

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Hirt, Robin; Kühl, Niklas; Martin, Dominik; Satzger, Gerhard

Article — Published Version Enabling inter-organizational analytics in business networks through meta machine learning

Information Technology and Management

Provided in Cooperation with: Springer Nature

Suggested Citation: Hirt, Robin; Kühl, Niklas; Martin, Dominik; Satzger, Gerhard (2023) : Enabling inter-organizational analytics in business networks through meta machine learning, Information Technology and Management, ISSN 1573-7667, Springer US, New York, NY, pp. 1-25, https://doi.org/10.1007/s10799-023-00399-7

This Version is available at: https://hdl.handle.net/10419/305783

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.



WWW.ECONSTOR.EU

https://creativecommons.org/licenses/by/4.0/

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



RESEARCH ARTICLE



Enabling inter-organizational analytics in business networks through meta machine learning

Robin Hirt¹ · Niklas Kühl² · Dominik Martin³ · Gerhard Satzger³

Accepted: 28 April 2023 © The Author(s) 2023

Abstract

Successful analytics solutions that provide valuable insights often hinge on the connection of various data sources. While it is often feasible to generate larger data pools within organizations, the application of analytics within (inter-organizational) business networks is still severely constrained. As data is distributed across several legal units, potentially even across countries, the fear of disclosing sensitive information as well as the sheer volume of the data that would need to be exchanged are key inhibitors for the creation of effective system-wide solutions—all while still reaching superior prediction performance. In this work, we propose a meta machine learning method that deals with these obstacles to enable comprehensive analyses within a business network. We follow a design science research approach and evaluate our method with respect to feasibility and performance in an industrial use case. First, we show that it is feasible to perform network-wide analyses that preserve data confidentiality as well as limit data transfer volume. Second, we demonstrate that our method outperforms a conventional isolated analysis and even gets close to a (hypothetical) scenario where all data could be shared within the network. Thus, we provide a fundamental contribution for making business networks more effective, as we remove a key obstacle to tap the huge potential of learning from data that is scattered throughout the network.

Keywords Meta machine learning · Data confidentiality · Business network · Distributed analytics

1 Introduction

Businesses are becoming increasingly connected to enhance collaboration and co-create value. The notion of "business networks" describes two or more linked businesses that act as "collective actors" [33]. Since the emergence and growth of the internet and the digitization of many aspects of a company, digital interaction has become more important in those networks [43]. Throughout every interaction and

 Niklas Kühl kuehl@uni-bayreuth.de
 Robin Hirt robin.hirt@prenode.de
 Dominik Martin dominik.martin@kit.edu
 Gerhard Satzger gerhard.satzger@kit.edu

- ¹ prenode GmbH, Karlsruhe, Germany
- ² University of Bayreuth, Bayreuth, Germany
- ³ Karlsruhe Institute of Technology, Karlsruhe, Germany

collaboration, masses of data are produced-data that can be analyzed to generate valuable insights, for example, to optimize processes [60] or build innovative services on top of existing offerings [41, 75, 87]. Davenport [24] describes data analytics as one of the most important activities to have gained competitive advantage. Combined with the need to better understand relations and interactions in business networks [6], the need for inter-organizational analysis of distributed data sources is evident. Previous work presents concepts for centralized data analytics for distributed data sources [30, 67].

Therefore, the topic of machine learning¹ across different entities² within a value chain or business network is of high relevance [36]. However, as recent work points out, "a substantial potential for utilizing AI across company borders has remained largely untapped" [34, p. 1]. For the manufacturing

¹ We define machine learning as a set of algorithmic methods used to solve real-world problems, which can learn to computationally solve a problem instead of being explicitly programmed [53, 55].

² We define an entity as an organizational grouping of various types, e.g., a department, a company, or a consortium. An entity may have legal borders which prevent it from sharing (sensitive) data with outside parties.

Design requirement	Description
DR1	Preserve data confidentiality of the individual entities in the machine learning and execution process ³
DR2	Minimize the amount of transferred data volume during learning and execution of the machine learn- ing process between entities
DR3	Improve the prediction over a meaningful benchmark

³We define the machine learning and execution process as the lifecycle of a machine learning model from initiation to training and performance estimation up to final deployment [54]

industry, the World Economic Forum estimates the potential value of sharing analytical knowledge and associated data at over \$100 billion [14]. Other sources confirm the economic and/or public benefits of inter-organizational machine learning for other domains like health care [65, 80, 90], mobility [71], or smart cities [48].

To address this research gap, Bach et al. [9] state that for the novel challenge of machine learning in business networks, "several approaches need to be extended or rethought" [9, p. 1]. Consequently, the centralized analysis of data across businesses faces several challenges—and real-world examples unlocking the potential of cross-entity learning, i.e., acquiring analytical knowledge across different (legal) organizations, remain scarce [14]. Hirt et al. [46] analyze typical barriers for machine learning in systems—and carve out three main requirements for successful inter-organizational learning as depicted in Table 1. In the course of this work, we will design a machine learning artifact for business networks focusing on these requirements: ensuring data confidentiality (DR1) and reducing data volume (DR2)—while still ensuring superior prediction performance (DR3).

Companies are afraid of exposing sensitive information throughout the process of data analysis (DR1). The need to protect sensitive data is subject to research in the area of business networking [50, 88] or customer privacy protection and advertising [37, 42, 66]. In complex business networks, collaboration happens between multiple organizations of different legal units, hence data confidentiality is required.

As more and more data is produced, the respective transfer of large volumes of data (e.g., to a central analysis unit) can be challenging and should be addressed (DR2). Techniques like complex event processing or fog computing offer solutions to cope with growing data streams [16, 67], but still lack convincing concepts for data confidentiality preservation [91]. Additionally, as sensor sensitivity increases, not all data produced can be centralized [2]. In practice, this leads to selective centralization and/or collection of data and, thus, to a major loss of potentially relevant information.

An artifact addressing these previous requirements also needs to ensure that the resulting performance of a method leveraging the network is superior to cases where a single company would only analyze its own data (DR3). Especially the trade-off between ensuring confidentiality while allowing superior performance is worth exploring and of raised interest—and will be analyzed in detail in the course of this article.

In our work, we propose inter-organizational meta machine learning, a method that addresses all three requirements for machine learning in business networks. The kernel theory of meta machine learning [18] informs the design of our artifact. Meta machine learning combines the prediction of several base classifiers (multiple entities in a business network, e.g., suppliers) to create one aggregated prediction (single entity in a business network, e.g., original equipment manufacturer (OEM)). To demonstrate the feasibility of meta machine learning as a viable solution within business networks, we instantiate our proposed method within a working prototype and evaluate it regarding the three criteria of data confidentiality, transferred volume, and achieved predictive performance based on an industrial use case. We highlight that analytics within organizations is often a trade-off between full data confidentiality, centralization of data, and overall predictive performance. In summary, we contribute to the body of knowledge by showing that our meta machine learning method is suited to inter-organizational machine learning in terms of general technical feasibility, addressing the requirements of data confidentiality (DR1), data volume reduction (DR2), and performance (DR3). Additionally, we demonstrate its usefulness within the application context at our industry partner.

The remainder of this work is structured as follows: We first set the fundamentals for our work by elaborating on business networks and meta machine learning (Sect. 2), as well as on our methodology and research questions (RQs) (Sect. 3). We then review state-of-the-art literature as part of the theoretical background (Sect. 4). With these pre-requisites, we present our concept of inter-organizational meta machine learning and explain its architecture, the data streams as well as the necessary processes in detail (Sect. 5). This concept is then applied to a real-world case and evaluated in a technical experiment for its usefulness (Sect. 6). We conclude with a summary, limitations, and an outlook for future research (Sect. 7).

Fig. 1 Simplified business network between two or more units of a business network, based on Anderson et al. [6]



Common context described by distributed data

2 Fundamentals

In this section, we first introduce business networks and distributed data sources. Then we describe meta machine learning as a foundation for comprehensive analyses across these networks.

2.1 Business networks

We base our conceptualization of a business network on Anderson et al. [6] for a common understanding. Every business network consists of two or more units-representing for example companies or other organizations-that have a dyadic relationship. Kambil and Short [49] describe the relationship as a linkage that can have different forms, such as an alliance or hierarchy. As businesses increasingly move towards digitalization to make processes more intelligent, data is produced at each company, leaving the network with various distributed heterogeneous data sources. Such networks can be described as smart business networks [43].

In a connected world, every unit in a network possesses a piece of the puzzle in the form of distributed data sources of a common context [10], the "big picture" as illustrated in Fig. 1. To identify this big picture, those distributed data sources must be analyzed comprehensively to derive holistic insights. In an ideal setting, all units would exchange their data and freely communicate with each other. In reality, practical barriers such as the sheer volume of data that is required to be transmitted and, foremost, the exposure of data outside company boundaries, prohibit such an analysis, leaving huge potential untapped [14, 34]. While machinelearning-based solutions exist to enable secure data centralization and analysis in business networks (e.g., AWS Amazon Macie), this centralization is often not happening as data is kept confidential and is not exposed to other parties. Appropriate methods that still allow learning from a distributed, but not shared dataset are lacking. However, there is a lack of methods and providers enabling a machine-learning-based analysis on data which itself is confidential and therefore cannot be directly accessed by the analyzing party itself.

2.2 Meta machine learning

Basic machine learning techniques are commonly used to solve various real-world problems. Machine learning describes computational methods that use a series of examples ("past experience") to learn about a given task [61]. Although statistical methods are used in the learning process, a manual adjustment or programming of rules or solution strategies to solve a problem is not required. In more detail, basic machine learning uses a model that is built by applying an algorithm on a set of known data to gain insight about an unknown set of data [18, 61].

The term "meta machine learning" describes methods that employ more than one layer of learning and is "concerned with accumulating experience on the performance of multiple applications of a learning system" [18]. Džeroski and Żenko [32] argue that meta machine learning enables to "learn about learning" [18, 83]. Based on Lemke et al. [56], we define meta learning as a system that includes a learning sub-system that builds meta knowledge. Meta knowledge is extracted by a previous learning episode on one or more data sets [56]. We further differentiate between two categories of meta learning: ensemble learning and stacked generalization. Ensemble learning methods such as bagging [19] or boosting [35] propose varying data selection and processing and building different sub-models. The output of these submodels is then combined by a meta-model (e.g., majority voting).

