

# Grade Level Filtering for Learning Object Search using Entity Linking

Ratan J. Sebastian<sup>1,\*</sup>, Ralph Ewerth<sup>1,2</sup> and Anett Hoppe<sup>1,2</sup>

<sup>1</sup>TIB Leibniz Information Centre for Science and Technology

<sup>2</sup>L3S Research Center, Leibniz Universität Hannover

## Abstract

More and more Learning Objects like lessons, exercises, worksheets and lesson plans are available online. Finding them, however, is a challenge as they often lack metadata concerning format, content and, in the K-12 context: grade-levels or age ranges for which they are appropriate. This work studies the automatic content-based assignment of this last aspect of Learning Object metadata. For this purpose, we (a) collected a dataset of physics lessons, (b) explored a set of text-based features for their automatic analysis (derived from both dense vector representations and entity linking methods) and (c) trained a machine learning model with different subsets of these features to predict a resource's target grade level. We compare and discuss the results.

## Keywords

information retrieval, learning object, search, metadata enrichment, machine learning, classification

## 1. Introduction

Learning Objects are defined by the IEEE Learning Object Metadata standard [1] as "[entities] which can be used, re-used or referenced during technology supported learning". These include textual or multimedia instructional content, worksheets, software tutors, lesson plans, etc. The quantity of these objects that are available online and the variety of sources that produce them lead to challenging search problems for teachers and students who search for them to augment existing learning material. Learning Object Repositories like OERCommons<sup>1</sup>, Elixier<sup>2</sup> and Wir Lernen Online<sup>3</sup> store Learning Objects and provide keyword and faceted search to find them.

Faceted search, however, is only as good as the metadata accompanying the Learning Object [2]. Common facets include subject domain, resource type, source and grade level [3]. Studies find that the available metadata are often incomplete and lack information in these fields [4] [5]. Consequently, several researchers explore the automatic enrichment of available metadata, focusing on quality estimation [5], error correction [6] and extracting tags or improving textual

---

*Joint Proceedings of the 10th International Workshop on News Recommendation and Analytics (INRA'22) and the Third International Workshop on Investigating Learning During Web Search (IWILDS'22) co-located with the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'22), July 15, 2022, Madrid, Spain*

\*Corresponding author.

✉ Ratan.Sebastian@tib.eu (R. J. Sebastian); Ralph.Ewerth@tib.eu (R. Ewerth); Anett.Hoppe@tib.eu (A. Hoppe)

🆔 0000-0002-8034-3382 (R. J. Sebastian); 0000-0003-0918-6297 (R. Ewerth); 0000-0002-1452-9509 (A. Hoppe)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

<sup>1</sup><https://www.oercommons.org/>

<sup>2</sup><https://www.bildungsserver.de/elixier/>

<sup>3</sup><https://wirlernenonline.de/>

metadata fields using knowledge graphs [7, 8].

In this study we address one of the metadata fields commonly absent in and very important to Learning Object search: grade level i.e. which school grade a certain object is appropriate for. We found that only 32% of Physics related material on Elixier (a central referatory run by the German federal government) have grade level information. For an initial study, we limit ourselves to Learning Objects that provide some sort of textual content. We acquire a dataset by scraping [LeifiPhysik.de](http://LeifiPhysik.de), a popular learning resource for German school students which provides textual instructional content covering a large portion of the physics syllabus for each grade.

The dataset has 539 lessons labelled with a grade level for each state. We experiment with different types of features, machine learning algorithms and machine learning problem formulations to test their relative performance. We discuss in detail the features provided by entity linking with a focus on their interpretability.

