# Chapter 13
# Data Ecosystems: A New Dimension of Value Creation Using AI and Machine Learning

**Dirk Hecker, Angelika Voss, and Stefan Wrobel**

**Abstract** Machine learning and artificial intelligence have become crucial factors for the competitiveness of individual companies and entire economies. Yet their successful deployment requires access to a large volume of training data often not even available to the largest corporations. The rise of trustworthy federated digital ecosystems will significantly improve data availability for all participants and thus will allow a quantum leap for the widespread adoption of artificial intelligence at all scales of companies and in all sectors of the economy. In this chapter, we will explain how AI systems are built with data science and machine learning principles and describe how this leads to AI platforms. We will detail the principles of distributed learning which represents a perfect match with the principles of distributed data ecosystems and discuss how trust, as a central value proposition of modern ecosystems, carries over to creating trustworthy AI systems.

D. Hecker (✉)
Fraunhofer-Institut für Intelligente Analyse- und Informationssysteme IAIS, Schloss Birlinghoven, Sankt Augustin, Germany

ML2R Kompetenzzentrum Maschinelles Lernen Rhein-Ruhr, Schloss Birlinghoven, Sankt Augustin, Germany
e-mail: dirk.hecker@iais.fraunhofer.de

A. Voss
Fraunhofer-Institut für Intelligente Analyse- und Informationssysteme IAIS, Schloss Birlinghoven, Sankt Augustin, Germany

Kompetenzplattform KI.NRW, Schloss Birlinghoven, Sankt Augustin, Germany
e-mail: angelika.voss@iais.fraunhofer.de

S. Wrobel
Fraunhofer-Institut für Intelligente Analyse- und Informationssysteme IAIS, Schloss Birlinghoven, Sankt Augustin, Germany

ML2R Kompetenzzentrum Maschinelles Lernen Rhein-Ruhr, Schloss Birlinghoven, Sankt Augustin, Germany

Institut für Informatik III, Rheinische-Friedrich-Wilhelms-Universität Bonn, Bonn, Germany
e-mail: stefan.wrobel@iais.fraunhofer.de

## 13.1 Introduction

In recent years, the effective use of machine learning and artificial intelligence (AI) has become a decisive factor for the competitiveness of individual companies and entire economies. Due to fundamental advances in available algorithms and computing power, given enough data, tasks that previously would have required human intelligence can now be automated, dramatically increasing efficiency and enabling completely new products and services. The availability of data, however, has proven to be a major roadblock for all but the largest and most digitally affine corporations. Typically, AI applications have used within-company data resources only, severely limiting the quality and feasibility of AI deployment in many cases.

The availability of a digitally sovereign, technically secure, and economically viable architecture for shared data spaces and digital ecosystems across different partners therefore is a fundamental game changer for the use of machine learning and artificial intelligence almost everywhere. Using a data space along the value chain or with complementary partners, new data sources can be brought into the construction of machine learning applications in a well-organized and, most importantly of all, self-determined fashion. Since availability and accessibility of data are always governed by the participants in the ecosystem, companies remain in control, thus enabling significantly more data to be made available to others.

Secondly, due to the semantic interoperability and unified data modeling of the data space, a sensible integration of different data sources, which is especially crucial for machine learning, can be performed with orders of magnitude less effort than in traditional project-wise approaches. With more and more complementary data sources becoming available, the quality and scope of machine learning results increase dramatically.

Furthermore, due to the inherently distributed nature of a data space, companies can now leverage the full potential of novel, distributed machine learning approaches. With these approaches it is no longer even necessary to combine all data at a central place; instead, they can be processed locally, at the edge of the data space nodes. This not only guarantees confidentiality of data but also affords considerable savings and potential for scalability.

Embedding machine learning and AI applications into data ecosystems using data space architectures thus not only makes AI applications accessible to organizations outside of the classical data-intensive digital sector, but also brings about significant benefits for those who have already deployed AI by allowing them to naturally move towards a distributed perspective on AI systems [1].