The same principle can be applied to perform comprehensive analyses on different data sources, using stacked generalization. A dedicated sub-model is built for every data source. Their prediction is then combined through a meta model (e.g., another trained machine learning model)



Fig. 2 Possible scenarios of comprehensive analyses in business networks

to get an aggregated prediction [89]. Through the combination of predictions, the uncorrelated error between all models can be minimized, which leads to a performance increase [79].

Meta machine learning is often used to combine heterogeneous types of data to perform a comprehensive analysis. Hirt et al. [45] use a stacked generalization approach to combine different types of data (e.g., pictures and text) by employing different sub-models and combining them through a meta model, mimicking a cognitive paradigm to predict attributes of Twitter users. In the area of financial fraud detection, Abbasi et al. [1] propose a meta learning method to combine heterogeneous data sources to improve prediction performance. They use meta learning to reduce the declarative and procedural bias [83] of classifiers working on company-internal and publicly available data in one specific use case. It is often used to enhance prediction performance and combine different data sources [1]. Similarly, in the course of this work, we consider stacked generalization and its potential to solve practical problems in business network analytics. In contrast to prior work in the area of meta machine learning, we do not solely focus on its performance-enhancing properties but utilize an underestimated characteristic: the information abstraction between the sub-layer and the meta layer.

In the context of business networks, this has two advantages. Considering that sub-models are deployed at different units of a business network and send their prediction to any desired unit that inherits a meta model, we suppose only a fraction of data needs to be transmitted, compared to a transfer of raw data. Additionally, confidential information is already (and possibly irreversibly) masked through the abstraction and pre-analysis of data, making a meta machine learning analysis confidentiality-preserving.

3 Methodology

The general research is based on evaluation-centric design science according to Venable et al. [82]. To guide the design of our artifact, meta machine learning [18] and service-oriented computing [47] act as kernel theories throughout the design process for construction [39, 84].

Prior studies focus on solving the issue of disclosing sensitive data during analysis by proposing to only exchange encrypted data or masking sensitive information in data sources before exchanging it. Architectures and principles to reduce or handle the amount of transferred data during analysis do not ensure data confidentiality. Existing methods are prone to disclose sensitive information, limiting analytical methods or significantly decreasing predictive performance. As outlined in the upcoming theoretical background in detail, we identify a need for inter-organizational machine learning approaches which preserve data confidentiality [91] while reducing volume [73] and still allowing for reasonable performance [31]. For instance, there are methods capable of not exposing any raw data, but they lack in performance [7] or are not suitable to machine learning endeavors [8]. Thus, we pose our general research question (GRQ):

General Research Question (GRQ):

How can we design a well-performing meta machine learning approach allowing the holistic analysis of distributed entities within business networks while preserving data confidentiality and reducing transferred data volume?

To better understand the effectiveness and efficiency of our proposed method, we consider three scenarios, as depicted in Fig. 2. In all three scenarios, data is distributed across different units of a business network. All units collaborate in some way with each other and could potentially optimize their own output, or the output of the overall network. In scenario 1, we assume that there is no communication of data or insights of any kind between business units caused by the stated obstacles ("isolated analysis"). Every unit only performs an isolated analysis of its own data to gain insights. In contrast, in scenario 2 we consider a situation where an analysis is performed through the proposed meta machine learning method that ensures data confidentiality and reduces the volume of transferred data ("inter-organizational meta machine learning method", short IOMML). Lastly, in scenario 3, we depict an "ideal world" where obstacles such as data confidentiality and volume are non-existent, and all data is accessible by all units of a network ("shared data pool").

The first main challenge that we address is the technical evaluation of whether the three design requirements (data confidentiality, reduction of data volume, performance evaluation) are met by the proposed method, thus stating RQ1 as follows:

Research Question 1 (RQ1):

Is the proposed method effective and efficient with regard to data confidentiality, volume reduction, and prediction performance?

We expect an increase in the predictive performance of an analytics method from scenario 1 (isolated analysis) to scenario 2 (inter-organizational meta machine learning method (IOMML)), as scenario 2 has more information available than scenario 1. By comparing scenario 2 (IOMML) with scenario 3 (shared data pool), it is expected to yield a lower predictive performance of the meta machine learning method than in a scenario with complete data exchange and all raw data at disposal. Narayanan and Shmatikov [62] describe this trade-off between anonymizing/masking data and prediction accuracy: While most public datasets revealed by companies are anonymized to protect user privacy, researchers hint that perfect anonymization is not possible without damaging the utility of the data. However, distributed analysislike the one suggested in this work—yields the advantage of separate, specialized models [32]. We are interested in the performance of the meta machine learning method in comparison to a case with complete data exchange and a case with an isolated analysis.

Apart from the technical effectiveness and efficiency, we aim to gain insights on the perceived usefulness of the method in the field, more precisely, in the organizational context where it could be established. We measure perceived usefulness in our case company with the respective sub-construct from the well-established technology acceptance model (TAM) [25], similar to related work [20, 26, 44]. Thus, we state the second RQ as follows:

Research Question 2 (RQ2):

How is the proposed method perceived within its application context in terms of usefulness?

To answer both questions, we instantiate the proposed artifact within a real-world production line case with our industry partner. To strengthen generalizability, we implement an additional robustness check within a distributed sensor group system (see Appendix A.3 on page 27).

4 Theoretical background

Within the body of knowledge, we can identify two research streams in the context of enabling an analysis of distributed data sources within a business network, which is closely related to two design requirements of this work: preserving data confidentiality (DR1) and reducing the amount of transferred data in the process (DR2). To outline the research gap, we describe work in the area of data confidentiality, often called privacy preservation—an established field of research—as well as the distributed analysis of large data streams.

4.1 Preserving data confidentiality

Data privacy and confidentiality can have multiple facets and are driven by different motives and in different domains, such as social media [93], healthcare [36], industrial applications [69] or others. Belanger and Xu [12] describe privacy in online social networks and propose a multi-dimensional privacy concept fit to online social interactions. Wohlgemuth et al. [88] describe the role of security and privacy in business networking. Kieseberg et al. [50] propose an algorithm for collusion-resistant anonymization and fingerprinting of sensitive microdata. Especially the involvement of end users requires data privacy. Riquelme and Román [66] assess the influence of privacy and security on online trust for consumers. Goldfarb and Tucker [37] elaborate on privacy regulation and online advertising, while Hann et al. [42] develop a theoretical approach to overcoming online information privacy concerns.

Methods to preserve data confidentiality and privacy can be distinguished based on their main principle: masking, noising, or encryption of data (see Table 2). Additionally, there are approaches combining the previously described principles. The field of privacy-preserving data mining aims to build accurate models without disclosing an individual data record. In the following, we provide an overview of related work in the area of preserving data confidentiality in general and then describe approaches to realize confidentiality-preserving analyses and their suitability for our task at hand.

Data masking and noising are approaches that originate from the statistical sciences that strive to perform analysis

Table 2 Overview of learning methods and strategies for data-confidential learning in business networks. Legend: \bullet = fully applies, \bullet = partially applies

Strategy/method	No machine learning	Machine learning
Masking	• Armstrong et al. [7]	O Asenjo [8]
Noising	 Kocabaş and Soyata [52] 	© Dwork and Feld- man [31]
Encryption	Asenjo [8]	C Graepel et al. [38]
Aggregation	• Anagnostopoulos et al. [5]	• This work

without compromising security and privacy [29]. These approaches reduce the problem to that of extracting usable information from noisy data [22, 28]. While data noising is fairly robust to standard security attacks like the man-inthe-middle attack or an structured query language (SQL) injection, the accuracy of the analysis result often suffers from the amount of noise introduced in the initial data [4].

Besides masking and noising, encryption is another key method for preserving private information. As a comprehensive analysis of distributed data requires the transport of all data sources to a central analytics unit, encryption could be used to secure the transmission. For analysis, this data needs to be decrypted, which might already disclose data to the central analytics unit. The efficacy of this approach, therefore, depends on the safety of data in "safe" zones.

In our work, similar to masking or noising techniques, we strive to transform data to preserve data confidentiality. However, in contrast to the mentioned techniques, we aggregate the data as a part of the desired analysis to minimize the loss of information during the process ("aggregation"). By shortly elaborating on the drawbacks of existing approaches, we discuss the suitability of an aggregation technique like meta machine learning.

Compared to methods relying on encryption, our technique is able to leverage any machine learning during analysis. Although there are novel approaches that perform mathematical and rudimentary learning techniques on encrypted data [38], those do not allow for flexible use of various machine learning methods. Furthermore, performing operations on encrypted data is known for causing a high computational effort [15]. The higher computational effort and the incapability to perform machine learning on encrypted data make encryption not a suitable technique for the task at hand. Recent reports suggest that these techniques might also be vulnerable to external attacks [86]. Bhattacharya et al. [15] describe a method for privacy preserving analytics using homomorphic encryption of data among peers, enabling them to perform analyses. Their key proposition is to perform analysis on encrypted data, deducing the desired insight and, therefore, never exposing data to a third central party. Although they extend the tool set of analytical capabilities, their approach is limited to only performing basic mathematical operations (i.e., calculating the sum of products). Additionally, the computational costs are increasingly high due to the necessity of homomorphic encryption.

In the case of masking techniques, critical fields of data entry are masked to ensure confidentiality. Especially in cases where only single elements of a data entry are critical (e.g., the name of a data entry about a person), masking might be a viable option to consider. However, in the case at hand the critical data itself is the one which needs to be analyzed, making masking techniques a non-viable option.

Noising techniques strive to preserve confidentiality by adding noise to the critical data element. This enables exposing the noised data and, then, performing every machine learning technique on it. However, with increasing noising of data and, therefore, increasing data confidentiality, the predictive performance also drops significantly. Therefore, noising could be applied but has major drawbacks in terms of performance for the task at hand.

The proposed method in this work can be characterized as an aggregation technique to group and summarize critical data in a form, where the result is not exposing any private information. However, this technique aims towards realizing the aggregation by a subordinate layer of machine learning, leaving only relevant information for further analysis in the aggregated result.

As we are striving towards realizing inter-organizational machine learning, in the remainder of this paper we focus on comparing our method with the noising technique.

While masking, noising, and encryption manipulate the source data but try to keep the information content and data structure as similar as possible, we propose to only transmit information that has a direct impact on the target of the analysis.

