

# Extending semantic context analysis using machine learning services to process unstructured data

Anja Wilhelm<sup>1\*</sup>, and Wolfgang Ziegler<sup>1</sup>

<sup>1</sup>Karlsruhe University of Applied Sciences, Faculty of Information Management and Media, 76133 Karlsruhe, Germany

**Abstract.** The primary focus of technical communication (TC) in the past decade has been the system-assisted generation and utilization of standardized, structured, and classified content for dynamic output solutions. Nowadays, machine learning (ML) approaches offer a new opportunity to integrate unstructured data into existing knowledge bases without the need to manually organize information into topic-based content enriched with semantic metadata. To make the field of artificial intelligence (AI) more accessible for technical writers and content managers, cloud-based machine learning as a service (MLaaS) solutions provide a starting point for domain-specific ML modelling while unloading the modelling process from extensive coding, data processing and storage demands. Therefore, information architects can focus on information extraction tasks and on prospects to include pre-existing knowledge from other systems into the ML modelling process. In this paper, the capability and performance of a cloud-based ML service, IBM Watson, are analysed to assess their value for semantic context analysis. The ML model is based on a supervised learning method and features deep learning (DL) and natural language processing (NLP) techniques. The subject of the analysis is a corpus of scientific publications on the 2019 Coronavirus disease. The analysis focuses on information extractions regarding preventive measures and effects of the pandemic on healthcare workers.

## 1 Introduction and background

This paper illustrates the introduction of AI into the field of TC and examines the potentials of ML to gain insight into data and to enhance or replace manual feature extraction and classification tasks, to provide a deeper semantic analysis of large amounts of unstructured data.

The implementation of a ML model for semantic context analysis and insight into unstructured data shall be examined regarding its potential for automating and extending manual analysis processes.

For this purpose, a ML model with DL techniques was implemented using *Watson*, a cloud-based MLaaS algorithm by *IBM*. The domain-specific feature extraction was based on the semantic tagging of content according to a customized supervised learning model created in *IBM Watson Knowledge Studio (WKS)*. By applying DL and NLP techniques, an unstructured data corpus, consisting of scientific publications concerning the Novel Coronavirus Disease (Covid-19) gathered from public databases, was examined to provide insight into the contents of the individual corpus publications. The information classification and feature extraction focused mainly on information concerning healthcare personnel and protective measures in patient care. Subsequently, the customized model for the Covid-19 domain was deployed to the *Watson Natural Language Understanding (WNLU)* service to actively conduct feature extractions and to measure the success of a domain-specific model for

context analysis tasks in contrast to the performance of a standard model provided by *IBM Watson*.

The field of TC has advanced considerably in the past decades, progressing from standardized content creation towards increasingly complex information modelling. Today, information modelling does not only serve to ensure and maintain the quality of information but has grown towards a bigger vision: Starting from semantic metadata application to the modelling of ontologies and semantic networks and, nowadays, leading to first advances in the utilization of AI in TC. Content modelling and management increasingly strives to produce 'intelligent content', content that can be processed automatically, or that can at least be highly systematized. Intelligent information in TC relies heavily on the application of semantic metadata, enabling the targeted retrieval of relevant information for specific user needs. For instance, context sensitive information access can be achieved by utilizing metadata for the implementation of navigational, searching, and filtering options in content delivery portals (CDP). The introduction of ontologies provides an extension to the classification of metadata in taxonomies by modelling relations between metadata which are not limited to hierarchical relations. They provide more explicit semantics [1].

Intelligence in TC is dependent on granular, semantically structured content labelled with metadata. Those requirements fall within the scope of native intelligence. Native intelligence depends on manual

\* Email: [anja.wilhelm@hs-karlsruhe.de](mailto:anja.wilhelm@hs-karlsruhe.de)

knowledge input to provide a basis for the development of applications on a higher intelligence level. Augmented intelligence entered TC with the addition of relations between semantically created content in the form of ontologies and semantic networks. AI introduces the next step in the progression of TC [1]. In the context of technical information analysis, generation and delivery, AI applications can be applied for automated extraction of metadata and knowledge, for automated content classification, or even automated content generation.

## 2 ML modelling with IBM Watson

Companies specialized in AI development play a key role in shaping the future of AI applications. *IBM* is one of the leading companies that is commercially successful and has invested significantly in AI technologies [2]. *IBM*'s supercomputer *Watson* combines AI and analytical software for optimal performance in generating responses to digital prompts in natural language. It provides a semantic search engine for information retrieval [3].