In the following sections, we will first (Sect. 13.2) give a brief overview of how artificial intelligence and machine learning applications are built today in the development cycle that centers on the use of the right data combined in the right way. We will then widen our perspective in Sect. 13.3, moving away from the technical development cycle and focusing instead on the platforms that are used to put these applications into deployment. In Sect. 13.4, we will then take a deep dive into one particular machine learning technology, distributed machine learning, that is

especially suited for AI systems and federated ecosystems and promises to deliver all the advantages of powerful machine learning modeling without centralized data storage. From thereon, in Sect. 13.5, we will generalize and discuss a more general architecture for machine learning in (distributed) digital ecosystems. Since trust is a core value proposition of these ecosystems, in Sect. 13.6 we will conclude the chapter with a brief discussion of the aspects of trustworthiness and AI applications.

## 13.2   Big Data, Machine Learning, and Artificial Intelligence

Big data became a hot topic in Europe around 2013 and was soon followed by data science, machine learning, deep learning, and artificial intelligences as trending topics. *Artificial intelligence* (AI) is a subfield of computer science, and machine learning and deep learning are now its most successful areas. M*achine learning* produces knowledge in the form of statistical models. There are different types of models, each equipped with a learning algorithm that optimizes the model automatically from training data. The most popular *models* are decision trees for classification tasks, regression curves for quantitative prediction tasks, clusters of similar data for pattern recognition, and artificial neural networks from deep learning [2, 3].

The recent success of machine learning is due to the volumes of data, and especially unstructured data, that can be stored in big data architectures: text, speech and audio, images and video, and streams of sensor data. *Deep learning* takes such data and trains artificial neural networks. They can make predictions and give recommendations. They can develop strategies for actions in games or to control a robot. They can generate images, texts, and language [4].

*Data science* combines methods from mathematics, statistics, and computer science for the discovery of knowledge in data. *Knowledge discovery* is also called *data mining*. Already in 1999 a process model, called CRISP-DM, was published as a step-by-step data mining guide. CRISP-DM stands for "Cross-Industry Standard Process for Data-Mining" and became an industry standard. The six phases of the process, shown in Fig. 13.1, are interlaced because data mining is an explorative process. Problem statements and the hypotheses initially formulated during *business and data understanding* may have to be adjusted when it turns out that the available data do not yield good enough models. Therefore it is important to identify many sources of high-quality data. Apart from data catalogs, data understanding can be supported by tools for descriptive statistics and visual analytics to find outliers, gaps, missing data, or underrepresented cases.

*Data preparation* involves data cleaning and transformations: generation of missing or derived data and removal of outliers, possibly also annotations and semantic enrichments. The latter transform the data into a standard vocabulary and guarantee seamless and ambiguous exchange between companies [5]. Data

**Fig. 13.1** The CRISP-DM process

preparation is still the most time-consuming phase and may take up to 70 % of the entire process [6].

Modeling and evaluation are supported by dedicated machine learning tools. For training, many *learning algorithms* require data which is labelled with the correct results or with other semantic annotations. This often drives the cost of data preparation. Machine learning is a statistical approach, and its models only return approximate results. Therefore *evaluating* a model is mandatory. It is done by running the model on extra data, which has not been used for training, and comparing predicted and correct results. Data which is not representative and contain errors or prejudices lead to deficient models.

Not for all learning tasks there is enough data to train good machine learning models. Therefore, researchers are now investigating algorithms that can learn causal relations and supplementary models that can represent facts and symbolic knowledge [7, 8].

A trained and evaluated model is not yet an AI application or intelligent solution. Usually, the model is placed into a workflow and linked to components for data fetching, data preprocessing, and model invocation, and the workflow is embedded in coded components that act upon the model's results. For instance, in an email

filter, a model may classify an email as junk or no-junk, but ordinary code must be written to move the email into the corresponding folder. A trained model embedded in traditional code constitutes the *deployable* application or solution.

## 13.3   An Open Platform for Developing AI Applications

In the preceding section, we have been focusing on the abstract data science process. When developing and deploying applications, this process is typically carried out with the help of software toolkits with the view towards deployment platforms which we will have a look at in the following paragraphs.

In the marketplace, there are several toolkits and platforms for machine learning and data mining, both open source and free to use as well as commercial. They provide the discussed means to analyze and preprocess data, to generate and test models, and to wrap them with workflows for deployment in smart applications. But to industrialize machine learning, such platforms need to be combined with data sharing platforms like the IDS on the one hand and digital business platforms with AI services on the other hand. All big ICT enterprises provide such integrated environments. But each of them supports a particular language and a set of compatible libraries, such as SciKit Learn, TensorFlow, H2O, and RCloud. This imposes specific standards and interfaces which lock the users into the provider's particular ecosystem.

