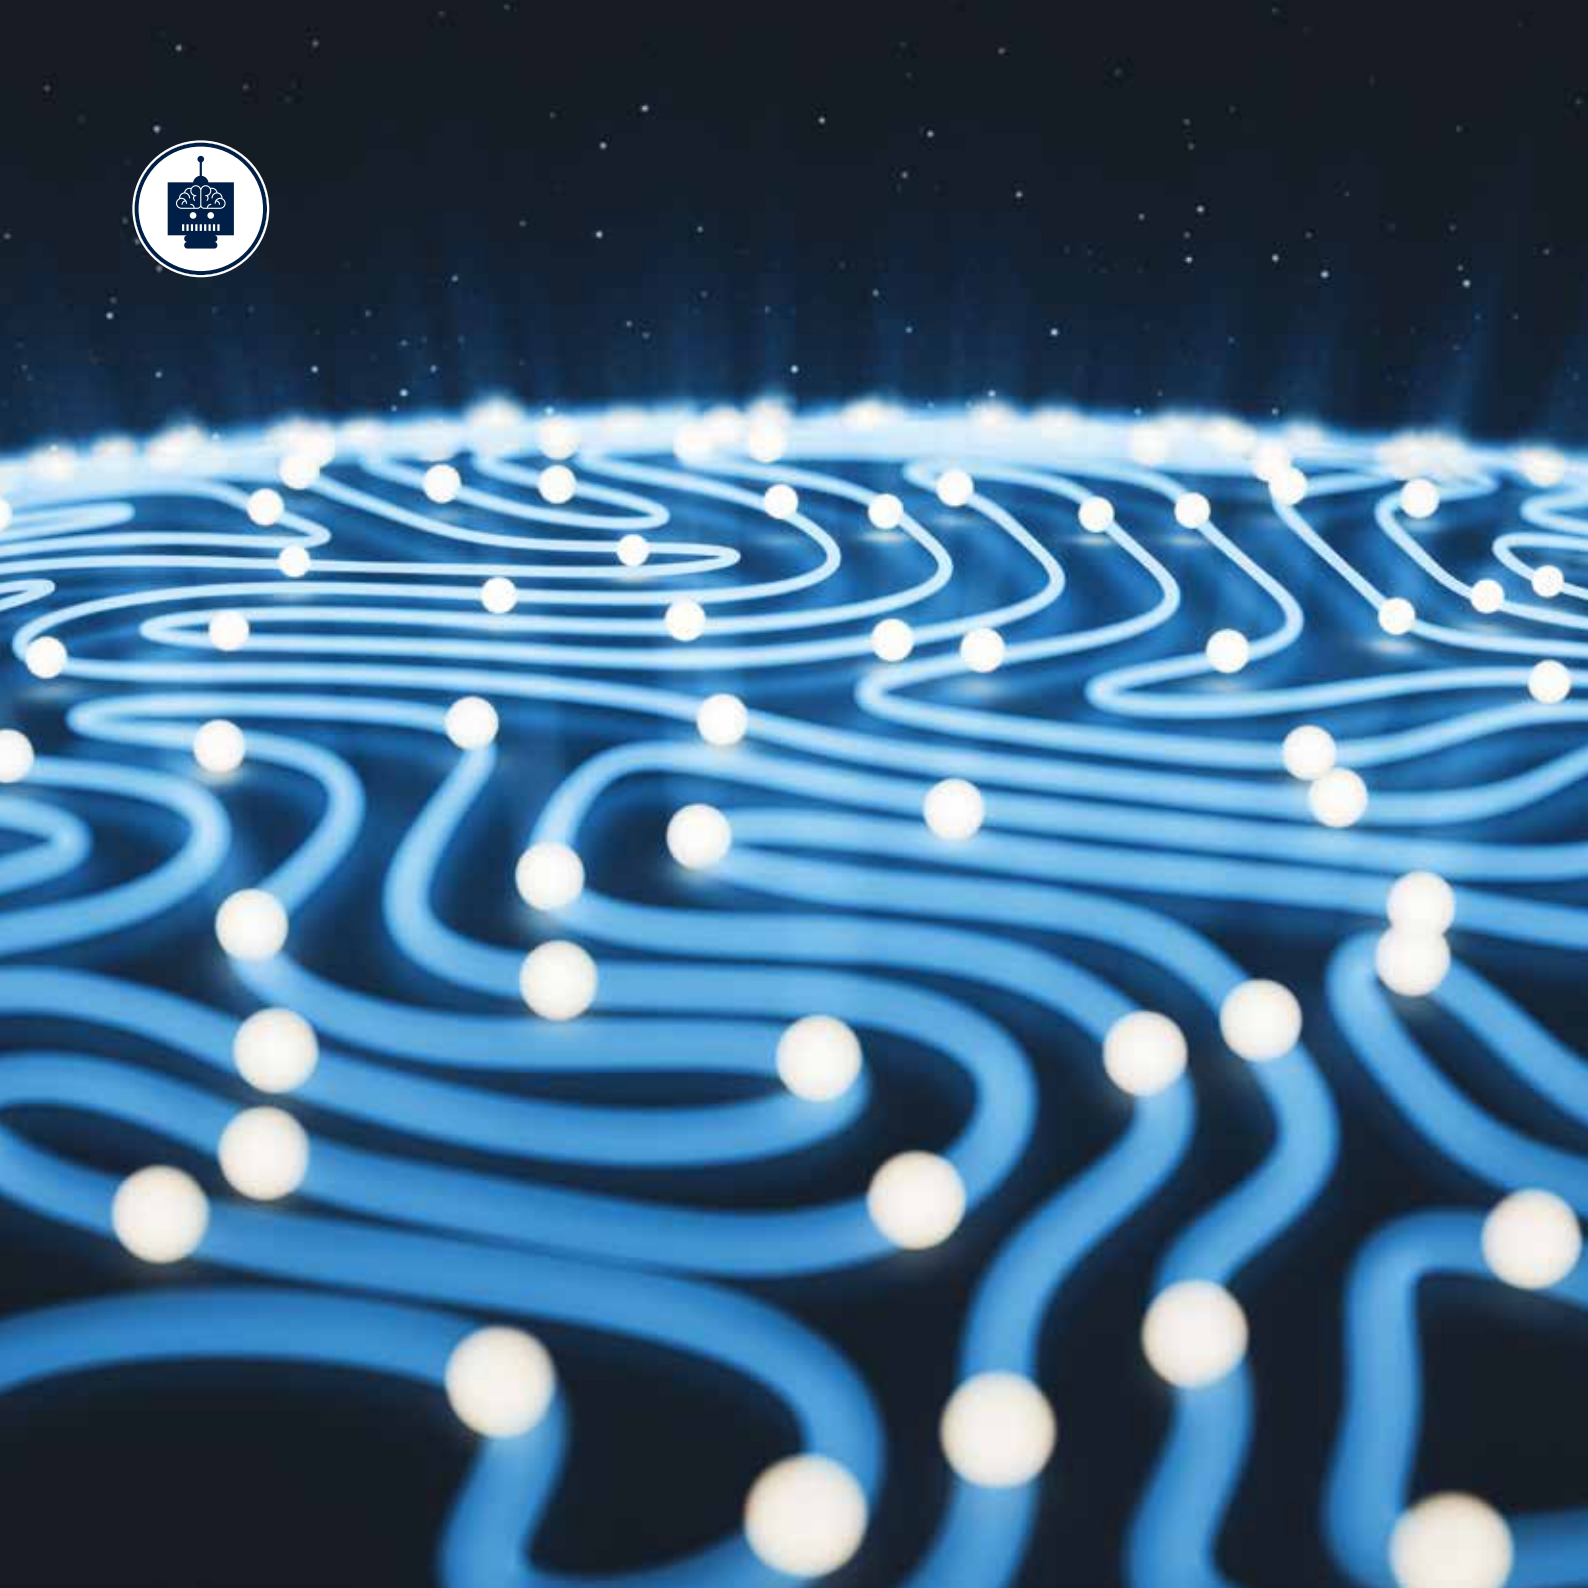

Which learning approach does an artificial intelligence use?

In terms of an AI learning from data, a distinction is drawn between three fundamentally different approaches:

Supervised learning
As well as the input data, the learning dataset also contains information on the expected outcome. A learning algorithm can thereby compute what errors an AI would make when processing this data and thereby adjust the parameters to minimize these errors.