In summary, the contributions of this paper are three-fold:

1. Introduce the grade level prediction task with the goal of enriching Learning Object metadata with a grade level.
2. Use entity linking to generate interpretable and filterable features and examine their effectiveness.
3. Report empirical results on the accuracy of text-based grade level prediction using a dataset collected from [LeifiPhysik.de](http://LeifiPhysik.de)

## 2. Related Work

### 2.1. Learning Object Search

The lack of effective search systems have been identified as a barrier to teachers finding good Learning Objects online [2][9]. Anderson et al. [2] identify specifically the "poor incorporation of educational standards and learning outcomes into metadata" as a reason for existing solutions being ineffective. A study in the German context [9] finds that most teachers use the search functions in Learning Object Repositories with keywords but often find this insufficient and fall back to using Google. A 2018 survey of the field [10] identifies several search systems that can use metadata like educational standards or grade level for search but these systems are limited by missing metadata. Ochoa et al. (2008) [11] chart a history of learning object search in which metadata-enabled, precision-favouring faceted search gives way to full-text, keyword search which favours recall. They point out pros and cons to both approaches and propose a learning to rank approach that addresses the latter but do not discount the effectiveness of complete metadata to deliver on the precision of results promised by the former.

### 2.2. Metadata Enrichment

Many learning material platforms allow authors and curators to describe material with metadata. Typical fields used are: content type, subject area, source, grade level and tags. For learning

objects, there are a number of commonly used metadata standards such as IEEE LOM, ADL SCORM and LRMI [6].

Ochoa et al. (2009) [5] identify metadata quality in Learning Object Repositories to be a problem and propose quality measures for them. A quantitative analysis [12] of all resources annotated with a popular metadata standard, LRMI, found that most repositories that use this standard do not provide fields like `educationalAlignment`, `alignmentType` or `typicalAgeRange` all of which indicate the age bracket or grade level a resource is appropriate for.

To address the problem of missing metadata, Ochoa et al. (2005) [13] propose a framework for creating Learning Object metadata from course metadata. However, the problem of inferring metadata values in cases when no clear source exists in the course metadata is out of the paper's scope. Classifying Learning Objects using Machine Learning was explored by Hassan et al. [14]. However, the study aimed to predict a Learning Object's "educational value". The authors find this can be reliably assigned using ML methods. Other work that tries to solve the missing metadata problem focuses on linking to knowledge graphs to improve keyword search [7], [8] or error correction based on existing metadata [6].

All of these works either focus on enriching metadata fields other than grade level or propose general methods that might infer a grade level but a.) do not study its performance in detail and b.) do not take advantage of all the sources of information available in a learning object to perform this inference.

### 2.3. Entity Linking

Entity Linking has been used to help teachers author learning material in [15]. The authors use DBpedia Spotlight [16] to link fragments of text to DBpedia Resources and in turn use those Resources to recommend similar content from other textbooks to teachers. Similarly, Alpizar-Chacon et al. [17] use DBpedia to link textbooks to each other. Entity linking has been used in other domains [18] as a component in a larger machine learning system for document retrieval and classification.

## 3. Study Design

### 3.1. Problem Definition

In grade level prediction, the input is formatted textual content and the output is an ordinal variable or a list of ordinal variables that indicate which grade level(s) this text is suited for. The grade level is an integer that ranges from the lowest to the highest grade applicable for the educational context.

In this work, we restrict the context to Physics Education in grades 5 and up in all 16 German states. The states vary widely in their curricula. Lessons are taught at different grades in different states. Some lessons are taught across more grades in some states as opposed to others. For instance a lesson on string vibration <sup>4</sup> is taught at grades 5,6,9 and 10 in Berlin but only at grade 8 in Hessen.

---

<sup>4</sup><https://www.leifiphysik.de/akustik/akustische-phaenomene/grundwissen/saitenschwingung>

### 3.2. Data Collection

Data is collected by crawling the LeifiPhysik website. The website organises lessons by state and then by grade level for that state. These two attributes as well as the HTML content of the page are collected during the crawl. There are 539 lessons with 11,667 annotations.