4.2 Reducing and processing large data streams

To realize analytics in distributed systems, large volumes of data that originate from heterogeneous sources are required to be processed: sensors, transactional or social networks, or company-internal information systems. The reasons for performing analytics in those systems can be the detection of undesired behavior or other specific patterns (e.g., misconduct, unusual events, runtime errors) and to derive higher-level information from them [27, 57]. However, the increasing number of devices with access to the internet and the increasing interconnectedness of data-producing units pose challenges for data analysis infrastructure and techniques.

In general, there are two opposing literature strands of data processing and computing and, i.e., analytics in networked systems in the literature [73]: centralized and decentralized paradigms. Centralized approaches aim towards processing real-time, possibly fluctuating, data streams generated by heterogeneous, distributed units gathering low-level information in the cloud [57, 78]. In contrast, there are approaches to directly process decentralized data, like edge or fog computing [68].

Centralizing analytics requires the data to be in one place. Driven by the need to transfer and process large data streams in real time at once in order to detect undesired behavioral patterns, complex event processing (CEP) has emerged. CEP represents a set of techniques to analyze event-driven information systems. According to Luckham [58], an event is a record of an activity within a system that may depend on other events. A set of events, including their dependencies, results in a complex event containing valuable, higher-level information. In order to be able to continuously process the events gathered on distributed units, a technology must be able to apply complex analyses in parallel on several data streams [67]. For instance, CEP is used in finance to detect fraud [74] and make automatic trading decisions [3]. In addition, CEP is also used to analyze time series collected by sensors to perform real-time analyses of complex interactions measured by independent sensors [30, 85].

In contrast, decentralized analytics is driven by the utilized unused processing capabilities [72] or a need to reduce the transferred data volume [81]. Sarlis et al. [72] propose a decentralized analytics system for network traffic data, dynamically distributing parts of the decentralized data for processing and orchestrating an analysis. Uhlmann et al. [81] describe a decentralized data analytics framework for maintenance for connected manufacturers. The described system pre-analyzes sensors on site and sends the status of a machine to a central platform. Then the information is distributed and made accessible via a dashboard. Pournaras and Nikolic [64] describe on-demand self-adaptive data analytics in large-scale decentralized networks. They focus on the automated allocation of computational capacity in a network of multiple processing nodes. To strike a compromise between cloud-based and edge computing, cloudlets have evolved. A cloudlet represents the middle layer between a cloud and a mobile device, addressing latency issues of cloud architecture as well as centralization endeavors [73].

Centralized analytics imposes the necessity to transfer data to a central unit, which prohibits the analysis of sensitive data and, in fact, prohibits any analysis. Hitherto, mechanisms for decentralized analytics might impose first characteristics similar to this work. In most mechanisms, there is no additional meta-analysis of the pre-analyzed content to gain further insights. Furthermore, in the area of processing large streams of data, confidentiality is often not considered a problem. As decentralized approaches attempt to significantly reduce the network load by pre-processing low-level information on the edge device, our approach processes data where it is produced, compresses it but derives high-value information. At a central level, i.e., in a meta entity, the information from distributed units is aggregated. The information is then combined on this (central) level and analyses across multiple organizations can be performed. Thus, our proposed concept combines the advantages of central information processing with low latency due to low data volume.

5 Inter-organizational meta machine learning

To address the challenges of realizing a data confidentiality preserving method that minimizes the amount of transferred data volume in decentralized business networks, we propose an IOMML. This method uses data aggregation techniques in order not to disclose data and reduce their volume.

In the following, we suggest and describe our artifact in a general way before we instantiate it in a use case in Sect. 6. We perform this description along three perspectives: the intended architecture, the data and model output exchange during analysis, and the life cycle of an instantiation.

5.1 Architecture

As elaborated, every business network consists of different units that interact with each other. Every unit might possess its own data sources, might be owned by different legal organizations, or might be geographically distributed. The architecture should therefore enable to preserve data confidentiality and reduce the required volume of data that is transferred during analysis.

We distinguish between two unit types: sub-units and meta units. A sub-unit possesses one or many data sources (e.g., a customer data base and corresponding transaction history data) that-in combination with data of other unitsmight reveal a network-wide pattern after analysis. The meta unit represents a virtual unit that analyzes the overall situation based on the sub-units. All involved units-regardless of whether they are sub or meta-have a common understanding of the goal of the analysis. Every sub-unit possesses a certain data source that requires processing. By nature, data sources might store heterogeneous types of data that require an individual analysis. Although these separate data sources might reveal information and insights on their own, the core assumption is a pattern that is distributed throughout more than one unit. We aim to analyze and learn about these patterns-the "bigger picture". For further consideration, we assume that every sub-unit's data source reveals such a piece of the puzzle.



input



In Fig. 3, we depict a simplified structure of a meta analysis between multiple sub-units. Hereby, every sub-unit analyzes its own data, using a customized sub-model. As a result, each sub-unit prepares a collection of an item identifier, a (categorical) result, and a corresponding certainty value for the analysis. This output of the sub-model is sent to a virtual meta unit that uses a meta model to perform a comprehensive analysis. The final output of the meta analysis could then be sent back to every sub-unit as a basis for further activities or could be used by other units to build upon the insight. Note that the meta unit is just a virtual construct and could be represented by any (sub-) unit.

The data source of a sub-unit gets processed by a submodel that is highly customized to the corresponding data structure and predicts a certain attribute that might be an indicator of the big picture we are trying to reveal. As depicted in Fig. 4, after processing the sub-unit data point, the prediction output gets sent to a meta unit that aims to aggregate all incoming, subordinate predictions. Note that at no point throughout this process does raw data get transferred or exposed outside the processing unit. The meta unit does not need any information about the input data of the sub-models or how the sub-models perform their analysis. No bare information or raw data that was not intended to be

shared is distributed, thereby preserving intellectual property and confidentiality of data. To aggregate all incoming predictions and to retrieve insights from them, the meta unit employs a meta layer of machine learning that learns which sub-model's prediction is of importance in which situation. The output of such a meta model prediction is an accumulated prediction towards a distributed database. To make this prediction, the meta model draws on the stacked generalization paradigm from meta machine learning as a kernel theory [40].

The meta unit collects information of different, distributed, and independent sources to make a holistic prediction as an insight that is latently present in these data sources. It uses machine learning to gain information about the significance, relevance, and validity of each sub- model prediction and their interdependencies. Without communicating the meaning of a sub-model output, or even sub-unit data, to the meta unit, the stacked generalization meta model can still identify the desired big picture. The meta model prediction output can then be included in analytics applications or other smart services to create value.



Fig. 5 Exemplary communication in an inter-organizational meta learning landscape: data and prediction flow across involved units

5.2 Data and model output exchange during analysis

During the process of a meta prediction, data is analyzed by different sub-models, and sent to the meta unit for aggregation by the meta model. The final result is then sent back to every sub-unit, e.g., to optimize local business processes. In Fig. 5, we depict the data and model output exchange during a meta analysis in a business network. Hereby, a sub-unit possesses confidential unit data. That data is analyzed on the sub-unit's site by a sub-model, generating an abstract sub-model output. The sub-model output is then transferred to a meta unit. That step repeats for every sub-unit that is part of the analysis in a business network.

The architecture aims to maintain data confidentiality while minimizing the volume of data transmitted during analysis. In this context, it is important to consider the type of data being processed and transferred, as the reduction of volume can have a different impact on preserving confidentiality for structured and unstructured data. For structured data, such as tables, reducing the volume depends on the number of columns a row or observation possesses. Given the prediction output of a sub-classifier stays the same, the more columns a row has, the larger the reduction in volume and the higher the abstraction and security can be. This can help maintain data confidentiality by minimizing the amount of sensitive information that is transferred. For unstructured data, such as images or videos, naturally, more information is reduced by just sending a classifier output instead of the complete visual information. In this case, the data confidentiality can be high in comparison to considering the raw data. In conclusion, it is crucial to take into account the type of data being processed and the methods available for reducing its volume while maintaining its integrity.

At the meta unit's site, all incoming abstract sub-model predictions are combined by a feature aggregation, forming the input for a meta model. The meta model then processes its input and generates an output. As that output is based upon all underlying sub-model predictions, based on sub-unit data, comprehensive insights can be derived. Afterward, the meta model output can be consumed by the meta unit, every participating sub-unit, or a possible third party unit as a basis for any action or decision.

6 Evaluation in production line quality prediction

In order to evaluate the proposed IOMML artifact, we follow the FEDS framework and its application to a real-world use case according to Venable et al. [82] as depicted in Fig. 6. The FEDS framework is a framework for evaluating decision-making systems. The evaluation episodes (EEs) within the framework consist of a series of tests or scenarios that are designed to assess the performance and behavior of the system under different conditions. These episodes are used to evaluate the system's ability to handle various types of uncertainty, its robustness to different types of failures, and



Fig. 6 Evaluation episodes according to the FEDS framework [82]





its overall effectiveness in making decisions—by moving from artificial to more naturalistic evaluations with each episode as well as forming more summative than formative knowledge. The goal of these evaluations is to identify any weaknesses or limitations in the system so that they can be addressed and improved upon. We conduct two EEs in alignment with our two RQs: Evaluation episode 1 (EE1) covers the technical feasibility aspects of the artifact and its characteristics to meet our design requirements (RQ1). These requirements include privacy preservation (DR1), data volume reduction (DR2) as well as prediction performance (DR3). In the subsequent evaluation episode 2 (EE2), we cover the potential users of the artifact and the assessment of its usefulness, thus, addressing our second research question (RQ2).

In the remainder of this section, we start by thoroughly describing the industrial use case that serves as the basis for our evaluation. We showcase the technical instantiation and describe its technological foundation and a possible user interface for demonstration. To understand the technical effectiveness and efficiency of the presented meta machine learning method as our design science research (DSR) artifact, we conduct three evaluations based on an industrial use case. First, we elaborate on the data confidentiality preserving capability of our approach. Then we show the reduction of data volume that needs to be transmitted during analysis. Third, we measure the performance of the approach and compare it to two reference scenarios, as described in Sect. 3. Fourth and finally, we evaluate the artifact with experts from the related application field to assess its potential usefulness.