In this section, we want to focus on ACUMOS AI [9], a platform that avoids the lock-in effect. It is an open-source development by the Linux foundation, was adopted by the European flagship project AI4EU [10], and thus will set a standard in the way AI applications are governed in Europe. The key idea of the AI4EU ACUMOS platform is that many models are reusable. Models pre-trained with big data sets often can be transferred to specific tasks by post-training with additional, task-specific data. In artificial neural networks, for instance, this can be achieved by retraining an emptied last layer, which is responsible for the final results.

To maximize the reusability of machine learning models, ACUMOS AI separates the work of the machine learning specialists from that of the application developers as shown in Fig. 13.2.

Machine learning specialists who have explored, trained, and evaluated a model in their preferred machine learning toolkit can *onboard* it to ACUMOS AI (1 in Fig. 13.2). That means, the model is packaged into combinable micro-services and described in a catalog. The catalog also contains components for data access, data transformation, and complementary software. ACUMOS provides a design studio, where application developers can graphically connect components from the catalog without any coding and without knowledge of the components' interna. Thus, they can build training workflows to retrain models into so-called predictors and put them into application workflows (2). Workflows can also be *published* in the catalog where others can rate them or give more specific feedback (3). Application workflows, packaged into a docker image, can be *deployed* in an execution
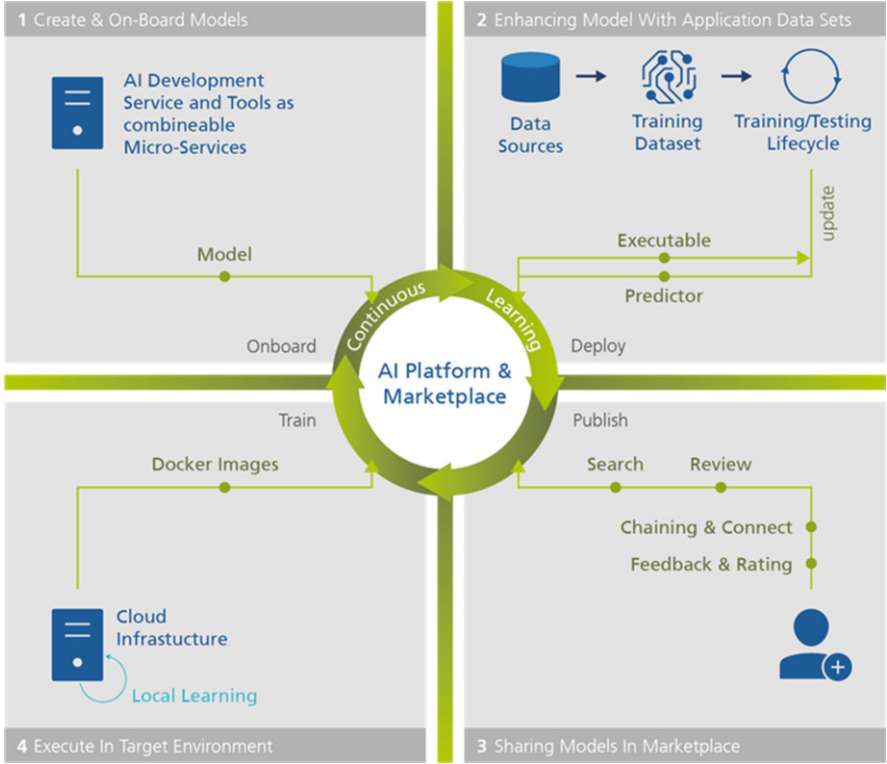
**Fig. 13.2** Application building in ACUMOS AI

environment such as Azure, AWS, other popular cloud services, or any corporate data center or any real-time environment (4).

Within AI4EU, ACUMOS AI was fed with various AI models. Application developers can find them in the catalog and combine them into hybrid models [11]. ACUMOS AI is a federated platform, which means that users can access catalogs from different ACUMOS instances. The German node of the AI4EU-platform will be hosted by KI.NRW at Fraunhofer IAIS. KI.NRW acts as an umbrella for the transfer of AI from science into companies in North Rhine Westphalia [12].

In summary, ACUMOS AI facilitates collaboration of data scientists with different competences and roles: Experts who experimentally develop AI models with dedicated toolkits and application developers who possibly retrain them and graphically compose them into deployable AI applications. The data needed for machine learning models is identified, selected, and preprocessed by data managers. The IDS with its data connectors, data transformation apps, and semantic vocabularies provides an environment where this can be done in a controlled way.