The main text of each web page is extracted using Trafilatura [19], a software package intended to create text corpora from web crawls. Equations are stripped out using a simple regex, as the presented analysis is limited to the textual content of the lessons. Of the remaining lessons, only those with at least 25 words are included in the dataset. This leaves 515 lessons with a median of 181 words and a max of 1,227. Per state, the number of lessons ranges between 308 and 474. Further details about the data broken down by state and grade level are available in Tables 1 and 2.

**Table 1**

Number of grade level annotations per state

State	Annotation Count
Berlin	1163
Baden-Wuerttemberg	1046
Hamburg	1032
North Rhine-Westphalia	914
Brandenburg	877
Thuringia	773
Bremen	754
Lower Saxony	708
Rheinland-Palatinate	673
Saxony-Anhalt	656
Saxony	585
Hesse	525
Schleswig-Holstein	513
Bavaria	506
Saarland	482
Mecklenburg-Vorpommern	460

### 3.3. Entity Linking

Entity linking is used in order to a.) filter out entities based on their Wikipedia categories and b.) get interpretable features.

Two different entity linkers explored in this study:

1. **Wikifier** [20] is a web service<sup>5</sup> that annotates a document with relevant Wikipedia concepts. Wikifier uses PageRank [21] to find a coherent set of inter-related concepts to annotate the document. Entities returned have a PageRank score which indicates how central they are in the network of links between concepts. Default values are used for all

---

<sup>5</sup><https://wikifier.org/>

**Table 2**

Mean and total number of annotations (across all states) per grade level.

Grade Level	Mean	Total
Grade 5	29	471
Grade 6	41	665
Grade 7	89	1426
Grade 8	97	1567
Grade 9	96	1539
Grade 10	108	1742
Grade 11	114	1835
Grade 12	125	2000
Grade 13	25	415

**Table 3**

Physics-related DBpedia categories that are used for entity filtering. The suffix dbc refers to the namespace <http://DBpedia.org/resource/Category>

dbc:Subfields_of_physics	dbc:Physical_quantities	dbc:Physics
dbc:Concepts_in_physics	dbc:Classical_mechanics	dbc:Metrology
dbc:Physical_sciences	dbc:Engineering_disciplines	dbc:Universe
dbc:Electrical_engineering	dbc:Applied_sciences	dbc:Electromagnetism

parameters except language which is set to German. Various settings of the PageRank threshold are experimented with and an optimal one is chosen (details in Section 4.2).

2. **Babelfy** [22] is a web service <sup>6</sup> that similarly annotates text with entities in the BabelNet knowledge base which in turn – when possible – link to DBpedia Resources. It focuses on multi-lingual word disambiguation and was chosen since it is able to link more German words to entities than Wikifier.

In the context of the LeifiPhysik dataset, relevant entities are ones that are linked to one of the physics related categories in Table 3. The link could be a direct one between the resource and the category through the <http://purl.org/dc/terms/subject> property. Or it could be an indirect link from one of those categories to an ancestor category through a chain (up to 5 links long) of `skos:broader` properties. `skos:broader` connects a category to a broader parent category.

### 3.4. Modelling

Lesson documents are first encoded into a vector representation. The following vector representations are used:

---

<sup>6</sup><https://babelfy.org/>

### 3.4.1. Features

1. **TF-IDF**: Term Frequency-Inverse Document Frequency generated using sklearn's <sup>7</sup> TfidfVectorizer and the standard German tokenizer and stop word list.
2. **SentenceTransformers (SBERT)**: Dense vector representation generated using SBERT<sup>8</sup>.
3. **YAKE**: Keywords are extracted from documents using YAKE [23], a one hot encoding of which is used to represent the document. YAKE only uses syntactic features to identify keywords and does not rely on an external knowledge base. It is used as a baseline to compare the entity linker features to.
4. **Babelify & Wikifier**: The services mentioned in section 3.3 are used to perform entity linking and then the linked entities for each document are encoded into vectors using a one hot encoding.