*IBM Watson Cloud* offers services to train the *Watson* ML algorithm for the language of a specific domain to enhance semantic analysis tasks, such as the extraction of domain-specific information (concepts, entities, relations). *Watson* adapts to new inputs and learns through data analysis. It uses DL, as a ML technique, layering algorithms to create an artificial neural network. By utilizing DL, *Watson* improves its accuracy with increasing amounts of input data and can also learn from unstructured data. DL also enables *Watson*'s NLP. *Watson* learns by deconstructing sentences, analysing, and identifying concepts and relations within those sentences, and understanding how everything fits together to work out the context and intent of the input.

For business application, the algorithm also needs to understand the terminology and language of a specific industry (domain). With traditional AI, this task would demand for a multitude of additional data and computing power. *Watson*, however, applies a method called 'transfer learning'. With transfer learning the algorithm does not need to be trained from scratch but can be fed prior knowledge to speed up the process. This knowledge transfer is based on *Watson*'s three-layered AI model: The bottom layer is made from off-the-shelf general knowledge, e.g., consulting the *DBpedia* database for concept understanding [4]. The middle layer contains knowledge tailored to specific industries (domain knowledge) and defines how *Watson* understands specific terms. The top layer customizes the algorithm to specific business needs or use cases by building a domain-tailored model for domain and company specific terminology, language, and concepts [5].

*IBM WKS* is a cloud-based application to build such customized domain entity-relation models to train *Watson* in understanding linguistic nuances and in identifying domain-specific entities and relationships in unstructured text data whilst operating under a supervised learning method [5].

The workflow of the modelling process in *WKS* can be summarized into the following steps [5]:

- First, a data corpus fitting the domain of Covid-19 must be gathered to supply a training and test data set for the ML model.
- Second, the model assets to customize the model to the domain must be created.
- After this, the training data corpus must be manually annotated (semantically tagged) by the human annotator.
- Once a certain amount of data has been tagged manually, the annotation of the training data can be semi-automated with support from pre-annotation processes. The human annotator gradually changes roles from 'tagger' to a supervisor of the ML annotation.
- The completed tagging of the training data establishes the 'ground truth' of the ML model. Now, a version of the model can be trained and evaluated for improvements and for deployment to other *Watson* services or applications.

### 2.1 Data corpus collection

To create a manageable but also suitable corpus for the model training in the context of this research, the overall domain of Covid-19 had to be narrowed down to specific subtopics. Because of the absence of expert knowledge in highly medical content, the chosen subtopics dealt with parameters of the Covid-19 disease which were more relevant in the everyday experiences of the pandemic. Therefore, this model focused on the areas of (1) *Healthcare Personnel* and (2) *Preventive Measures*. The gathering of the corpus documents should also include overlaps in content and detectable entities to enable the extraction of semantic relations between entity types for the entity-relation design in *WKS*.

The following public data sets and databases for Covid-19 research and scientific publications served as the main sources to gather natural language training data with domain-specific terminology and phrasings:

- 'WHO Data COVID-19', a database compiled by *The World Health Organization*.
- 'LitCOVID', a database compiled by *The National Library of Medicine* from Covid-19 articles in *PubMed*.
- 'CORD-19', an open Covid-19 research data set compiled by *The Allen Institute for AI*.

### 2.2 Preparing model assets

Once the data corpus had been gathered, the model customization for the domain could be implemented by setting up the *WKS* model assets. The assets in *WKS* consist of documents, dictionaries, and the Type System [5].

**Documents** for *WKS* are uploaded from the previously gathered data corpus. They serve as the training data for annotation and for building a test set for model evaluation. They must meet specific requirements to be conforming to *WKS* requirements, such as the word

count, file format, and structure (only unstructured, natural language texts are supported and no tables or graphics). Therefore, some pre-processing (raw text extractions, document splitting) of the data corpus had to be conducted before the final upload to *WKS* [5].

**Dictionaries** in *WKS* define terms that must be considered equivalent by the model. They improve the understanding of the domain language and enable automated pre-annotation of domain terminology as entities [5] (e.g., containing terminological variations of the naming for “Novel Coronavirus Disease” and annotating them with the “Disease” entity type). They provide a quick and reliable way to pre-annotate training data if, for example, a terminology database, an ontology, a metadata taxonomy, or even a simple word list for the domain has already been designed in previous working procedures.