## 13.4 Machine Learning at the Edge

Having discussed that today, machine learning algorithms will typically operate on data residing in distributed and federated platforms, a natural question is whether then the classical way of using machine learning by first centralizing all data can be improved by using the distributed data for machine learning where they are. Indeed this is possible and has been proven to work extremely well even for complex models [13], so let us have a brief look at how this works and what the benefits are.

*Deep artificial neural networks* can cope with complex learning tasks because they have many parameters, namely, a weight at every link (or synapse) between two nodes at connected layers. Optimizing many weights requires lots of data [14]. Therefore, learning complex models usually takes place on a central server in the cloud, meaning that all data have to be transferred to this server. Usually, even the trained model remains in the cloud so that all application-specific data is moved to a central server. A popular example are speech assistants and translators.

However, there is a range of applications where such an approach is infeasible for legal reasons, to protect business secrets, or due to technical restrictions. An example for legal restrictions is the protection of personal data, specifically in the health sector. An example for business secrets are production data from machines, which the producing company does not want to disclose to a central service that is operated by the machine manufacturer for predictive maintenance or quality control. An example for technical restrictions are autonomous vehicles or driving assistants, which cannot transfer all data from the various cameras and sensors into a cloud in order to obtain piloting and navigation instructions.

*Distributed machine learning* solves this problem. A key idea is to learn at each of the distributed local data sources, which means at the edge of the cloud. Of course, a local node does not see much data; therefore, the second idea is to transmit the local model, rather than the local data, to the central server. Here the models are aggregated and redistributed to the local nodes. Thus all nodes indirectly profit from all data without ever exchanging it. For artificial neural networks, model aggregation means calculating the average of every parameter, which is a simple matrix operation. Figure 13.3 illustrates the communication between the central and local nodes.

All nodes continue training their model with new local data. For all nodes to continuously profit from their respective improvements, the local nodes need to exchange their updated models repeatedly. This can be done at fixed intervals or dynamically. The dynamic approach is more ambitious because here the nodes must somehow signal their progress. Distributed learning terminates when there is no more significant progress.

In 2016 researchers at Google were the first to successfully train a deep network with fixed synchronization intervals [15]. A team at Fraunhofer IAIS and Volkswagen elaborated this for dynamic synchronization and could show that the quality of the models suffered marginally while the communication effort could be reduced considerably [13].
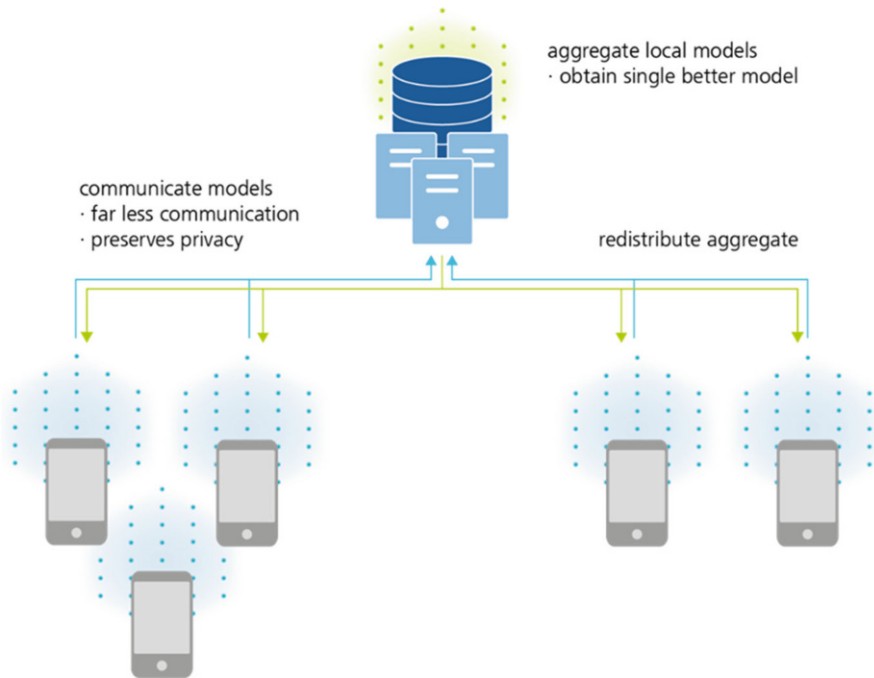
**Fig. 13.3** Communication of models in distributed machine learning

## 13.5 Machine Learning in Digital Ecosystems

Having powerful algorithms for distributed and non-distributed machine learning and the right platforms for deploying their applications as discussed in the preceding sections, let us now zoom out again and look at the overall structure of digital ecosystem that results when these technologies are used and explain the different "spaces" in which digital value is then created.