### 3.4.2. Experimental Settings

The XGBoost [24] library was used for all model training. XGBoost was used since performance was comparable with or better than other modelling methods like Random Forests and Support Vector Machines (see Table 4) and the focus of this paper is on examining the impact of various features and formulations rather than ML algorithms. Hyperparameter tuning is employed in each case to decide an optimal `n_depth` hyperparameter for the model. In all experiments, the train-validation-test split is 66.66%-16.66%-16.66%. Hyperparameter tuning is done on the validation set and accuracies are reported on the test set. All experiments were run with 3-fold cross validation and results were averaged across folds.

## 3.5. Experiments

### 3.5.1. Feature Comparison

In order to see which features capture information relevant to the grade level prediction task, the performances of the feature sets described in Section 3.4.1 are compared. Results are shown

<sup>7</sup><https://scikit-learn.org/stable/index.html>

<sup>8</sup><https://www.sbert.net/>

**Table 4**

Accuracy of various models on SBERT, an entity linking feature (Babelify) and TF-IDF in the Categorical formulation. We see that choice of model does not improve performance significantly and in the case of SVM is worse.

Model	SBERT	Babelify	TF-IDF
XGBoost	0.249	0.235	0.202
SVM	0.241	0.205	0.186
Random Forest	0.250	0.233	0.191

in Table 5. The specific evaluation metric used is described based on the problem formulation discussed in Section 3.5.2.

### 3.5.2. Problem formulation comparison

Since the grade level prediction task can be formulated as different machine learning tasks, we explore the mapping from grade level prediction task to machine learning task and the trade-offs of each choice. We then compare them in terms of performance to see if one formulations models the problem better than any other. The machine learning formulations considered are:

- **Categorical:** This formulation treats grade level prediction as a single-label categorical classification. Each of the grade levels is treated as a separate category with no ordering relationship. Since each document can have only one label in this formulation, lessons that are assigned to more than one grade are filtered out. Since this formulation requires omission of data and does not use ordering constraints on the classes, it is intended as a baseline which which the ordinal and multi-label formulations can be compared.
- **Ordinal:** This formulation takes into account the ordering relationship between grades. A series  $N - 1$  of categorical binary classifiers are combined to predict  $N$  ordinal variables as outlined in [25].
- **Multi-Label:** This formulation accommodates the fact that lessons can be assigned to multiple grades. It treats grade level prediction as a multi-label categorical classification problem by training one classifier for each label using SKLearn's `MultiOutputClassifier`<sup>9</sup>.
- **Coarse-grained:** Practical deployment of grade level classification could benefit from higher accuracy at the cost of a coarser definition of grade-level. Hence in this formulation, 5<sup>th</sup> grade is considered Primary, 6<sup>th</sup> - 10<sup>th</sup> is considered Secondary-I and 11<sup>th</sup>-13<sup>th</sup> is Secondary-II making this a 3-label categorical classification problem rather than a 9-label one.
- **W/ states:** In all previous settings the input was just the vector encodings of the lessons. A model is trained separately for each state and the average performance across states is reported. Here, the state is encoded as part of the input using a one hot vector and the problem is treated as a multi-label categorical classification.

The performance metric in all formulations is class-weighted accuracy. This rewards higher accuracy for grades that appear more frequently in the data. This is chosen for two reasons: (a) we found that accuracy is correlated with the number of annotations available for a class (See Section 4.5) and (b) The focus in this experiment is on comparing various feature sets and formulations. Hence, it did not make sense to penalize an experiment setting for a lack of data.

---

<sup>9</sup><https://scikit-learn.org/stable/modules/generated/sklearn.multioutput.MultiOutputClassifier.html>

**Table 5**

Accuracy scores for the best performing hyperparameter-tuned models with each feature set and each modelling setting.

Features	Ordinal	Categorical	Multi-label	Coarse-grain	W/ states
SBERT	0.265	0.249	0.289	0.557	0.112
Babelfy	0.235	0.218	0.240	0.542	0.100
Wikifier	0.242	0.226	0.130	0.536	0.118
YAKE	0.211	0.194	0.209	0.507	0.104
TF-IDF	0.202	0.199	0.222	0.522	0.104

**Table 6**

Comparison of most distinguishing features between YAKE Keyword Extraction and Babelfy and Wikifier entity linking. Note: Yake and Babelfy terms have been translated from German for comprehension.