6.1 Use case description and suitability

The use case originates from industrial manufacturing and serves as a basis for simulating a business network with

different units. We deliberately chose a network within one legal entity which enables us to compile a benchmark case where all data is available in one place (see Fig. 2 on page 5).

A global industrial manufacturing company has provided us with a dataset that inherits data about 1,183,747 parts as well as information on whether each part has passed the quality control ("no scrap") or not ("scrap"). During the production process, each part goes through a varying sequence of several lines and their stations. The present dataset comprises 52 stations across four lines. The dataset includes 968 numeric features, 1156 date features, and 2140 categorical features. In addition to the large number of existing features, the sparse nature of the data poses an additional challenge. Most of the data instances contain empty values for more than half of the features because a part only passes through a fraction of a number of the stations. Figure 7 illustrates the paths of the parts through the different lines during the production process. Each horizontal bar represents an independent entity-in our case a production line. Each lane represents a subset of parts that undergo a production step in the respective line. As depicted in the graph, most of the parts pass lines 0 and 3 (77.4%), while only a small share includes all four lines (< 0.1%). The second and third most frequently passed paths comprise lines 1, 2, and 3 (20.7%), and lines 0, 2, and 3 (9.6%), respectively. The data itself is very imbalanced. The complete set contains only 6,879 parts labeled as faulty, which corresponds to a failure rate of 0.58%. For the overall production, it is desired to reduce the number of faulty parts by predicting future failures in time and to intervene. Often, as data is not accessible, there is a lack of quality prediction mechanisms that help to increase overall production quality by intervening and improving the production during the process. Hereby, an intervention could be done either during an ongoing production or afterwards. In the first case, potentially faulty parts could be inspected separately as they flow through the production line. This could help to detect causes for the quality issue, such as degraded production gear, or to prevent a faulty part overall. In the second case, quality issues could be detected after production but before shipment. In some cases, quality checks are also a cost driver and are only performed on a sample of the overall produced parts. Therefore, having a predictive model which can pre-determine the sample for that quality check could decrease the cost of quality management but still increase the overall quality of production.

To perform a comparative evaluation between an isolated scenario and one where data can be freely shared, we choose a use case where data sharing is possible in a test setting but not a productive system. This enables us to create a measurable benchmark for our proposed method. The production lines may be distributed geographically and even owned by different legal units.

The use case is well suited for our instantiation of interorganizational meta machine learning, as it allows us to develop an artifact addressing design requirements (DR1-3). By encapsulating the lines, we can simulate completely independent entities (DR1). In interviews with experts working in the application context, they reveal that data transfer, especially in rural areas of production sites, can be quite challenging, as the amount of data produced exceeds 60 terrabytes (TBs) per day (DR2). Our industry partner provided the data set as part of a "Kaggle competition"³ with the aim of benefiting from a community-driven increase of prediction performance (DR3).

6.2 Artifact instantiation

In Fig. 9, we depict the instantiation of the proposed meta machine learning method in our use case. For each line, a sub-model is generated that produces a sub-prediction. After receiving all sub-predictions, the meta unit-in our case a production control-receives all sub-predictions, aggregates them into a single feature array, and then analyzes it throughout the meta model. The result is a holistic quality prediction that can be used to improve production.

To evaluate the technical performance of the prediction, we use the Matthews correlation coefficient (MCC) as a metric for evaluation, which is particularly robust to class imbalance [17].

The MCC is calculated directly from the results of the binary predictions and lies in the interval [-1, 1], with values of 1 denoting perfect classification, values of -1 denoting complete disagreement and values of 0 denoting an uncorrelated relation between prediction and ground truth.

Due to its high sparsity (> 99%), previous work suggests omitting the categorical data [92]. Additionally, as we are interested in comparing methods across several scenarios, we do not include categorical data in our analysis and rather focus on numerical data. We cope with missing values (81%, cf. Fig. 8) in the numerical data by replacing them with a marker value [63]. The remaining dataset is adopted unchanged. In addition, the date information is compressed into four representative features. As shown in Fig. 9, we compare the inter-organizational meta learning approach (scenario 2) to a separate isolated analysis of data in each unit (scenario 1) and a comprehensive analysis with a shared data pool and all data in one model (scenario 3). We choose the random forest classification as it offers good results on this dataset with comparatively little training time [92]. For training, the parameter search through the parameter grid shown in Table 3 and the validation of the regular approach, we use threefold nested cross-validation to avoid overfitting [21]. Similarly, in the case of meta machine learning, we apply an adapted threefold nested cross-validation, which we have altered towards the conditions of the two-stage process to prevent data leakage. The nested cross-validation uses three outer and two inner folds. The test set of the inner fold is once again based on a three-fold cross-validation. The training set of the inner fold is used to train all sub-models, while the meta model is trained and evaluated on the predictions of the test set of the inner fold by another threefold cross-validation.

In addition to the meta machine learning classification model, we develop a microservice-based web service. This web service simulates the data generated in the individual lines, classifies these by the sub-model and transfers the results to the meta model service. This classifies the data originating from the sub-models and makes them available to the frontend.

The microservice pattern is an architectural style for software applications, whose basic idea is to split a heavyweight monolithic application into several independent, usually smaller, self-contained parts. This architecture is well-suited to the concept of meta machine learning, as the individual components are loosely connected and easily expandable. Each model, as well as an additional web frontend that visualizes the meta results in a web browser, simulating the production control, is encapsulated as a standalone microservice. Each service provides a uniform representational state transfer (REST) application programming interface (API) with exactly one endpoint. This endpoint accepts hypertext transfer protocol (HTTP) POST requests with attached JavaScript object notation (JSON) formatted text. The incoming data is processed within the service and passed to the subsequent service.

The result is a frontend (cf. Fig. 10) in which the classification results of the individual lines and the result of the meta model are displayed. For each part, the sub-model outputs are

³ Kaggle is an online community of data scientists allowing users to publish data sets, which are then part of so-called "competitions" with the aim to be "solved", which typically means an increase of prediction performances.

Fig. 8 Sparsity matrix depicting the large ratio of missing values in numeric data. Rows represent observations/parts and columns numeric features across four lines



shown as they come in. After at least one sub-model output is available, the meta model predicts an output that is also shown in the production control dashboard. This prototype illustrates the two-layer architecture of the meta machine learning approach and the dependencies between sub-models and the meta model.

6.3 Evaluation episode 1: technical evaluation

6.3.1 Preserving data confidentiality (DR1)

On the basis of the artifact (IOMML) and its instantiation, we evaluate the confidentiality aspect of meta machine learning in business networks. We define that a system is confidential when it ensures "that only authorized users access information" [23]. In our case, the users are the units. Each unit, regardless of being sub or meta, should only be able to access its own raw data. raw data can only be accessed by the unit it originates from, we compare the different scenarios of business network analyses as depicted in Table 4 based on our research framework (Fig. 2 on page 5). In the scenario of isolated analysis in separate units (scenario 1), no data is exchanged, therefore data confidentiality is preserved. In the other extreme of comprehensive network-wide analysis using all available data based on a shared data pool (scenario 3), data is, by definition, distributed among all units in the business network. Data confidentiality is therefore violated. The scenario of the meta machine learning method is of interest for further evaluation, as all subunits only have access to their own data, but the meta unit receives the output of the sub-units' machine learning models. The question remains whether data from the subunits can be reproduced from these abstract outputs-and,

To answer the first research question (RQ1), whether



Fig. 9 Instantiation of an inter-organizational meta learning method in an industrial use case with four production lines representing the subunits and a production control as a meta unit

Table 3 Parameter search space for random forest model

Parameter	Values
Number of estimators	25, 50, 100, 200, 300
Max depth	25, 50, 100, 200, 300

consequently, whether data confidentiality is preserved or violated.

To answer this question, we first need to regard the raw data in each sub-unit. The dataset contains an extremely large number of anonymized features. Features are named according to a convention that tells reports on the production line, the station on the line, and a feature number. For example, L3_S36_F3939 is a feature measured on line 3, station 36, and is feature number 3939. An example of an observation is depicted in Table 5. Every row represents one part that is described by different features at each station. Every

feature represents measurements performed for the specific part at the respective station during the production process.

Now each sub-unit builds its own model, based on the goal of predicting the target value (scrap, no scrap), and communicates this prediction to the meta unit. The communicated output of prediction and its probability is depicted in Table 6. For each part, the sub-set of data is analyzed at each line towards the attribute "scrap" or "no scrap". Each sub-prediction also contains a probability score of the respective prediction.

There is no possibility of reconstructing the raw data from Table 5 with the abstract predictions of Table 6. The machine learning model of each sub-unit is highly complex and a reconstruction from a binary value and a probability is impossible, as the nature, amount and type of the raw features are unknown to the meta unit. We can therefore positively answer RQ1, as data confidentiality is preserved in the scenario of meta machine learning.

		Line 1	Line 1		Line 2			Meta /combi	ned
Part ID	SoP	Prediction	Confidence	Prediction	Confidence	Prediction	Confidence	Prediction	Confidence
232	10:57:55	No Scrap	63%	No Scrap	56%	No Scrap	23%	No Scrap	93%
234	10:57:56	No Scrap	33%	No Scrap	42%	No Scrap	24%	No Scrap	82%
332	10:57:57	Scrap	63%	No Scrap	26%	Scrap	77%	Scrap	87%
678	10:57:58	No Scrap	33%	No Scrap	42%	No Scrap	24%	No Scrap	82%
325	10:57:59	No Scrap	63%	No Scrap	72%	No Scrap	25%	No Scrap	92%
097	10:58:00	No Scrap	73%	No Scrap	42%	No Scrap	23%	No Scrap	51%
834	10:58:01	Scrap	77%	Scrap	67%	Scrap	23%	Scrap	82%
876	10:58:02	Scrap	63%	No Scrap	56%	No Scrap	23%	No Scrap	93%
212	10:58:03	No Scrap	33%	No Scrap	42%	No Scrap	24%	No Scrap	82%
331	10:58:04	No Scrap	33%	No Scrap	42%	No Scrap	44%	No Scrap	82%
464	10:58:05	No Scrap	33%	No Scrap	42%	No Scrap	24%	No Scrap	82%
122	10:58:06	No Scrap	33%	No Scrap	42%	No Scrap	34%	No Scrap	82%
146	10:58:07	No Scrap	23%	Scrap	42%	No Scrap	64%	Scrap	52%