Digital business ecosystems, in practically all domains such as mobility, healthcare, industrial production, logistics, and finance, will evolve around a shared data space and profit from AI to create value from the data by improving processes or products and by generating new business models.

In digital business ecosystems, there will be many data owners with similar kinds of data, like different hospitals with medical data records or different manufacturers of intelligent cars with data from the cars' cameras and other sensors. Probably, the data will come in more or less different formats. To benefit from the ecosystem means to pass one's data to the same applications. Therefore, the IDS provides vocabularies and data transformation apps so that data with the same semantics also ends up in the same format [16].

This is an ideal situation for machine learning, because data in standardized form will dramatically reduce the data preparation effort, which may take most of the

effort in first-of-a-kind projects. In a shared data space, machine learning can be industrialized. Not only are there input data in a standardized format; a model, which was originally trained on shared data, can also be reused and transferred. Data providers, who in principle are competitors, may collaborate in training a shared basic model and tune it with their individual data to their individual context and deploy it in their distinguished smart application. Such an industrialized way of machine learning calls for a good separation of work between the different kinds of data scientists: the machine learning and other AI experts who deliver models, experts for model adaptation and application building who deliver smart applications, and data managers and data engineers who find and transform the data into a ready-to-use form.

In Fig. 13.4, the realm of the data engineers is the data space. Data engineers use the data broker to find suitable data and the app store to find data transformation apps for the target vocabularies. The data transformation apps can be applied in connectors that expose data in a preprocessed standard format.

AI development environments are the realm of machine learning experts and application builders work. In Fig. 13.4 they operate in the development space. It sits on top of the data space layer because it can receive training data via connectors from the data space. The smart applications developed in the AI development space can be deployed in domain-specific digital business ecosystems. They reside at the top of Fig. 13.4 in the solution space. Smart applications can also be published as smart apps in the app store of underlying shared data space.

Machine learning at the edge is a special case, treated on the right-hand side of Fig. 13.4. Since here, learning is a continuous activity, smart applications must be "smart learning applications" which can learn in the execution environment by aggregating local models and redistributing them. The "edge space" in Fig. 13.4 is a specialized data space that extends down to special sensors, like IoT sensors for the internet of things, and energy-efficient hardware. Data must not leave the edge space. The edge space provides smart local learning apps and connectors to exchange models between the local learning apps and the smart learning application in the solution space.

## 13.6   Trustworthy AI Solutions

With descriptions of the basic machine learning technologies, the data science process, the platforms, and digital ecosystems in place in the preceding sections, let us now return to one of the core value propositions of modern federated data ecosystems, and in particular of the industrial data space (IDS) architecture [5] and GAIA-X [16]: Trust. While the ecosystem architecture focuses on trust in the data providers, the data consumers, brokers, and transport, here we want to focus on what it takes to make the AI applications built on top of these ecosystems trustworthy.