	YAKE	Wikifier	Babelfy
0	basic knowledge	Thermodynamic_free_energy	Electron
1	look	Force	Visual Perception
2	Energy	Fundamental_science	Electrical Current
3	Body	Animation	Magnetism
4	Light	Light	Energy
5	Fig.	Physical_body	Heat Conduction
6	important	Motion_(physics)	Electromagnetic Induction
7	Power	Particle	Light
8	Movement	Magnetic_field	Moon
9	Nature	Water	Atom

### 3.5.3. Feature Interpretability

In order to see what kinds of features the model uses to distinguish between grades, the 10 most important features as learnt by the XGBoost model are extracted and reported for the YAKE, Babelfy and Wikifier feature sets in Table 6. Further, the occurrence of these features at various grade levels is plotted in Figure 2 to make it possible to manually verify that the features that the model considers important are indeed key concepts taught at certain grade levels.

### 3.5.4. Entity Filtering

The confidence scores attached to linked entities are in some sense a measure of the centrality of the annotated entity in the graph of candidate entities for a given document. Since centrality might not directly relate to utility in grade level prediction we would like to see which confidence score threshold yields the best grade level prediction accuracy. The relation between these two quantities is plotted in Figure 1.



## 4. Results and Discussion

### 4.1. Feature Comparison

We see from Table 5 that Sentence Transformers consistently outperform entity linking features. Sentence transformers capture signals like authorial style and a broader array of content non-specific complexity features [26]. There is even some evidence to support the idea that entity relationships from open knowledge bases are also captured by these embeddings [27]; relationships that we hoped to capture and exploit by filtering entity linking features by domain. This result shows that entity linking and domain-adjusted filtering does not capture and boost enough information about domain terms in lessons to outweigh the general language model understanding of the same sentences in these lessons.

SBERT features still have the problem of interpretability and future work incorporating more features of related entities such as the number of relations between entities [28] and prerequisite relationships between them [29] could improve their performance while helping us construct a more understandable mapping from concepts to grade level.

However, the entity linked feature sets (Babelify and Wikifier) perform better than the TF-IDF or YAKE baselines for most problem formulations. TF-IDF uses all tokens as features and ranks them based on frequency and YAKE finds keywords based on frequency, sentence and document structure. The entity linking features represent an additional source of information by finding DBpedia nodes that can be linked to those keywords and ranking them by coherence to the main topic of the lesson. Further it allows entities to be filtered by domain. This additional step has an overall positive but complicated effect on grade level prediction performance that is discussed in the next section.

### 4.2. Entity Linking & Filtering


It is clear from Table 7 that the entity filtering step described above helps with overall performance. This is expected since it strips out noisy features that the model would otherwise have to deal with.

**Table 7**

Effect of filtering out non-physics related entities. Note: Results reported are from the Ordinal formulation.

Filtering	Babelify	Wikifier
All	0.230	0.232
Physics	0.235	0.241

However, using thresholds on confidence scores reported by the entity linkers to filter out incoherent entities works less well. Figure 1 shows that prediction accuracy largely decreases in the case of both Wikifier and Babelify as the threshold is increased. Optimal values found were 0.01 for the PageRank of Wikifier and 0.7 for the confidence score reported by Babelify.



**Figure 1:** Effect of keeping only physics related DBpedia Resources at different Entity Linking confidence thresholds.