The **Type System** is the basis for generating semantic tags (labels) for model training in *WKS*. It defines all relevant entity and relation types of the domain. Entity types are used for the classification of entity mentions [5]. For example, the mention “Covid-19” can be annotated with the “Disease” entity type and the mention “face mask” with the “PPE” (personal protective equipment) type. Subtypes can be defined to further classify entity types by adding a hierarchical level. Relation types define interrelations between entities. Relations in *WKS* are binary and can only be assigned within the boundary of a sentence. They can be symmetrical or asymmetrical. When defining a symmetrical relation, both entity types must be defined as a source and target entity to enable the relation to be labelled in both directions. For asymmetrical relations, the order of the source and target entity type of the relation must be defined. The relation of the type “Provides Protection From” between the entity type “PPE” and the “Disease” type can only start from “PPE” and end at “Disease” (asymmetrical).

Entity and relation type definitions are exportable in JSON format and contain information, such as the assigned ID, the label, and subtypes. For relations, the export also contains the source and target entity types of the relation, referenced by their IDs.

The entity type definitions for the custom model development of this research were based on a metadata model which had been created during the corpus generation of the training data. For metadata extraction, affiliated metadata from the source databases (2.1) and the individual documents have been extracted and compared based on a frequency analysis on the main subjects and keywords associated with the articles. A manual extraction process added to the extracted metadata by sighting several documents and detecting domain-relevant metadata and context information for later relationship definitions between metadata.

The final metadata taxonomy consisted of a four-level hierarchy which then had to be logically translated into the two-level setup of entity types and subtypes in *WKS*. For this process, the desired granularity of the final analysis tasks was the decisive factor. For a feature extraction and context analysis on a more general level

(e.g., full text classification) the upper levels of the metadata provided these broader classifications for entity types. For a more detailed extraction of text fragments (e.g., on a sentence basis for extracting specific context information within full documents), the entity type definitions could be transferred from lower-level metadata.

**Table 1.** Three-level metadata taxonomy.

Meta Level 1	Meta Level 2	Meta Level 3
Prevention and Protection Measure	PPE	Facial Protection Body Protection
	Preventive Measure	Social Distance Quarantine Personal Hygiene

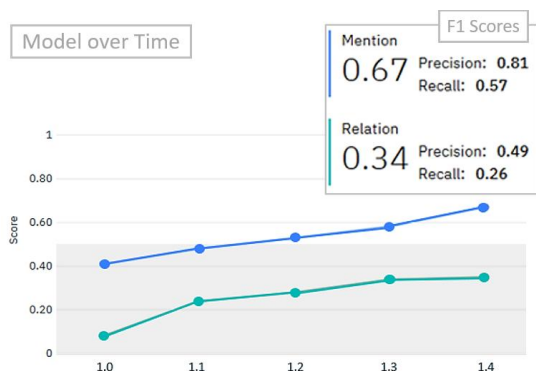
Take for example the metadata levels in Table 1: Depending on the desired granularity of the output they could have been transferred to the *WKS* Type System in different ways:

- a) Defining one entity type named “Prevention and Protection Measure” (level 1) with two subtypes “PPE” and “Preventive Measure” (level 2).
- b) Defining two entity types: “PPE” and “Preventive Measure” (level 2) with their respective subtypes “Facial Protection” and “Body Protection”, and “Social Distance”, “Quarantine”, and “Personal Hygiene” (level 3).

### 2.3 Model performance evaluation

The training of models in *WKS* is based on the annotation, the semantic labelling, of the documents with the type labels resulting from the Type System definitions.

Once at least ten documents have been annotated, a version of the model can be trained. Each training cycle of the model supplies statistics for model evaluation based on the performance of the model on the test set [5]. The performance statistics include f1 scores for entity and relation extractions, which are composed of the precision and recall values of the model analysis on the test set. Precision and recall are commonly used to evaluate classification and information retrieval systems [6]. Precision describes the fraction of relevant instances among all retrieved instances. Recall defines the fraction of retrieved instances among all relevant instances. The *WKS* model statistics provide overall scores on entity and relation extraction which are also visualized in a diagram showing the model progression over time regarding the ratio of the model version and their respective f1 scores for entity and relation extraction (Fig. 1). The *WKS* evaluation statistics also provide a breakdown on each individual entity and relation type in terms of precision and recall.