Unsupervised learning
If no information or only limited details of the expected outcome are available, learning often takes place on the basis of assumptions or requirements. For example, one might assume that similar data will form a group. New data points can then be allocated to one of these groups on the basis of their similarities. Clustering procedures are based on this assumption. If the general assumption is that the data stream will display regularities, an AI can then recognize anomalies – without having previously observed them in the training data.

Reinforcement learning
If an AI is to make decisions that can be evaluated during learning – in a separate training phase or even in live operation – the reinforcement learning approach is often applied. In this case, the AI learns entire decision-making strategies to interact with its environment. This type of AI can be training either in live operation or through simulations. For instance, it can learn to play board games or control robots. At this point, a familiar problem becomes particularly clear: balancing further exploration of the environment to improve the decision-making strategies learned to date against using what has been learned to date. This conundrum is often referred to as “exploration vs. exploitation”.

Supervised learning

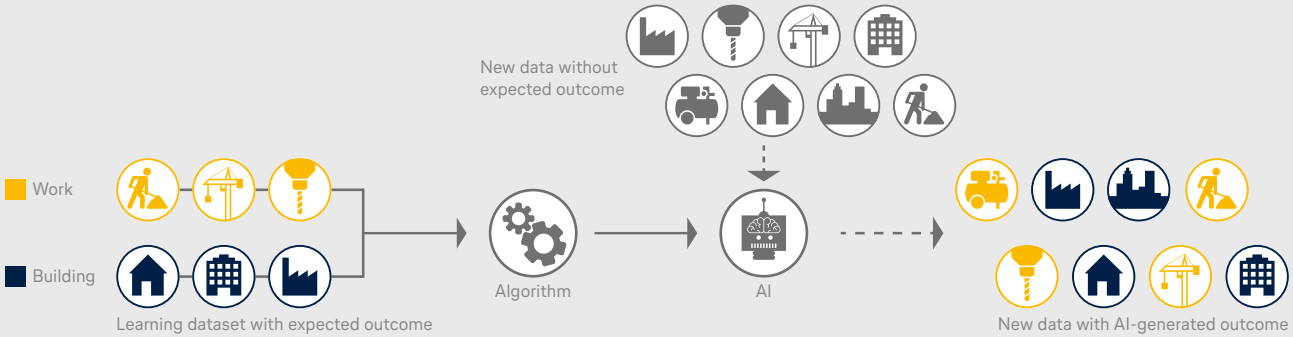

Unsupervised learning

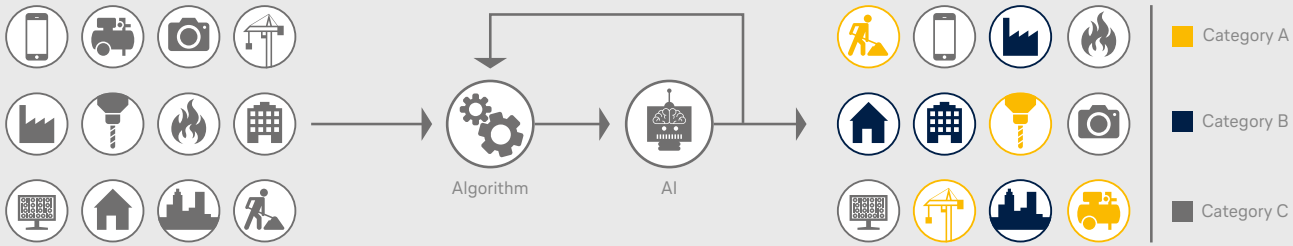

Reinforcement learning

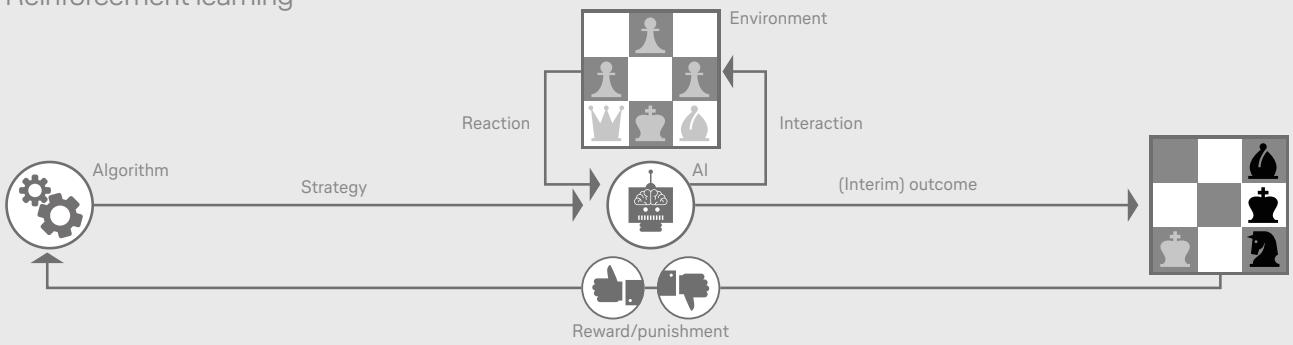

Can an artificial intelligence learn from small data volumes?

Yes, provided that:

- the data volume required for training is reduced through direct modeling of expert knowledge.
- experts determine the relevant features of the data points for the respective problem and reduce the search space by removing irrelevant features.
- learning achievements of an AI from a similar or related problem area are carried over. The existing data can adjust the pretrained AI to the new domain (known as transfer learning).
- additional synthetic data is used to supplement the original data volume. This could involve using entirely synthetic data or, alternatively, modifying existing data (making it “noisy”) or combining it in a manner that preserves its original nature.
- simulations are used to generate synthetic data and evaluate AI decisions through reinforcement learning.
- the problem or issue is simplified: this reduces the number of dependencies and thereby also the scope of data required.

Common mistakes

- Flawed model assumptions may train the AI to make connections that do not reflect real-world causality.
- **Solution:** AI experts collaborate with domain specialists and examine the models and parameters to ensure they are syntactically and semantically correct.

- Omnipresent computing power can create an excessively complex AI. Such an AI tends to memorize data more quickly rather than generalizing it.
- **Solution:** The connotations of model assumptions and hyperparameters must be understood to facilitate a sensible, appropriate selection.

- Raw data is treated in a domain-agnostic manner.
- **Solution:** Domain specialists and AI experts prepare existing data with reference to semantic and model-specific criteria.