### 4.3. Comparison of interpretable features

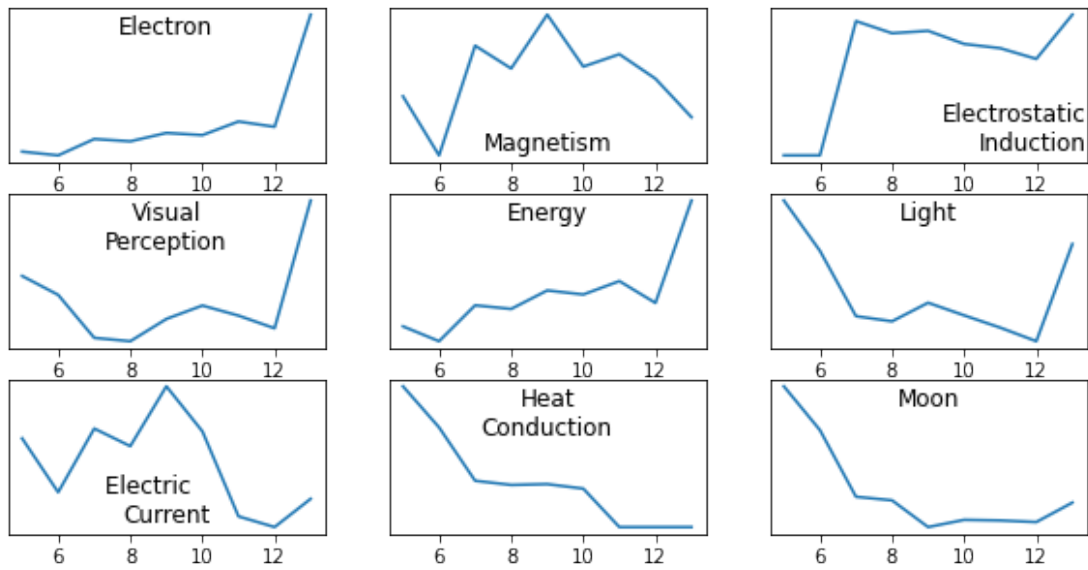
The best performing XGBoost model for each of the feature sets was queried for its most important features. The ranks of the features across each of the cross-validation folds were averaged and the top 10 features for each set are shown in Table 6.

- YAKE shows features unrelated to physics concepts such as "basic knowledge" and "Fig." (from figure titles). While Babelfy and Wikifier yield top features that are obviously more physics related even though these features were retrieved from a feature set that did not filter for domain specific features.
- Between Babelfy and Wikifier, there are more specific physics concepts found from the Babelfy features. This could be explained by Babelfy performing better on multilingual entity link benchmarks [22].
- Figure 2 shows how frequently different features show up at different grade levels. Visualization such as these could help domain experts verify that the model is learning useful distinctions and improve confidence in such a metadata enrichment step. For instance, the figure shows that the model thinks that Electrons are discussed at higher grade levels and Heat Conduction is discussed earlier which can be easily verified by an expert.

### 4.4. Problem formulation comparison

**Ordinal classification** performs better than the baseline categorical as expected since it takes into account the ordering of classes.

**Multi-label classification** shows very interesting results. There is a wide divergence in the performance of Babelfy and Wikifier which is not the case in any of the other formulations. Wikifier performs the worst in fact. This could reflect the fact that Babelfy performs better at German entity linking than Wikifier [22].



**Figure 2:** Prevalence of terms at different grades among documents in the test set. The x-axis of each graph represents the grades levels and the y-axis shows the prevalence of each entity at that grade. The features charted are the top 9 Wikifier features from Table 6

The use of **coarse-grained** categories unsurprisingly leads to the best overall accuracy and might be the most viable solution in a practical metadata enrichment system. A more detailed discussion about this model’s fitness for purpose follows in Section 4.5.

The **W/ states** formulation shows considerably worse performance and leads us to believe that the interactions between the content based features and the state encoding features are complex and are best not handled within a machine learning solution but rather separately.

#### 4.5. Error Analysis

**Table 8**

Detailed performance of model using SBERT features in Ordinal formulation broken down by grade level.