**Fig. 1.** Model development over time with the top graph visualizing the entity (mention) f1 score, and the bottom graph visualizing the relation f1 score.

The overall evaluation of the model training showed that the most frequently annotated types performed expectedly well compared to the types with fewer training events. However, scores of entity types could be drastically improved by creating dictionaries for ill-performing types regarding missed entities (recall) and confused, falsely classified types (precision).

The occurring errors during entity extractions could be divided into three categories based on their severity for the final analysis task:

- Missed annotations, where the entity was not extracted at all and therefore no information extraction would have been possible in a final application of the model.
- Falsely labelled entities, where the entity mention was correctly identified in the test set but labelled with the wrong entity type, leading to false information extractions.
- Cases where the extracted entity was not an exact match on a token basis, e.g., the mention “strict home isolation” of the “Preventive Measure” type was annotated in the test set as “home isolation”, missing the token “strict”. This error type could be considered as a minor error in comparison to error types (a) and (b), as it would still have provided a mostly correct information extraction.

Errors during the machine labelling of entities on the test set frequently occurred in cases of annotated phrases and compound mentions, whereas single word entities performed above average. This was most likely due to the model’s need for a larger amount of training data to reliably identify more complex linguistic patterns based on the applied techniques of DL and NLP [7].

### 3 Deploying the model to Watson NLU for context analysis

Once the custom domain model had been trained in *WKS*, it could be deployed to *WNLU* for feature extraction and context analysis [5].

*WNLU* is a service focused on the analysis of unstructured content by enabling the extraction of text analysis features. *WNLU* provides a text analysis for 13 languages, though some extended features are currently only available for English documents [8].

### 3.1 Analysis features

Out of the possible text analysis features of the *WNLU* service, this research focused on the extraction of entities and relations, the features which were customizable by the *WKS* model, and on concepts and categories.

The **concepts** feature returns high-level concepts of the analysed text. This way, a research paper on “deep learning” might return the concept “artificial intelligence” even though the term may not have been directly mentioned in the paper. Concepts are identified by *WNLU* by consulting *DBpedia*, a semantic web with linked data from *Wikipedia* articles [4].

The **categories** feature is based on *Watson*’s default taxonomy, a five-level hierarchy for categories, such as “Automotive and Vehicles”, “Education”, “Finance”, “Health and Fitness”, and “Science” [8].

**Entity and relation types** can either be defined by a deployed *WKS* custom model or by the Type System of *Watson*’s standard model for the English language. In both cases, *Watson* conducts its analysis based on pre-defined entity and relation types [8].

Depending on the extracted feature types, *WNLU* assigns relevance and/or confidence scores to the analysis output (Fig. 2). The relevance score assesses the importance of the feature within the source text. The confidence score (for entities only) indicates how secure the intelligent tagging is in the type it assigned to the entity [8].

```

"type": "Organization",
"text": "CDC",
"relevance": 0.9711,
1 "confidence": 1.0,
-----
2 "type": "Disease",
"text": "COVID-19",
"relevance": 0.683086,
"confidence": 1.0,

```

**Fig. 2.** Excerpt of the analysis results of the entity feature extraction with the custom model: Extracted entities “CDC” of the “Organization” type, and “COVID-19” of the “Disease” type.

### 3.2 Performance comparison of the standard model and the custom model

To evaluate the benefits of developing a custom domain model, a comparative analysis on domain-specific source documents was conducted. The first analysis applied the *WNLU* standard model for feature extraction, the second analysis applied the custom model on the same source documents.

The documents which were object to this comparative analysis contained information on the topic of supply shortages of personal protective equipment for healthcare workers in the U.S. during the Coronavirus pandemic.

#### 3.2.1 Category and concept extraction with the *Watson* standard model

Category extractions with *WNLU* provide explanations on the categorization reasoning of the model in the form of extracted keywords from the source document. The model

extracted categories, such as “Health and Fitness” with the subcategories “Men’s Health” and “Education/Teaching and Classroom Resources”. These results showed that *Watson*’s five-level taxonomy did not include a category specifically for medical content. It categorised the source documents as “Men’s Health” based on keywords, such as “public health” and “healthcare system”, and as “Teaching and Classroom Resources” because of extracted keywords, such as “supplies” and “shortages” in the source document, which actually referenced shortages in the supply of PPE for healthcare workers, not in school supplies. Unless the category taxonomy of *Watson* will be adaptable to specific content-domains in the future, it is likely that an extraction of categories will not provide valuable insights into highly specific domain data.