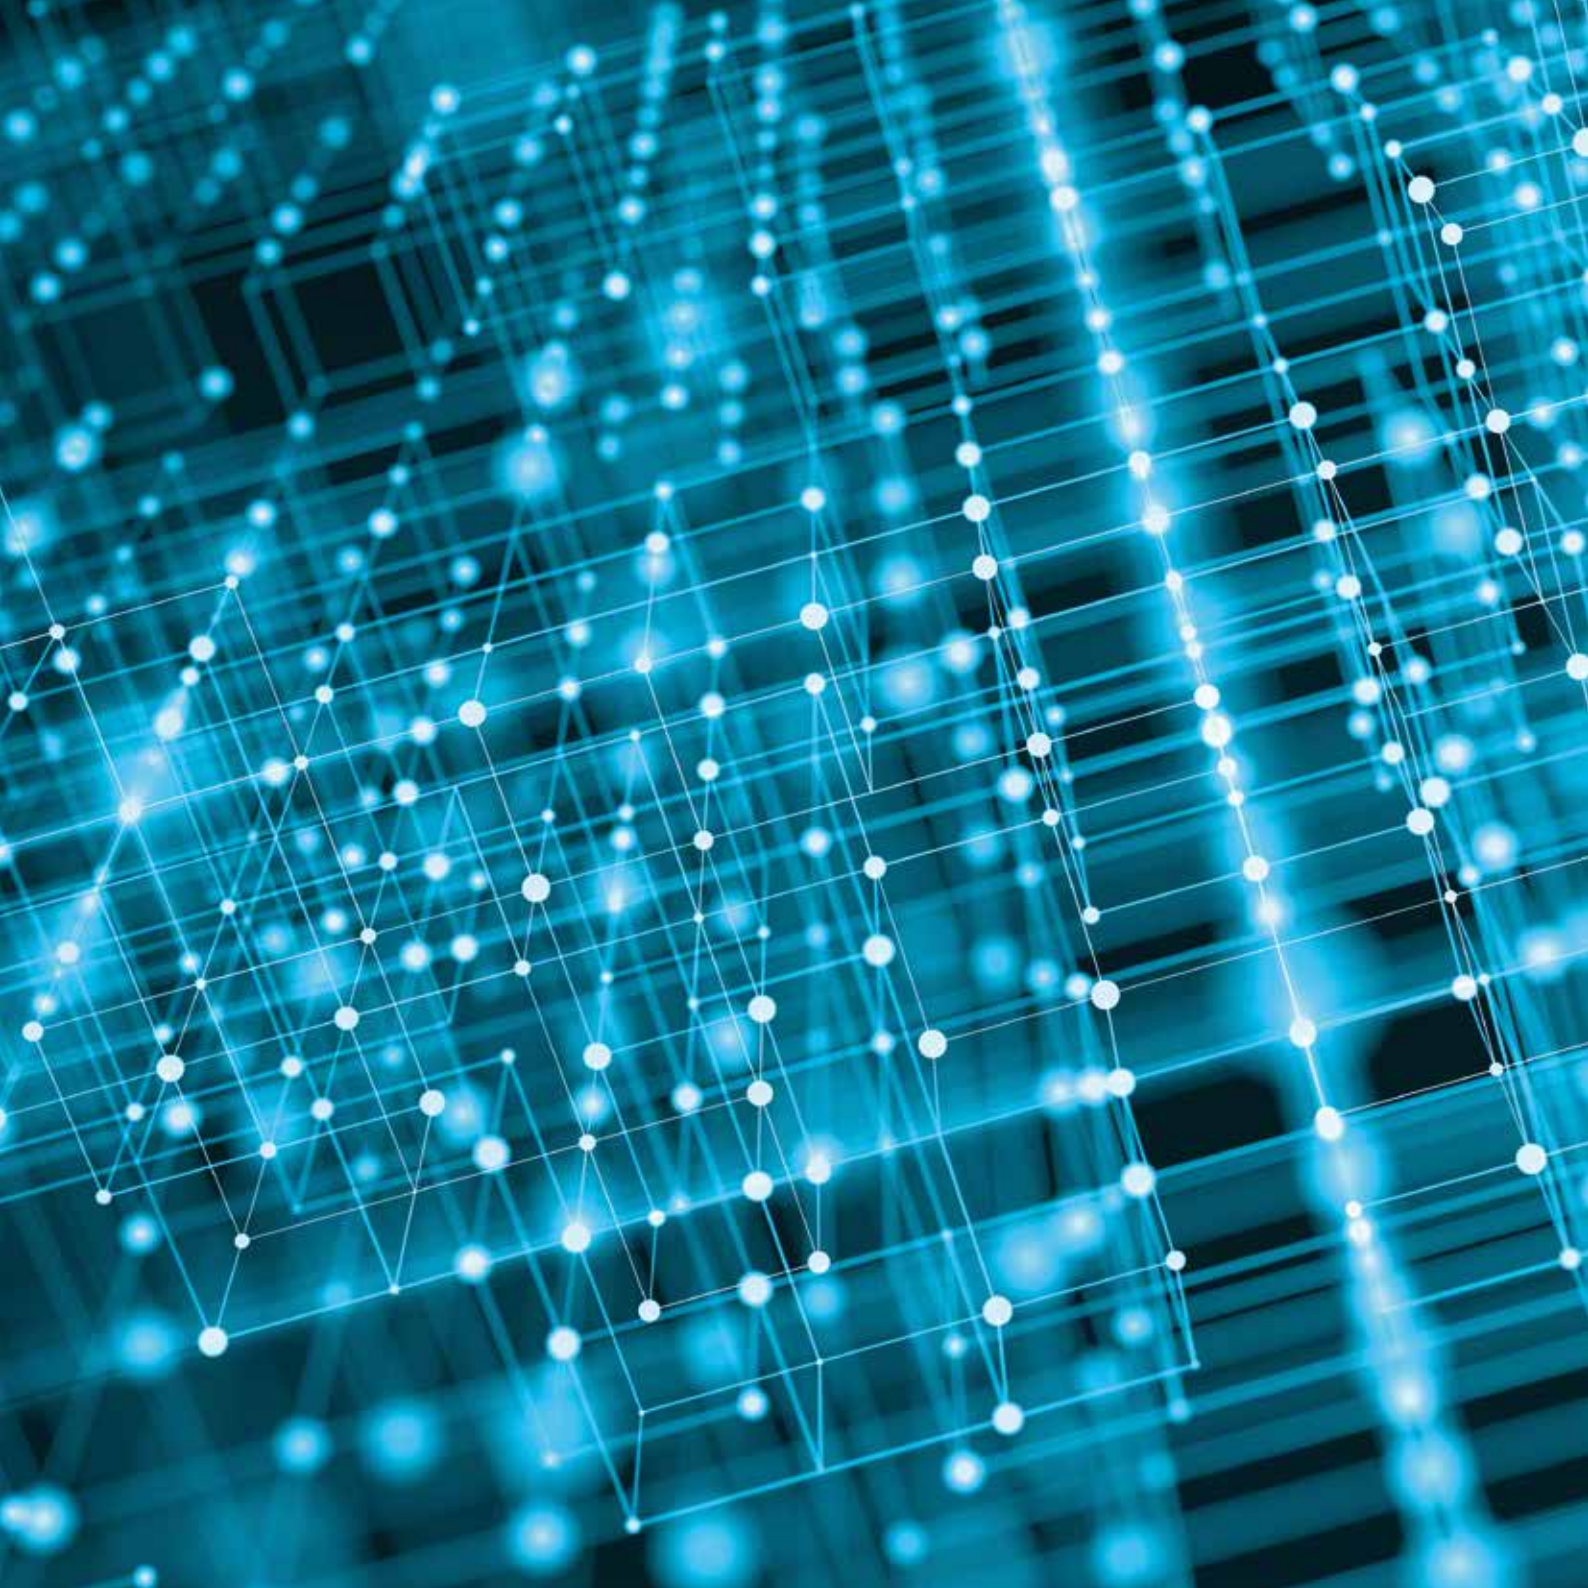

Artificial intelligence at LHIND

We implement AI to support people and their work processes.

- The starting point for this support is always your specific situation and existing digital strategy, which we work together to enhance in order to achieve your vision.
- Our focus on business cases enables us to implement AI at suitable junctures and avoid technological overcomplication and gimmickry.
- Tailored AI solutions integrate seamlessly in existing infrastructure and rely on cost-efficient solutions, either on-premises or in the cloud.

We work with a sustainable, future-oriented focus.

- We evaluate existing processes and identify bottlenecks that an AI could resolve. This makes it possible to support people in recurrent tasks.
- Our solutions combine continual updates, high quality AI forecasting or decision-making, and professional operationalization.
- We develop strategies to monitor operationalized AI at all times, as well as strategies to improve AIs, adapt them to new framework conditions, and ensure all decisions can be fully audited.
- We use established components and supplement these with in-house developments where necessary, taking the latest research developments into account.