Fig. 10 Web front-end of the instantiated artifact

 Table 4
 Comparison of scenarios in regard to data confidentiality

Data availability	Scenario 1: Isolated analysis in separate units	Scenario 2: Meta machine learning method	Scenario 3: Comprehensive network-wide analysis using all available data based on a shared data pool
Sub-units	Access to own data	Access to own data	Access to all data
Meta unit	n/a	Access to output of sub-units machine learning models	n/a

Table 5 Excerpt of raw data for sub-units 0 and 3

(a) Line 0								
Part ID	Line 0 sub-unit							
	L0_S0_F0	L0_S0_F2	L0_S0_F4		L0_S0_F6			
# 001	-0.042	-0.049	-0.015		0.003		No scrap	
# 002	-0.023	-0.049	-0.015		-0.016		No scrap	
(b) Line 3								
Part ID	Line 3 sub-unit				Target			
	L3_S35_F3889	L3_S35_F3894	L3_S35_F3896					
# 001	-0.079	0.030	-0.200		No scr	ар		
# 002	0.049	-0.030	0.072		No scr	ар		

Table 6 Excerpt of sub model output and probabilities for sub-units 0 to 3

(a) Lines 0 and 1				
Part ID	Line 0 sub-unit		Line 1 sub-unit	
	Prediction	Probability	Prediction	Probability
# 001	No scrap	97.24%	No scrap	98.67%
# 002	No scrap	99.29%	No scrap	99.71%
(b) Lines 2 and 3				
Part ID	Line 2 sub-unit		Line 3 sub-unit	
	Prediction	Probability	Prediction	Probability
# 001	No scrap	100%	No scrap	100%
# 002	No scrap	100%	No scrap	98.28%

6.3.2 Reduction of transferred data volume (DR2)

We also ask whether transferred data volume can be drastically reduced during comprehensive analyses. In this section, we evaluate the reduction of transferred data volume in business networks between sub-units and a meta unit by comparing the different scenarios depicted in Fig. 2. In this case, units are represented by organizational business units, sites, or companies.

Assuming that all features require the same amount of space s and there are k sub-units with a varying number of features, the transferred volume for a comprehensive network-wide analysis using all available data in a shared data pool (scenario 3) is composed of the sum of all sub-units' features multiplied by the feature size. However, no data is transferred in the case of isolated analysis within individual units (scenario 1) without data exchange. In comparison, applying the proposed architecture of inter-organizational meta machine learning (scenario 2), every sub-unit individually analyses its own data (i.e., features produced by a certain organizational unit) and only transfers the output to the meta unit. These three scenarios with their transferred data volume between sub-units and the meta unit is depicted in Table 7. Thereby, the volume of data to be transferred in scenario 2 is reduced compared to scenario 3, assuming that the number of output features m of a sub-model is smaller than its number of input features n. This leads to savings of $(\sum_{i=1}^{k} n_i - k * m) * s$ when considering scenario 2 compared to scenario 3. Accordingly, the reduction ratio is described by $\frac{k*m}{\sum_{i=1}^k n_i}$.

Regarding the industrial use case from our evaluation, we have four production lines as sub-units with different numbers n of features or columns per data instance: $n_i \in \{173, 519, 48, 251\}$. Each sub-model predicts a certain output based on its input features. Due to the very small number of output features (m = 2) compared to the number

of input features (scenario 2), the data volume to be transferred to the meta unit is reduced to 0.81% of the volume in the case of complete information in a shared data pool (scenario 3) considering our presented industrial use case. We can therefore answer RQ2 and demonstrate that our method enables the drastic reduction of the required amount of transferred data volume.

6.3.3 Performance of method (DR3)

Finally, we are interested in the performance of our method in comparison to meaningful benchmarks-and estimate the "loss of privacy" of a scenario with meta learning and distributed data sources in comparison to one shared data pool. In Sect. 3, we give an overview of our research design and consider three scenarios that require comparison: In the first scenario, units in a network perform an isolated analysis. In the second one, we consider our meta machine learning method to realize comprehensive analysis. In the third scenario, we draw on a complete analysis of all data available in one shared data pool. By comparing the performances between scenarios 1 and 2, we expect to see a performance increase due to the comprehensive meta learning approach. Between scenarios 2 and 3, two effects could occur: increased performance (performance gain) due to the application of stacked generalization or performance loss due to the processing of prediction outputs rather than raw data (loss of abstraction).

 Table 7 Comparison of scenarios regarding data volume with k - amount of sub-units; n - number of input features; m - number of output features of sub models; s - volume of a feature

Scenario 1	Scenario 2	Scenario 3
0	<i>k</i> * <i>m</i> * <i>s</i>	$\sum_{i=1}^k n_i * s$





 Table 8
 Technical performance of method compared to other scenarios

Model	MCC	Model parame	ers
		#estimators	Max depth
Sub-model line 0	0.1935	100	25
Sub-model line 1	0.2151	50	50
Sub-model line 2	0.2326	200	50
Sub-model line 3	0.1950	100	50
Meta model	0.2822	50	25
Complete model	0.2965	300	200

During the meta machine learning classification process, one sub- model per line is trained. However, in the chosen dataset, not all parts pass each of the four lines. Figure 11 depicts the relative amount of parts passing a certain line. Accordingly, for parts that pass only a subset of lines, only predictions of sub-models of these lines are used as input features for the meta model.

We present an overview of our results in Table 8. The sub-model performances in the form of an MCC range from 0.1935 to 0.2326 (for additional metrics see Appendix A.1 on page 25). As depicted in Figs. 5 and 9, not every part passes every line, making a comparison of results of the sub-models difficult. However, we can see that the meta model yields a performance increase of 21.32% compared to the best performing sub-model by reaching an MCC of 0.2822. We can conclude that the meta model aggregates the information of the sub-model outputs and is able to draw comprehensive conclusions that are superior to the ones of the sub-models (performance increase).

These results are consistent with the findings of Džeroski and Ženko [32]. We can therefore already partly answer RQ2, as we observe a significant increase in statistical performance when comparing an isolated scenario with the applied meta machine learning method.

As expected by Narayanan and Shmatikov [62], in scenario 3 (shared data pool) we reach a slightly superior performance of 0.2965 compared to all regarded baselines, surpassing the meta model's performance by 5.07% (loss of abstraction). Despite the extremely low information content of the training data visible to the meta model compared to the complete classification, the performance deteriorates only slightly. We can therefore fully answer RQ2, as we regarded the baselines of all scenarios.

Table 8 depicts the optimal model parameters, number of estimators, and maximum tree depth of the respective models. Especially the maximum tree depth ranges between 25 and 50 for all sub-models and the meta model. The complete model trained on all lines performs best with a maximum tree depth of 200. The estimators' parameter representing the number of trees used for a model varies between 50 and 200 for sub-models and meta models, and is 300 for the complete model.

Summarizing the results, we show the technical feasibility of our method regarding data confidentiality preservation (DR1) and data volume reduction during a comprehensive analysis (DR2). We identify a performance gain (DR3) that is enabled by our method in comparison to an isolated analysis (scenario 1 vs. scenario 2), but also a performance loss (scenario 2 vs. scenario 3) due to the analysis of abstract prediction outputs (loss of abstraction). Although this performance loss seems rather small, it depicts a consideration between performing a comprehensive analysis of all raw data sources at once or a confidentiality-preserving one. We denote the "price of privacy" as the difference between the effectiveness of a scenario with perfect data availability (but a violation of privacy) and a distributed meta-analysis without the exposure of sensitive data. In our case, the price is Table 9Overview of workshopparticipants

Workshop participant	Position at industry partner	Scope of duties	Time with industry partner
α	Project manager	Project organization and line rollout	5 years
β	Expert team leader	Research and production management	5 years
γ	Head of Department	Multi-project management, IT and architecture governance, Software development and opera- tions	8 years

rather small (5% loss of MCC), but we gain the possibility to hide the raw data from other units in a business network-and still allow them to cooperate in terms of holistic analyses. Compared to our proposed approach, noising as an alternative shows a significantly higher price of privacy (12% loss of MCC, cf. Appendix A.2 on page 27). In general, the entire business network can profit from such analyses, as the comparison of performance to isolated analyses is remarkable and the scenario with a shared data pool is highly improbable for different legal units [51]. Furthermore, we show the increased performance of our method also in an additional industrial use case (see Appendix A.3 on page 27 for more details).

6.4 Evaluation episode 2: usefulness

After the technical evaluation of the artifact, we now aim to evaluate its usefulness within its designated application field (RQ2). To this end, we discuss the developed artifact with practitioners from our industry partner as part of a workshop. The aim is to gain feedback on the artifact in general as well as its perceived usefulness. The workshop participants are from different divisions with different roles in the company. An overview of their characteristics is depicted in Table 9.

We elaborate on the artifact's capabilities, demonstrate it and let them interact with it. We discuss advantages and disadvantages and provide the experts with a short questionnaire on the perceived usefulness using the measures developed by Davis [25]. As the artifact is in an early stage and usability aspects were not of interest, we omit measures of ease of use in this evaluation episode and focus on the more general aspect of artifact adoption, regardless of the detailed user interface choices [77]. The perceived usefulness measure prompted participants to indicate their level of agreement on six items about how the artifact would enable them to perform tasks quicker, increase their performance on the job, increase their productivity, increase their effectiveness, increase their easiness in the job environment as well as an assessment on the general usefulness. Responses range from "very unlikely" (1) to "very likely" (5) on a 5-point Likert-type scale. Several studies have indicated satisfactory reliability for perceived usefulness in TAM for artifacts in an early development stage [70]. The results of the aggregated questionnaire are depicted in Table 10.

All participants (n = 3) demonstrate a positive attitude towards IOMML with a median of "4" in all six questions. In discussion with the experts, multiple aspects arise. First of all, β mentions that fast analyses are often important in their daily work: "With over 60 TB of transferred sensor data per day, any abstraction that still allows analyses is beneficial to us". Participant γ tributes that the incorporated process model also contains the training phase, which is often neglected when implementing IT artifacts. However, he is doubtful about the necessary incentive of the affected employees within an organization to implement a system that first has to be trained for a certain amount of time before it can be put into production. Both, α and β note that the aspect of the live analysis of distributed data sources with meta machine learning would be highly beneficial, because in the current state such analyses (if possible at all) could only be done after something went wrong, e.g., a part not being within quality. Then the department typically starts an intensive investigation, which becomes very complicated once it leaves company borders. When discussing a possible productive implementation, α notes that some suppliers would even be open to sharing data for analyses to increase their unique selling point towards an OEM. Within the same legal entity, access to both the raw data or abstracted predictions would not be an issue (α and β).