The extraction of concepts, however, promised a new approach in identifying a variety of domain-relevant terminology and possibly entity types as they utilized a wide semantic web for concept extractions. The analysis results included concepts, such as “personal protective equipment”, “health care”, public health” and “protection” as associated concepts of the source documents. “Personal protective equipment” was already a defined entity type in the custom model. These results revealed the potential to utilize the *Watson* standard model on unstructured data to assist the identification of entity types in the early stages of the customization process by analysing training data documents automatically in *WNLU*.

### 3.2.2 Comparing entity and relation extractions of both models

Even though the *WNLU* Type System for the English language supports the extraction of a variety of entity and relation types, it is not specialized in one specific domain but rather offers a broad coverage of a multitude of subjects. The standard model correctly recognized and classified entities of types, such as “Organization” (CDC), and “Location” (U.S.) with high confidence and relevance scores but failed to extract any domain-specific entity types. Furthermore, it showed an error-proneness for entities which were mentioned in abbreviated forms, such as “HCP” (healthcare personnel) and falsely classified them as organizations, locations, and facilities.

The recall for relation instances was overall rather low. It recognized only two generalized relations of the type “Part of Many” between entities of the type “Person”, e.g., “patients” (source “Person” entity) is “Part of Many” “others” (target “Person” entity). The relations could be considered correct as they did not provide false information, but the algorithm failed to identify more domain-relevant relations between entities.

The entity and relation extraction with the custom model showed rather different results. The overall recall, the number of extracted entities and relations, was much higher than the recall of the standard model. It extracted all the entities identified by the standard model and classified them correctly. In addition, it extracted several domain-specific entities, such as “Covid-19” of the

“Disease” type (Fig. 2) and “personal protective equipment” of the “PPE” type. The precision of the entity classification, as well as the higher recall of domain-specific entities, showed a clear advantage of the application of a customized model for analysis purposes.

The relation extraction of the custom model provided far more domain-specific results than the standard model analysis. Due to the specialized entity types of the domain-specific model, it was able to extract relations, such as “Provides Protection For” between “personal protective equipment” (source “PPE” entity) and “healthcare personnel” (target “Person” entity).

Due to the Covid-19-focused terminology in the source documents for this analysis, the standard model was not expected to perform similarly insightful entity and relation extractions. On the other hand, the extraction of concepts with the standard model could provide a new angle in identifying possible entity types for expanding or improving the Type System of a domain model. Overall, the custom model clearly outperformed the standard model regarding these domain-specific context analysis tasks.

## 4 Conclusion and outlook

One of the biggest advantages of using MLaaS solutions, like *IBM Watson* services, for AI-related tasks is that the responsible person for creating and training the ML algorithm does not have to possess a deep knowledge in coding and computer linguistics. This way, companies can employ personnel for ML creation processes whose field of expertise is more heavily focused on use case analysis, domain knowledge representation, and information retrieval and delivery. Cloud-based solutions also provide the necessary computing power and data storage capacities to utilize DL methods.

The tailoring of *Watson* for specific domain content, via a supervised learning method, plays to the potential of integrating expert domain knowledge into the model training process to attain better performances.

The results of the research on implementing a custom model showed that the deployed version of a domain-specifically trained model provided deeper data insights through more specific feature extractions during analysis processes than relying on a more generalized model by AI providers (3.2).

Directly importing pre-existing information structures, models, or other semantic information into MLaaS systems can be a challenging task due to the often-missing insights into the workings of the algorithm, or due to provider-specific data properties, such as *IBM Watson*’s system-generated and assigned IDs for entities and relations. Nevertheless, knowledge graphs, taxonomies and ontologies can serve as templates for the modelling of entities and relations for semantic tagging of the ML training data, even though some adjustments must be made to the modelling depth of hierarchies to fit them to the limited modelling depth of *Watson*’s custom models.

Overall, *IBM Watson* provides a quick and intuitive entry into ML modelling and the training of the algorithm for the understanding of a specific domain language. It

presents a starting point to get familiar with the workings and concepts of AI technologies regarding context analysis, and to gain insight into the demands on training data and NLP-based entity and relation extraction processes. But, once the basics have been established, *IBM Watson* does provide little scope to adjust the analysis for individual needs (e.g., modelling depth/granularity, integration of knowledge from different providers). Additionally, restricted insight due to the '*Watson Blackbox*', resulting from the nondisclosure of the algorithm, can make it difficult to identify areas of improvement for the model, aside from adding more training data.