Modern machine learning methods are extremely powerful, but due to their data-driven nature and the extremely high dimensionality of their models, establishing
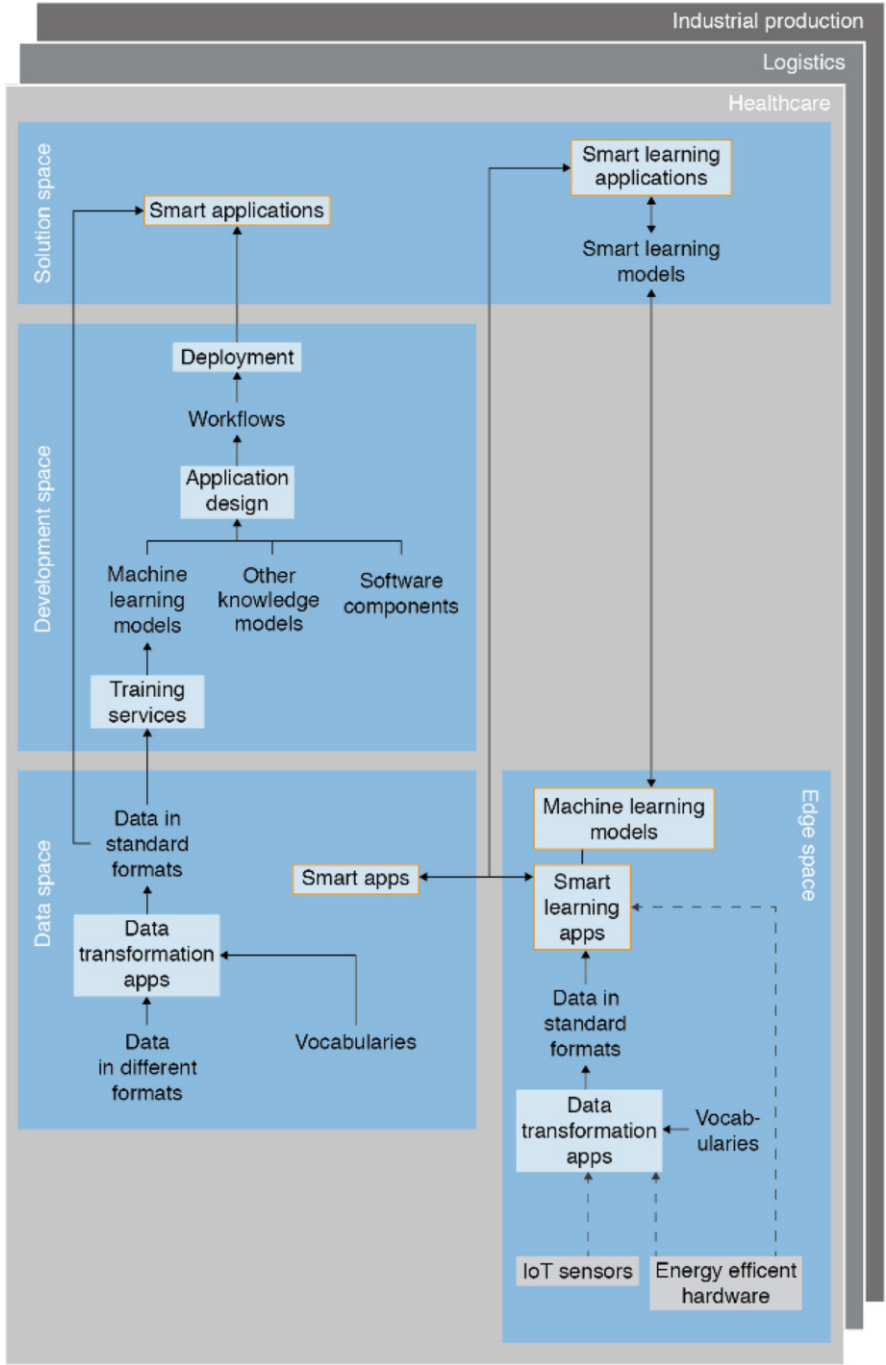
**Fig. 13.4** Ai application, machine learning, and shared data spaces

trust in their results presents particular challenges. Artificial neural networks are a particularly good example of this. These networks are "black boxes" because they do not contain any code nor rules that humans could easily inspect. Moreover, all machine learning models return results that are not perfectly true or completely false, but more or less correct. Therefore, many models can also be made to output a confidence value which reflects their uncertainty. Finally, AI applications can be built to *learn continuously*, where the model is updated from new data or feedback during operation—with possibly unforeseen effects. Edge machine learning is only one type of continuous learning.

So it is difficult to argue that an application with machine learning inside will behave as intended. Figure 13.5 gives an overview of important principles that a *trustworthy application* of AI should incorporate [17].

By choosing and configuring the type of model, the data scientist tries to achieve accurate predictions. If the model is not robust against small changes, it can be improved by adding noise and other systematic transformations of the input data. Correcter and more robust models are more *reliable*. A black-box model can be supplemented with more explicatory models to facilitate its interpretation or defend individual results. Both increase the model's *transparency*. Other methods prevent a model from exposing any *private* data which may be encoded in the millions or even billions of weights.

Quality and quantity of the training data have a huge impact on the *reliability, privacy,* and *fairness* of the model because training data may be unrepresentative in general and contain wrong features or examples of low quality. This must be investigated in particular with respect to gender, religion, ethnicity, disability, and age to improve the *fairness* of the model. Of course, all data must comply with the data protection regulations and laws.