In regard to other application areas within their company, they note that only critical processes would be of interest.

Table 10 Results of an expert workshop on the perceived usefulness of IOMML. Items are rated on a Likert scale of 1 ("unlikely") to 5 ("likely"). N=3

Item	Median	SD
Using IOMML in my company would enable us to accomplish tasks more quickly	4	0.58
Using IOMML would improve our job performance	4	0.58
Using IOMML would enhance our effectiveness on the job	4	0.58
Using IOMML would make it easier to do our job	4	0.58
I would find IOMML useful	4	1.00

All three experts raise legal concerns and elaborate that this aspect needs more attention.

7 Conclusion

This work aims to overcome the data confidentiality and transfer volume barriers caused by distributed data sources across different units in business networks. Specifically, we propose an inter-organizational meta machine learning method (IOMML) built on meta machine learning and service-oriented knowledge as kernel theories. In our setup, we differentiate between various scenarios in a business network, instantiate our method based on an industrial use case and evaluate it according to the feasibility of preserving data confidentiality and reducing the volume of transferred data during analysis and the overall prediction performance. We show, first, implications for its suitability in a production control interface implemented via a service-oriented architecture. Furthermore, we discuss the potential usefulness of the artifact with practitioners. Our contribution to the body of knowledge is threefold: First, we propose a flexible method that can be used in business networks to perform comprehensive analyses on a distributed data source and show its technical feasibility in terms of a prototypical instantiation, preserved data confidentiality, reduced data volume, and statistical performance. Second, we show that the artifact is perceived as useful within its application context. Third, we show that the method of IOMML could be well feasible compared to the two scenarios of either sharing all data or no data within a business network.

In addition to these theoretical contributions, concrete managerial implications are obvious: The proposed method allows units in business networks to share insights without exposing data-a possibility that has so far been limited in traditional settings. Especially in co-opetition networks [13] such a method can lower the barrier for individual units to collaboratively work on insights that are a shared interest among all parties. However, even if all units would (in theory) agree to share all data, it would be technically challenging to transfer all data, especially in production scenarios with large data streams [76]. With the drastic data volume reduction of the proposed method, analyses of large, distributed data sources become possible. Lastly, the application of the method would facilitate comparability among different units and drive standardization towards a uniform structure and schema of gathered data. This would be especially true for all platforms thriving on shared data, for example in the area of predictive maintenance.

While there is potential for theory and practice, our work also poses several limitations that need to be addressed in future research. As of now, we only instantiate the developed method in an artificial industrial use cases to test its feasibility. However, additionally, we conduct a robustness check on a second industrial case (see Appendix A.3 on page 27). In our main evaluation, the test performed with the artifact involved units of the same organization. To generalize and deduce insights on its projectability to other problems and domains, further evaluation and studies are needed. Future work requires researchers to elaborate on how the proposed method can be applied in a real-world business network. For example, a consortium of different value co-creating businesses could apply this method in an experimental setup to observe and size individual benefits. Furthermore, we do not include concrete aspects of the instantiation of our approach using IT systems or services, as we only address the conceptual aspects of the information flow between business entities, but not infrastructure-specific properties. Additionally, we only evaluated the perceived usefulness of the artifact, not its actual usefulness and usability in use [11]. We evaluate the technical efficiency of the proposed approach to preserve the confidentiality of data originating from subordinate entities. However, we have to acknowledge the possibility of information leakage through the sub-predictions. By analyzing the aggregated sub-predictions, one could for instance derive insights into the reliability of each entity. Thus, we can only account for preserving the raw information values of each entity and not overlying concepts or paradigms that might or might not materialize through abstract subpredictions. However, we also observe a continuum between the absence of inter-organizational analytics-and a full exchange and exposure of data. In this continuum, the level of shared information increases. Organizations have to make a trade-off: living with a fraction of analytical insights, or opening up-and potentially exposing information through the sub-predictions, but receiving system-wide insights.

Regarding the technical dimensions of the proposed method, we only reviewed stacking as a possibility of meta machine learning. It would be interesting to explore alternative types and algorithms, for example, distributed deep networks. As a basis for these algorithms, the features for meta learning could be altered and additional information could be communicated to the meta unit besides prediction and probability, such as the number of features, training parameters, or additional meta data. Apart from the technical aspects of our work, a thorough assessment of the organizational aspects of the proposed method is still required. This includes but is not limited to, questions on how the proposed method would perform in a real-world scenario, how a system would need to be designed to incentivize all entities to participate, and how and where the meta unit is governed. This includes the legal dimensions, questions of ownership, and liability. Finally, while the method is able to preserve the confidentiality of subunits' attributes to other units during analysis, it is not able to mask the existence of the instance itself, which limits its privacy-preserving characteristic. Despite these limitations, the proposed method could fundamentally change the way of communication between the units of a business network, foster system-wide analytics, and, therefore, improve overall network productivity.

Appendix

Production line quality prediction

Complete model

See Tables 11 and 12.

Table 11 Confusion matrix

		Predicted	
		No scrap	Scrap
Actual	No scrap	1180766	2981
	Scrap	5177	1702

sis
5

		#estimators				
		25	50	100	200	300
Max depth	25	0.2731	0.2822	0.2806	0.2752	0.2760
	50	0.2648	0.2660	0.2531	0.2621	0.2519
	100	0.2469	0.2447	0.2490	0.2498	0.2492
	200	0.2313	0.2328	0.2298	0.2262	0.2199
	300	0.2374	0.2310	0.2294	0.2327	0.2134

Bold indicates the highest MCC values

Sub model 0

See Tables 16 and 17.

Table 16 Confusion matrix

		Predicted	
		No scrap	Scrap
Actual	No scrap	1175324	8423
	Scrap	5221	1658

MCC	0.296544
Accuracy	0.993148
F1-Score (weighted)	0.992501
Precision (weighted)	0.991982
Recall (weighted)	0.993148
Cohen's Kappa	0.291095

Table 17	Metrics	

MCC 0.193483 Accuracy 0.988540 F1-Score (weighted) 0.989614 Precision (weighted) 0.990776 Recall (weighted) 0.988540 Cohen's Kappa 0.189955

Meta model

Table 12 Metrics

See Tables 13, 14 and 15.

Table 13 Confusion matrix

		Predicted	
		No scrap	Scrap
Actual	No scrap	1180558	3189
	Scrap	5231	1648

Sub model 1

Table 19 Metrics

See Tables 18 and 19.

Table 18 Confusion matrix

		Predicted	
		No scrap	Scrap
Actual	No scrap	1176646	7101
	Scrap	5164	1715

Table 14 Metrics

MCC	0.282242
Accuracy	0.992928
F1-Score (weighted)	0.992315
Precision (weighted)	0.991805
Recall (weighted)	0.992928
Cohen's Kappa	0.277880

MCC	0.215102
Accuracy	0.989699
F1-Score (weighted)	0.990330
Precision (weighted)	0.991002
Recall (weighted)	0.989699
Cohen's Kappa	0.213436

Sub model 2

See Tables 20 and 21.

Table 20 Confusion matrix

		Predicted	
		No scrap	Scrap
Actual	No scrap	1178458	5289
	Scrap	5243	1636

Table 21 Metrics			
	Metrics	MCC	0.232585
		Accuracy	0.991154
	F1-Score (weighted)	0.991169	
	Precision (weighted)	0.991184	
		Recall (weighted)	0.991154
		Cohen's Kappa	0.232584

Sub model 3

See Tables 22 and 23.

Table 22 Confusion matrix

			Predicted	
			No scrap	Scrap
Actual		No scrap	1177577	6170
		Scrap	5432	1447
Table 23	Metrics		MCC Accuracy F1-Score (weighted) Precision (weighted) Recall (weighted)	0.195008 0.990256 0.990502 0.990755 0.990256

Noising

Noising techniques strive to preserve confidentiality by adding noise to the critical data element. The predictive performance also drops significantly with increasing noising of data and therefore increasing data confidentiality. Figure 12 shows this effect, where a noise term of 0 means that there is no confidentiality at all (no noising) and a term of 1 means that the original data is obscured by white



Fig. 12 Sensitivity of model performance with an increased ratio of additive noise term

noise of the order of the standard deviation of the data itself. We assume that a sufficiently strong preservation of confidentiality accompanies this. We applied ascending noise terms to the data in 0.1 increments to highlight the trade-off between these two extremes.

Robustness check: distributed sensor groups

We evaluate the robustness of the proposed instantiation of IOMML by means of an additional use case in the field of operation and maintenance.

Use case description

As a second data set we consider an example, where the status of a function-critical component (seal) in a hydraulic application as well as corresponding sensor measurements are available. It is technically not possible to observe the condition of the component of interest directly—as, for instance, sensors cannot be mounted at the component. Therefore, data from sensors in the nearby environment of the component could be leveraged to infer the state of the seal by means of machine learning [59].

The data set consists of 2.230.992 instances of time-independent recording intervals. each with an associated state description (*no failure. assembly failure and damage*) which are almost balanced (*no failure: 48.36%; assembly failure:* 26.1% and damage: 25.51%). Overall, the data set contains 46 features which can each be assigned to one of six physically and logically separated groups of sensor measurement points. These groups are each assigned to different legal units due to structural separation and connection to separate gateways, as each sensor group originates from a different manufacturer with its own, proprietary IoT platform. The use case is well suited for our instantiation of IOMML, as it complies with design requirements 1 to 3: The separate sensor groups each represent independent entities (DR1). Because the gateways transmit data through bandwidth-limited communication channels such as CAN bus, the smallest possible transmission size is absolutely necessary (DR2). Likewise, preliminary interviews with the industry partner, who provides the data, have shown that the best possible prediction performance is an essential requirement to be able to initiate maintenance measures at an early stage (DR3).

Artifact instantiation

Equivalent to the procedure described in Sect. 6, a sub-prediction and a corresponding certainty value are generated for each sensor group, which is subsequently received by a meta unit and analyzed in aggregated form by a meta model. The result is a holistic state description of the functionally critical component. Also, here we compare the inter-organizational meta learning approach (scenario 2) to a separate isolated analysis of data in each unit (scenario 1) and a comprehensive analysis with a shared data pool and all data in one model (scenario 3). All classification models utilize the random forest algorithm and are validated in a nested cross-validation.