The introduction of AI into the field of TC is still in its early stages. AI-assisted content creation, in the form of grammar and syntax checkers, and neural machine translations, provides promising opportunities and challenges for linguistic areas of TC. The fields of machine translation and language checkers shape up well for a wider integration of DL technologies as they show high learning capabilities regarding NLU [9]. Though, to make such systems more employable for the field of highly technical content, the adaptability for specific domains, text types, and terminology integration must be ensured in the progress of these technologies [10].

Information extraction by ML applications in the form of terminology, metadata, and entity extraction introduces a new approach to rival statistical methods. DL algorithms can extract more context-sensitive semantic data for further tasks, such as the building of domain classifiers, and to improve NLU by tackling ambiguity and colloquialism problems via deep text analytics methods [11].

Content delivery provides a point of intersection where AI can support processes, by means such as automated rule-writing instead of purely hand-crafting rules for content delivery. Systems like automated ontologies and delivery solutions such as semantic correlation rules (SCR) [12] could aim for the inclusion of ML to extract rules from pre-existing information models, statistics, and analytics to ease the workload of information architects.

NLG in TC faces the challenge of a high demand for accuracy to guarantee the safe guidance of users in their tasks and to meet legal regulations, standards, and norms. Because of the high demand for correct information, rule-based solutions should always be considered alongside or instead of AI solutions in the decision-making process of introducing an improved system. At the current state of research, rule-based approaches can show a higher degree of accuracy in their performance, depending on the expertise and precision incorporated into the rule-writing process. They can also operate on a much smaller data corpus and are more easily generalized than ML models. To decide whether an AI-based or a rule-based approach is more fitting for a specific use case, the project parameters regarding available data, computing power and demand for accuracy must be analysed in detail.

We would like to give thanks to the Faculty of Information Management and Media, and the Karlsruhe

University of Applied Sciences, Germany, for financially supporting this conference participation and publication.

## References

1. W. Ziegler, *Drivers of Digital Information Services: Intelligent Information Architecture in Technical Communication*, in ACM Chapter Proceedings on Educational Technology, Language & Technical Communication, ACM 2019, pp. 51-52 (2019)
2. A. Patrizio, J. Maguire, *Top 100 Artificial Intelligence Companies 2020*, <https://www.datamation.com/artificial-intelligence/top-100-artificial-intelligence-companies-2020/> (2020)
3. E. Burns, *IBM Watson Supercomputer*, <https://searchenterpriseai.techtarget.com/definition/IBM-Watson-supercomputer> (2020)
4. A. Stevens, *Wikipedia Factoid Bot: Extract Answers from DBpedia Using SPARQL*, <https://medium.com/ibm-watson/wikipedia-factoid-bot-5-of-6-extract-answers-from-dbpedia-using-sparql-eab1b268bfa> (2017)
5. IBM Cloud, *Watson Knowledge Studio Documentation*, <https://cloud.ibm.com/docs/watson-knowledge-studio> (2020)
6. C. Goutte, E. Gaussier, *A probabilistic interpretation of precision, recall and F-score, with implication for evaluation*, in European conference on information retrieval, Berlin, Heidelberg, pp. 345-359 (2015)
7. T.M. Georgescu, B. Iancu, M. Zurini, *Named-entity-recognition-based automated system for diagnosing cybersecurity situations in IoT networks*, in Sensors, **15**, 3380 (2019)
8. IBM Cloud, *API Docs: Natural Language Understanding*, <https://cloud.ibm.com/apidocs/natural-language-understanding> (2020)
9. C. Larssonneur, *Neural Machine Translation: From Commodity to Commons*, in When Translation Goes Digital, Palgrave Macmillan, Cham, pp. 257-280 (2020)
10. U. Reuther, P. Schmidt, *Schon nahe dran*, in technische kommunikation, **06/18**, pp. 18-25 (2018)
11. S. Mahapatra, *Why Deep Learning over Traditional Machine Learning*, <https://towardsdatascience.com/why-deep-learning-is-needed-over-traditional-machine-learning-1b6a99177063> (2018)
12. W. Ziegler, *Semantic Correlation Rules as a Logic Layer between Content Management and Content Delivery*, Proceedings of the ETLTC2021 ACM Chapter International Conference on Educational Technology, Language & Technical Communication. Aizuwakamatsu, Japan. (2021)