A model cannot sense, behave, or communicate. It is conventionally coded software wrapped around and controlling the model's invocation that determines the functionality and appeal of the application at the user interface. This software must be designed according to the desired level of control. The user's *agency and oversight* is high in smart assistant tools, can be decreased while keeping the user in the loop or on the loop, and is minimized in a fully autonomous system. Applications



**Fig. 13.5**  Values to be respected by trustworthy AI applications

can be designed so convenient that users over-rely on them and unlearn important competences. Intelligent devices and robots, in particular, can be made to look and feel so human-like that users get overly attached and dependent. The embedding software also contributes to *reliability*. It must invoke the model only when the application context fits to the training data, and it must override the model's output when its confidence is low. For the worst cases, fail-safe procedures must be invoked. A final job of the embedding code is data logging so that failures of the application are documented and can be investigated, thus contributing to *transparency*.

Unfortunately, the requirements on trustworthy AI may be conflicting. Especially reliability may suffer when transparency, privacy, and fairness are improved. Moreover, creating a trustworthy AI application will be costly. Therefore, for each principle of trustworthiness, the risks of ignoring it must be assessed. The effectiveness of improvement measures should correlate to the risk, with low risks requiring no measures at all.

A European standard for auditing trustworthy AI applications could be a competitive advantage for European providers of AI software. The standard would have to balance costs against risks so that an AI certificate would be a competitive advantage and promote innovative trustworthy solutions, without raising to high barriers for market entry [18]. North Rhine Westphalia is supporting such an endeavor by Fraunhofer and partners. The so-called Bonner Katalog [19] will provide framework for certification based on the principles from Fig. 13.5 that elaborates the recommendations of the European High Level Expert Group on AI [20].

## 13.7 Summary

In this chapter, we have described how the arrival of modern federated data ecosystems acts as a driver that pushes forwards the use of artificial intelligence and machine learning technologies across all application areas. By making larger volumes of data from multiple partners available to all participants in such ecosystems in a trustworthy fashion, more and more companies will be capable of developing and/or deploying successful artificial intelligence systems. Moreover, as we have described in this chapter, recent developments in particular in distributed machine learning are a particularly good match for the environment that is provided by federated data ecosystems. Thus, we can expect that in the future, AI and machine learning will be a core part of any digital ecosystem in the manner that we have discussed above. This opens up the exciting prospect that federated data ecosystems will be the basis for a thriving economy that is characterized by fairness, competition, market orientation, and, thus, best possible value creation for enterprises and citizens alike.

# References

1. Hecker, D., Koch, D. J., Heydecke, J., & Werkmeister, C. (2017b). Big-Data-Geschäftsmodelle – die drei Seiten der Medaille. *Wirtschaftsinformatik & Management, 8*(6), 20–30.

2. Döbel, I., Lies, M., Vogelsang, M. M., Neustroev, D., Petzka, H., Riemer, A., Rüping, S., Voss, A., Wegele, M., Welz, J. (2018). *Maschinelles Lernen* [online]. *Eine Analyse zu Kompetenzen, Forschung und Anwendung*. Sankt Augustin, Fraunhofer Gesellschaft. Accessed November 11, 2020, from https://www.bigdata.fraunhofer.de/de/big-data/kuenstliche-intelligenz-und-maschinelles-lernen/ml-studie.html

3. Hecker, D., Döbel, I., Rüping, S., & Schmitz, V. (2017a). Künstliche Intelligenz und die Potenziale des maschinellen Lernens für die Industrie. *Wirtschaftsinformatik & Management, 9*(5), 26–35.

4. Paass, G., & Hecker, D. (2021). *Künstliche Intelligenz - Was steckt hinter der Technologie der Zukunft?* Springer. ISBN ist: 978-3-658-30210-8.

5. Bader, S. & Oevermann, J. (2017). *Semantic annotation of heterogeneous data sources: Towards an integrated information framework for service technicians*. Proceedings of the 13th International Conference on Semantic Systems (pp. 73–80), SEMANTICS 2017, Amsterdam, The Netherlands, September 11–14.

6. Shearer, C. (2000). The CRISP-DM model: the new blueprint for data mining. *Journal of Data Warehousing, 5*(4), 13–22.