Evaluation episode 1: technical evaluation

Since training of sub-models condenses the complex features of the sensors into a prediction about the state of the functioncritical component, the original data cannot be reconstructed. Strictly speaking, the original features of the data set describe numerical values such as temperatures or pressures at certain locations within the system, while the results of the sub models only give a binary prediction result and its probability. Thus, data confidentiality is preserved in the scenario of meta machine learning (scenario 2) in contrast to scenario 3 (RQ1). Considering the calculation logic depicted in Sect. 7, scenario 3 results in a data volume of 46 times the volume of a single feature, while in scenario 2 this volume can be reduced to 12 (6 sub models times two output features) times the volume of a feature. Thus, the amount of data transferred in scenario 2 is reduced to 26% of the amount of data in scenario 3.

See Table 24.

In terms of predictve performance, we observe a similar effect as in the production line case. In Appendix 24, we present the results for scenario 1 to 3 in terms of the respective MCC. We observe a performance gain from scenario 1 to scenario 2 for every sub-model. Hereby, the sub model from "sensor group 4" only reaches an MCC of 0.0744, and the model which is trained on data originating from sensor group 1 performs best with an MCC of 0.7250. In comparison, the aggregated meta model reaches an MCC of 0.7920,

Table 24 Technical performance of method compared to other scenarios

Model	MCC
Sub-model sensor group 0	0.6677
Sub-model sensor group 1	0.7250
Sub-model sensor group 2	0.5393
Sub-model sensor group 3	0.5249
Sub-model sensor group 4	0.0744
Sub-model sensor group 5	0.1812
Meta model	0.7920
Complete model	0.9543

outperforming the worst sub-model by 970.27% and the best by 9.24%. Similarly, as in the previous case, we observe a performance loss from scenario 2 to scenario 3 of 17.00%. for this use case.

Complete model

See Table 25 and 26.

Table 25 Confusion matrix

		Predicted		
		No failure	Assembly failure	Damage
Actual	No failure	1048299	16513	12980
	Assembly failure	10170	562431	9687
	Damage	5490	9841	553349

Table 26 Metrics			
	Metrics	MCC	0.954285
		Accuracy	0.970979
		F1-Score (weighted)	0.971026
		Precision (weighted)	0.971148
		Recall (weighted)	0.970979
		Cohen's Kappa	0.954239

Meta model

See Table 27 and 28.

Table 27 Confusion matrix

		Predicted		
		No failure	Assembly failure	Damage
Actual	No failure	986600	38605	52587
	Assembly failure	66409	461343	54536
	Damage	32209	48990	487481

Assembly failure

108553

393083

135690

Damage

121540

89700

334449

0.524877

0.701127

0.698990

0.698634

0.701127

0.524224

Table 28 Metrics

MCC	0.792021
Accuracy	0.868386
F1-Score (weighted)	0.868094
Precision (weighted)	0.868397
Recall (weighted)	0.868386
Cohen's Kappa	0.791788

Sub model 2

Actual

Table 34

See Tables 33 and 34.

Table 33 Confusion matrix

No failure

Damage

Assembly failure

Sub model 0

See Tables 29 and 30.

Table 29 Confusion matrix

		Predicted		
		No failure	Assembly failure	Damage
Actual	No failure	926791	65712	85289
	Assembly failure Damage	100310 64201	415950 86228	66028 418251

MCC

Accuracy

F1-Score (weighted)

Precision (weighted)

Recall (weighted)

Cohen's Kappa

Metrics	MCC	0.539291
	Accuracy	0.706775
	F1-Score (weighted)	0.707651
	Precision (weighted)	0.709522
	Recall (weighted)	0.706775
	Cohen's Kappa	0.538895

Predicted No failure

847699

99505

98541

Sub model 3

0.667663

0.790122

0.789722

0.789413

0.7901220.667620 See Tables 35 and 36.

Table 35 Confusion matrix

		Predicted		
		No failure	Assembly failure	Damage
Actual	No failure	891411	100951	85430
	Assembly failure	132810	353791	95687
	Damage	97820	153418	317442

MCC

Accuracy

F1-Score (weighted)

Precision (weighted)

Recall (weighted)

Cohen's Kappa

Sub model 1

Table 30 Metrics

See Tables 31 and 32.

Table 31 Confusion matrix

		Predicted		
		No failure	Assembly failure	Damage
Actual	No failure Assembly failure Damage	986704 87540 65418	42098 439058 85331	48990 55690 417931

MCC	0.724988
Accuracy	0.827228
F1-Score (weighted)	0.825501
Precision (weighted)	0.825219
Recall (weighted)	0.827228
Cohen's Kappa	0.724221

Sub model 4

Table 36 Metrics

See Tables 37 and 38.

Table 37 Confusion matrix

		Predicted		
		No failure	Assembly failure	Damage
Actual	No failure	637025	254065	186702
	Assembly failure	231050	191655	159583
	Damage	218352	253680	96648

Table 38 Metric

able 38 Metrics	MCC	0.074382
	Accuracy	0.415176
	F1-Score (weighted)	0.411569
	Precision (weighted)	0.410816
	Recall (weighted)	0.415176
	Cohen's Kappa	0.074030

Sub model 5

See Tables 39 and 40.

Table 39 Confusion matrix

		Predicted		
		No failure	Assembly failure	Damage
Actual	No failure Assembly failure	670710 195684	198536 148040	208546 238564
	Damage	126985	203305	238390

Table 40 Metrics

MCC	0.181204
Accuracy	0.474318
F1-Score (weighted)	0.478521
Precision (weighted)	0.485576
Recall (weighted)	0.474318
Cohen's Kappa	0.180575

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

- 1. Abbasi A, Albrecht C, Vance A, Hansen J (2012) Metafraud: a meta-learning framework for detecting financial fraud. MIS Q 36:1293-1327
- 2. Abu-Elkheir M, Hayajneh M, Ali NA (2013) Data management for the internet of things: design primitives and solution. Sensors 13:15582-15612
- 3. Adi A, Botzer D, Nechushtai G, Sharon G (2006) Complex event processing for financial services. Services Computing Workshops, 2006. SCW '06. IEEE
- 4. Agrawal R, Srikant R (2000) Privacy-preserving data mining. In: Proceedings of the 2000 ACM SIGMOD international conference on management of data, pp 439-450
- 5. Anagnostopoulos C, Savva F, Triantafillou P (2018) Scalable aggregation predictive analytics. Appl Intell 48:2546-2567
- 6. Anderson JC, Hakansson H, Johanson J (1994) Dyadic business relationships within a business network context. J Mark 58:1
- 7. Armstrong MP, Rushton G, Zimmerman DL (1999) Geographically masking health data to preserve confidentiality. Stat Med 18:497-525
- 8. Asenjo JC (2017) Data masking, encryption, and their effect on classification performance: trade-offs between data security and utility. Ph.D. thesis. Nova Southeastern University
- Bach V, Gao J, Chen X (2020) Special issue: machine learning in business networks. Electronic Mark. http://www.electronic markets.org/call-for-papers/single-view-for-cfp/datum/2018/09/ 08/cfp-special-issue-onmachine-learning-in-business-networks/
- 10. Bach V, Vogler P, Österle H (2013) Business knowledge management: Praxiserfahrungen mit Intranetbasierten Lösungen. Springer, Berlin
- 11. Bagozzi RP (2007) The legacy of the technology acceptance model and a proposal for a paradigm shift. J Assoc Inf Syst 8:3
- 12. Belanger F, Xu H (2015) The role of information systems research in shaping the future of information privacy. Inf Syst J 25:573-578
- 13. Bengtsson M, Kock S (2000) Coopetition in business networks-to cooperate and compete simultaneously. Ind Market Manag 29:411-426
- 14. Betti F, Bezamat F, Fendri M, Fernandez B (2020) Share to gain: unlocking data value in manufacturing. World Economic Forum. https://www3.weforum.org/docs/WEF_Share_to_Gain_ Report.pdf
- 15. Bhattacharya P, Phan T, Liu L (2015) Privacy-preserving distributed analytics: addressing the privacy-utility tradeoff using homomorphic encryption for peer-to-peer analytics. In: Proceedings of the international conference on information systems, ICIS, pp 1-11
- 16. Bonomi F, Milito R, Zhu J, Addepalli S (2012) Fog computing and its role in the internet of things. In: Proceedings of the first edition of the MCC workshop on Mobile cloud computing, pp 13-16
- 17. Boughorbel S, Jarray F, El-Anbari M (2017) Optimal classifier for imbalanced data using matthews correlation coefficient metric. PLoS ONE 12:e0177678
- Brazdil P, Carrier CG, Soares C, Vilalta R (2008) Metalearning: 18. applications to data mining. Springer Science & Business Media
- 19 Breiman L (1996) Bagging predictors. Mach Learn 24:123–140
- 20. Bunde E (2021) Ai-assisted and explainable hate speech detection for social media moderators-a design science approach. In: Proceedings of the 54th Hawaii international conference on system sciences, p 1264