	Precision	Recall	F1
Grade 5	0.224	0.153	0.182
Grade 6	0.268	0.168	0.207
Grade 7	0.329	0.270	0.298
Grade 8	0.450	0.285	0.349
Grade 9	0.378	0.278	0.320
Grade 10	0.450	0.387	0.416
Grade 11	0.271	0.479	0.346
Grade 12	0.520	0.477	0.497
Grade 13	0.270	0.251	0.260

To better understand what the grades that the model has a hard time predicting we took a closer look at how the model performed with each grade in one of the ordinal setting with SBERT features (Table 8). The precision and recall values are significantly greater than a random guess or weighted guess classifiers showing that the model is clearly learning the task.

The F1 scores correlate very well with the number of annotations available for each grade (Pearson Coefficient: 0.942) which indicates that more data would be needed for the lower and higher grades to improve performance.

## 5. Conclusions

Learning Object metadata enrichment is important for precise Learning Object search. A key but under-investigated part of this is grade level prediction which is studied here. The effectiveness of various feature representations of the textual content, of various ML formulations and the impact of entity linking on this task are covered. We show that dense vector representations of textual content lead to the best performance. However, if interpretability is desired, entity linking approaches perform better than TF-IDF or keyword extraction. We show filtering entities by domain improves performance at this task. This is a first attempt at a solution for this task. There are many avenues for improvement including: enhancing entity linking features with pedagogically relevant information, assembling a more balanced dataset and representing more of the lesson like equations, images, animations and lesson structure. These improvements would ultimately raise metadata quality which directly contributes to search features in Learning Object Repositories being more effective.

## 6. Acknowledgments

This work has been partly supported by the Ministry of Science and Education of Lower Saxony, Germany, through the Graduate training network “LernMINT: Data-assisted classroom teaching in the MINT subjects”. We would like to thank the reviewers for their valuable feedback.

## References

- [1] The Learning Object Metadata standard – IEEE Learning Technology Standards Committee, URL: <https://www.ieeeltsc.org/working-groups/wg12LOM/lomDescription/>.
- [2] T. Anderson, C. Leachman, Strategies for supporting oer adoption through faculty and instructor use of a federated search tool, *Journal of Librarianship and Scholarly Communication* 7 (2019).
- [3] S.-H. G. . F.-F. N. . A. Ernest, Repositories of open educational resources: An assessment of reuse and educational aspects, *International Review of Research in Open and Distributed Learning* 18 (2017) 84–120. doi:<https://doi.org/10.19173/irrodl.v18i5.3063>.
- [4] M. Tavakoli, M. Elias, G. Kismihók, S. Auer, Quality prediction of open educational resources A metadata-based approach, in: *20th IEEE International Conference on Advanced Learning Technologies, ICALT 2020, Tartu, Estonia, July 6-9, 2020, IEEE, 2020*, pp. 29–31.

URL: <https://doi.org/10.1109/ICALT49669.2020.00007>. doi:10.1109/ICALT49669.2020.00007.