- Thanks to the scientific backgrounds of our AI experts, we understand the models and algorithms we use and select them where sensible and appropriate.
- We play an active role in scientific discourse and translate the results of research into innovative solutions for industry.
- You benefit from our interdisciplinary expertise using cutting-edge technologies from our partners, such as AWS, Azure and Google. In doing so, we take your preferences into account and also incorporate project-specific and legal requirements.

Our strengths

Computer vision

- This image processing science takes images as a starting point and uses them to identify objects, measure sizes, evaluate user inputs and even draw conclusions of meanings in a given context.
- **Examples:** Detecting damage to vehicles; performing QA checks on materials and production goods; sorting general cargo; identifying objects in chaotic or confusing situations, and evaluating user-generated data

Stream analytics

- Continuous data streams from sensors (e. g. audio data, vibration levels, temperatures, etc.) are analyzed to identify patterns or deviations. Incoming signals are converted according to the specific application's requirements in order to facilitate efficient processing.
- **Examples:** Identifying anomalies in case of malfunctions; filtering background noise, and intelligent smart home applications

Natural language processing

- Documents, emails, comments, and reviews in text form can be understood in a given context and classified accordingly to enable an appropriate response. Beyond the explicit information in the text itself, this process takes into account the implicit context and intention.
- **Examples:** Performing sentiment analysis of social media posts; classifying user comments; analyzing the similarities between different documents; paraphrasing documents

Industry expertise and domain knowledge

Our AI experts work hand-in-hand with domain specialists from our industry-specific departments and continuously discuss and exchange information and ideas in order to create innovative solutions. This allows us to create synergies between solutions from different sectors.

Lufthansa Industry Solutions
Marketing & Communications
Schützenwall 1
22844 Norderstedt
Germany
T: +49 40 5070 3000 0
F: +49 40 5070 7880
E: marketing.sales@lhind.dlh.de
LHIND.de