- Cawley GC, Talbot NLC (2010) On over-fitting in model selection and subsequent selection bias in performance evaluation. J Mach Learn Res 11:2079–2107
- Chen BC, Kifer D, LeFevre K, Machanavajjhala A (2009) Privacypreserving data publishing. Found Trends Databases 2:1–167
- Cherdantseva Y, Hilton J (2013) A reference model of information assurance & security. In: International conference on availability, reliability and security, pp 546–555
- 24. Davenport TH (2006) Competing on analytics. Harv Bus Rev 84:98
- 25. Davis FD (1989) Perceived usefulness, perceived ease of use, and user acceptance of information technology. MIS Q. 319–340
- Delibasic B, Vukicevic M, Jovanovic M (2013) White-box decision tree algorithms: a pilot study on perceived usefulness, perceived ease of use, and perceived understanding. Int J Eng Educ 29:674–687
- Demirkan H, Delen D (2013) Leveraging the capabilities of service-oriented decision support systems: putting analytics and big data in cloud. Decis Support Syst 55:412–421
- Duncan G, Stokes L (2009) Data masking for disclosure limitation. Wiley Interdiscip Rev Comput Stat 1:83–92
- Duncan GT, Elliot M, Salazar-González JJ (2011) Statistical confidentiality: principles and practice. Int Stat Rev 80:479–480
- Dunkel J (2009) On complex event processing for sensor networks. In: Proceedings—2009 international symposium on autonomous decentralized systems. ISADS 2009:249–254
- Dwork C, Feldman V (2018) Privacy-preserving prediction. In: Conference on learning theory, PMLR 75:1693–1702
- 32. Džeroski S, Ženko B (2004) Is combining classifiers with stacking better than selecting the best one? Mach Learn 54:255–273
- Emerson RM (1976) Social exchange theory. Annu Rev Social arXiv:1011.1669v3
- 34. Fink O, Netland T, Feuerriegel S (2021) Artificial intelligence across company borders. Commun ACM 65:34–36
- Freund Y, Schapire RE (1996) Experiments with a new boosting algorithm. Int Conf Mach Learn 148–156. arXiv:978-0-387-09823-4_45
- 36. Gao J (2018) Security and privacy protection for eHealth data. Future network systems and security. In: Proceedings of the 4th international conference, FNSS 2018, p 197
- Goldfarb A, Tucker CE (2011) Privacy regulation and online advertising. Manag Sci 57:57–71
- Graepel T, Lauter K, Naehrig M (2012) Ml confidential: machine learning on encrypted data. In: International conference on information security and cryptology. Springer, pp 1–21
- 39. Gregor S (2006) Research essay: the nature of theory in information systems. MIS Q 30:611–642
- Gregor S, Jones D (2007) The anatomy of a design theory. J Assoc Inf Syst arXiv:1011.1669v3
- 41. Hakanen T, Jaakkola E (2012) Co-creating customer-focused solutions within business networks: a service perspective. J Serv Manag.
- 42. Hann IH, Hui KL, Lee SYT, Png IPL (2007) Overcoming Online information privacy concerns: an information-processing theory approach. J Manag Inf Syst 24:13–42
- 43. van Heck E, Vervest P (2007) Smart business networks: how the network wins. Commun ACM 33:3–8
- Hew KF, Huang W, Du J, Jia C (2021) Using chatbots in flipped learning online sessions: perceived usefulness and ease of use. In: International conference on blended learning. Springer, pp 164–175
- 45. Hirt R, Kühl N, Satzger G (2019) Cognitive computing for customer profiling: meta classification for gender prediction. Electron Mark 29:93–106
- 46. Hirt R, Kühl N, Schmitz B, Satzger G (2018) Towards service-oriented cognitive analytics for smart service systems. In:

Proceedings of the 51st Hawaii international conference on system sciences

- 47. Huhns MN, Singh MP (2005) Service-oriented computing: key concepts and principles. IEEE Internet Comput 9:75–81
- Jiang JC, Kantarci B, Oktug S, Soyata T (2020) Federated learning in smart city sensing: challenges and opportunities. Sensors 20:6230
- Kambil A, Short JE (1994) Electronic integration and business network redesign-a roles-linkage perspective. J Manag Inf Syst 10:15
- Kieseberg P, Schrittwieser S, Mulazzani M, Echizen I, Weippl E (2014) An algorithm for collusion-resistant anonymization and fingerprinting of sensitive microdata. Electron Mark 24:113–124
- 51. Kitchin R (2014) The data revolution: big data, open data, data infrastructures and their consequences. Sage, Thousand Oaks
- Kocabaş Ö, Soyata T (2016) Medical data analytics in the cloud using homomorphic encryption. In: E-health and telemedicine: concepts, methodologies, tools, and applications. IGI Global, pp 751–768
- 53. Koza JR, Bennett FH, Andre D, Keane MA (1996) Automated design of both the topology and sizing of analog electrical circuits using genetic programming. Artif Intell Des 96:151–170
- Kühl N, Hirt R, Baier L, Schmitz B, Satzger G (2021) How to conduct rigorous supervised machine learning in information systems research: the supervised machine learning report card. Commun Assoc Inf Syst 48:46
- 55. Kühl N, Schemmer M, Goutier M, Satzger G (2022) Artificial intelligence and machine learning. Electron Mark 32:2235–2244
- Lemke C, Budka M, Gabrys B (2015) Metalearning: a survey of trends and technologies. Artif Intell Rev 44:117–130
- Liu G, Zhu W, Saunders C, Gao F, Yu Y (2015) Real-time complex event processing and analytics for smart grid. Procedia Comput Sci 61:113–119
- Luckham D (2008) The power of events: an introduction to complex event processing in distributed enterprise systems. In: Lecture notes in computer science (including subseries lecture notes in artificial intelligence and lecture notes in bioinformatics). arXiv:9780.20139
- 59. Martin D, Kühl N (2019) Holistic system-analytics as an alternative to isolated sensor technology: a condition monitoring use case. In: Proceedings of the 52nd Hawaii international conference on system sciences
- 60. McCormack KP, Johnson WC (2016) Supply chain networks and business process orientation: advanced strategies and best practices. CRC Press, Boca Raton
- 61. Mitchell TM (1997) Does machine learning really work? AI Mag 18:11
- Narayanan A, Shmatikov V (2008) Robust de-anonymization of large sparse datasets. In: Proceedings—IEEE symposium on security and privacy. arXiv:0610105v2
- Pavlyshenko B (2016) Machine learning, linear and Bayesian models for logistic regression in failure detection problems. In: 2016 IEEE international conference on big data (big data), pp 2046–2050
- Pournaras E, Nikolic J (2017) On-demand self-adaptive data analytics in large-scale decentralized networks. In: 2017 IEEE 16th international symposium on network computing and applications, NCA 2017, pp 1–10
- 65. Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, Bakas S, Galtier MN, Landman BA, Maier-Hein K et al (2020) The future of digital health with federated learning. NPJ Digit Med 3:1–7
- Riquelme IP, Román S (2014) Is the influence of privacy and security on online trust the same for all type of consumers. Electron Mark 24:135–149

- 67. Robins DB (2010) Complex event processing. In: 2010 second international workshop on education technology and computer science 10
- Roman R, Lopez J, Mambo M (2018) Mobile edge computing, Fog et al.: A survey and analysis of security threats and challenges. Fut Gener Comput Syst arXiv:1602.00484
- Sadeghi AR, Wachsmann C, Waidner M (2015) Security and privacy challenges in industrial internet of things. In: Proceedings of the 52nd annual design automation conference on—DAC '15, pp 1–6
- Saeed KA, Abdinnour-Helm S (2008) Examining the effects of information system characteristics and perceived usefulness on post adoption usage of information systems. Inf Manag 45:376–386
- Saputra YM, Nguyen D, Dinh HT, Vu TX, Dutkiewicz E, Chatzinotas S (2020) Federated learning meets contract theory: economic-efficiency framework for electric vehicle networks. IEEE Trans Mobile Comput 21:2803–2817
- 72. Sarlis D, Papailiou N, Konstantinou I, Smaragdakis G, Koziris N (2015) SIGMOD '15- proceedings of the 2015 ACM SIGMOD international conference on management of data. SIGCOMM Comput Commun Rev 45:21–28
- 73. Satyanarayanan M (2017) The emergence of edge computing. Computer 50:30–39
- Schultz-Møller NP, Migliavacca M, Pietzuch P (2009) Distributed complex event processing with query rewriting. In: Proceedings of the third ACM international conference on distributed event-based systems—DEBS '09
- 75. Schüritz R, Satzger G (2016) Patterns of data-infused business model innovation. In: Proceedings—CBI 2016: 18th IEEE conference on business informatics
- Shi W, Dustdar S (2016) The promise of edge computing. Computer 49:78–81
- Sturm B, Sunyaev A (2019) Design principles for systematic search systems: a holistic synthesis of a rigorous multi-cycle design science research journey. Bus Inf Syst Eng 61:91–111
- Talia D (2013) Clouds for scalable big data analytics. Computer 46:98–101
- 79. Todorovski L, Džeroski S (2003) Combining classifiers with meta decision trees. Mach Learn 50:223–249
- Tuladhar A, Gill S, Ismail Z, Forkert ND, Initiative ADN et al (2020) Building machine learning models without sharing patient data: a simulation-based analysis of distributed learning by ensembling. J Biomed Inform 106:103424

- Uhlmann E, Laghmouchi A, Geisert C, Hohwieler E (2017) Decentralized data analytics for maintenance in industrie 4.0. Procedia Manuf 11:1120–1126
- 82. Venable J, Pries-Heje J, Baskerville R (2016) FEDS: a framework for evaluation in design science research. Eur J Inf Syst 25:77–89
- Vilalta R, Drissi Y (2002) A perspective view and survey of metalearning. Artif Intell Rev 18:77–95
- Walls JG, Widmeyer GR, El Sawy OA (1992) Building an information system design theory for vigilant EIS. Inf Syst Res 3:36–59
- 85. Wang W, Sung J, Kim D (2008) Complex event processing in EPC sensor network middleware for both RFID and WSN. In: Proceedings—11th IEEE symposium on object/component/service-oriented real-time distributed computing. ISORC 2008
- Wilson D, Ateniese G (2014) To share or not to share in clientside encrypted clouds. In: Information security: 17th international conference, ISC 2014, Hong Kong, China, October 12–14, 2014. Proceedings 17, pp 401–412
- Wixom BH, Schüritz R (2017) Creating customer value using analytics. CISR Res Brief 17:1–4
- Wohlgemuth S, Sackmann S, Sonehara N, Tjoa AM (2014) Security and privacy in business networking. Electron Mark 24:81–88
- Wolpert DH (1992) Stacked generalization. Neural Netw 5:241–259
- Xu J, Glicksberg BS, Su C, Walker P, Bian J, Wang F (2021) Federated learning for healthcare informatics. J Healthc Inf Res 5:1–19
- Yi S, Li C, Li Q (2015) A survey of fog computing: concepts, applications and issues. In: Proceedings of the 2015 workshop on mobile big data—Mobidata '15
- Zhang D, Xu B, Wood J (2016) Predict failures in production lines: a two-stage approach with clustering and supervised learning. In: Proceedings—2016 IEEE international conference on big data, big data 2016
- Zhang J (2011) Data use and access behavior in escience-exploring data practices in the new data-intensive science paradigm. Philadelphia, PA: Drexel University

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.