7. ML2R. (2020). *ML2R*: *The competence center machine learning Rhine-Ruhr* [online]. Sankt Augustin: Fraunhofer IAIS. Accessed November 11, 2020, from https://www.ml2r.de/en/

8. von Rüden, L., Mayer, S., Beckh, K., Georgiev, B., Giesselbach, S., Heese, R., Kirsch, B., Pfrommer, J., Pick, A., Ramamurthy, R., Walczak, M., Garcke, J., Bauckhage, C., & Schuecker, J. (2020). Informed machine learning: A taxonomy and survey of integrating knowledge into learning systems. *arXiv, 1903*, 12394. arXiv preprint.

9. AT&T and The Linux Foundation. (2018). *Acumos*: *An open source AI machine learning platform* [online]. The Linux foundation. Accessed November 11, 2020, from https://www.acumos.org/wp-content/uploads/sites/61/2018/03/acumos_open_source_ai_platform_032518.pdf

10. AI4EU. (2020). *AI4EU*: *A European AI on demand platform and ecosystem*. THALES SIX GTS FRANCE SAS: Gennevilliers, France [online]. Accessed November 11, 2020, from https://www.ai4eu.eu/

11. Rehm, G., Bontcheva, K., Choukri, K., Hajic, J., Piperidis, S., and Vasiljevs A. (Eds.). (2020). *Towards an interoperable ecosystem of AI and LT platforms: A roadmap for the implementation of different levels of interoperability*. Proceedings of the 1st International Workshop on Language Technology Platforms (IWLTP 2020, co-located with LREC 2020) (pp. 96–107). International Workshop on Language Technology Platforms (IWLTP-2020) Marseille, France, 16.05.2020.

12. KI.NRW. (2020). *Kompetenzplattform KI.NRW* [online]: *Die zentrale Anlaufstelle für Künstliche Intelligenz in Nordrhein-Westfalen*. Sankt Augustin: Fraunhofer IAIS. Accessed November 11, 2020, from https://www.ki.nrw/

13. Kamp, M., Adilova, L., Sicking, J., Hüger, F., Schlicht, P., Wirtz, T., & Wrobel, S. (2018). *Efficient decentralized deep learning by dynamic model averaging*. Proceedings of ECML-PKDD, Dublin, Ireland, 14–16 September 2018.

14. Hecker, D., & Paaß, G. (2020). *Künstliche Intelligenz. Was steckt hinter der Technologie der Zukunft?* Springer. ISBN: 978-3-658-30211-5.

15. McMahan, H. B., Moore, E., Ramage, D., Hampson, S. & Agüera y Arcas, B. (2017). *Communication-efficient learning of deep networks from decentralized data*. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) 2017 (pp. 1273–1282), Fort Lauderdale, Florida, USA. JMLR: W&CP 54.

16. Bader, S., Pullmann, J., Mader, C., Tramp, S., Quix, C., Müller, A. W., . . . & Geisler, S. (2020). The international data spaces information model–An ontology for sovereign exchange of digital content. In *International Semantic Web Conference* (pp. 176–192). Cham: Springer.

17. Cremers, A., Englander, A., Gabriel, M., Hecker, D., Mock, M., Poretschkin, M., Rosenzweig, J., Rosstalski, F., Sicking, J., Volmer, J., Voosholz, J., Voss, A., & Wrobel, S. (2019). *Trustworthy use of artificial intelligence* [online]. Sankt Augustin, Fraunhofer IAIS. Accessed November 11, 2020, from https://www.iais.fraunhofer.de/content/dam/iais/KINRW/ Whitepaper_Thrustworthy_AI.pdf

18. Heesen, et al., (2020). https://www.plattform-lernende-systeme.de/aktuelles-newsreader/certifi cation-of-ai-systems-plattform-lernende-systeme-names-challenges.html

19. Poretschkin, M., et al. (2021). *Leitfaden zur Gestaltung vertrauenswürdiger Künstlicher Intelligenz - KI-Prüfkatalog*. Fraunhofer IAIS. https://www.iais.fraunhofer.de/de/forschung/ kuenstliche-intelligenz/ki-pruefkatalog.html

20. HLEG on AI. (2019). *Ethics guidelines for trustworthy AI* [online]. Brussels: European Commission. Accessed November 11, 2020, from https://ec.europa.eu/digital-single-market/ en/news/ethics-guidelines-trustworthy-ai