- [5] X. Ochoa, E. Duval, Automatic evaluation of metadata quality in digital repositories, *Int. J. Digit. Libr.* 10 (2009) 67–91. URL: <https://doi.org/10.1007/s00799-009-0054-4>. doi:10.1007/s00799-009-0054-4.
- [6] S. Dietze, D. Taibi, R. Yu, P. Barker, M. d'Aquin, Analysing and improving embedded markup of learning resources on the web, in: R. Barrett, R. Cummings, E. Agichtein, E. Gabrilovich (Eds.), *Proceedings of the 26th International Conference on World Wide Web Companion*, Perth, Australia, April 3-7, 2017, ACM, 2017, pp. 283–292. URL: <https://doi.org/10.1145/3041021.3054160>. doi:10.1145/3041021.3054160.
- [7] A. Behr, A. B. Mendes, P. Trigo, J. Cascalho, H. Guerra, A. Costa, M. Parente, A. Botelho, R. Vicari, Recommending metadata contents for learning objects through linked data, in: F. de la Prieta, A. E. Bolock, D. Durães, J. Carneiro, F. Lopes, V. Julián (Eds.), *Highlights in Practical Applications of Agents, Multi-Agent Systems, and Social Good. The PAAMS Collection - International Workshops of PAAMS 2021*, Salamanca, Spain, October 6-9, 2021, *Proceedings*, volume 1472 of *Communications in Computer and Information Science*, Springer, 2021, pp. 115–126. URL: [https://doi.org/10.1007/978-3-030-85710-3\\_10](https://doi.org/10.1007/978-3-030-85710-3_10). doi:10.1007/978-3-030-85710-3\_10.
- [8] J. F. H. Cubides, P. A. G. García, C. E. M. Marín, S. Sánchez-Alonso, Improving OER descriptions to enhance their availability, reuse, and enrichment, *Educ. Inf. Technol.* 27 (2022) 1811–1839. URL: <https://doi.org/10.1007/s10639-021-10641-w>. doi:10.1007/s10639-021-10641-w.
- [9] T. Richter, U. D. Ehlers, Barriers and motivators for using open educational resources in schools, *Open Ed* (2010) 1–12.
- [10] G. A. Osorio-Zuluaga, N. D. Duque-Mendez, Search and selection of learning objects in repositories: A review, in: *2018 XIII Latin American Conference on Learning Technologies (LACLO)*, IEEE, 2018, pp. 513–520.
- [11] X. Ochoa, E. Duval, Relevance ranking metrics for learning objects, *IEEE Trans. Learn. Technol.* 1 (2008) 34–48. URL: <https://doi.org/10.1109/TLT.2008.1>. doi:10.1109/TLT.2008.1.
- [12] R. D. C. M. Rueda, S. Lujan, A quantitative analysis of the use of microdata for semantic annotations on educational resources, *J. Web Eng.* 17 (2018) 45–72. URL: <http://www.rintonpress.com/xjwe17/jwe-17-12/045-072.pdf>.
- [13] X. Ochoa, K. Cardinaels, M. Meire, E. Duval, Frameworks for the automatic indexation of learning management systems content into learning object repositories, in: *EdMedia+ Innovate Learning*, Association for the Advancement of Computing in Education (AACE), 2005, pp. 1407–1414.
- [14] S. Hassan, R. Mihalcea, Learning to identify educational materials, *ACM Trans. Speech Lang. Process.* 8 (2011) 2:1–2:18. URL: <https://doi.org/10.1145/2050100.2050101>. doi:10.1145/2050100.2050101.
- [15] C. Grévisse, R. Manrique, O. Mariño, S. Rothkugel, Knowledge graph-based teacher support for learning material authoring, in: *Colombian Conference on Computing*, Springer, 2018, pp. 177–191.
- [16] J. Daiber, M. Jakob, C. Hokamp, P. N. Mendes, Improving efficiency and accuracy in

- multilingual entity extraction, in: Proceedings of the 9th International Conference on Semantic Systems (I-Semantics), 2013.
- [17] I. A. Chacon, S. A. Sosnovsky, Knowledge models from PDF textbooks, *New Rev. Hypermedia Multim.* 27 (2021) 128–176. URL: <https://doi.org/10.1080/13614568.2021.1889692>. doi:10.1080/13614568.2021.1889692.
- [18] R. Reinanda, E. Meij, M. de Rijke, Knowledge graphs: An information retrieval perspective, *Found. Trends Inf. Retr.* 14 (2020) 289–444. URL: <https://doi.org/10.1561/15000000063>. doi:10.1561/15000000063.
- [19] A. Barbaresi, Generic web content extraction with open-source software, in: Proceedings of the 15th Conference on Natural Language Processing, KONVENS 2019, Erlangen, Germany, October 9-11, 2019, 2019. URL: [https://corpora.linguistik.uni-erlangen.de/data/konvens/proceedings/papers/kaleidoskop/camera\\_ready\\_barbaresi.pdf](https://corpora.linguistik.uni-erlangen.de/data/konvens/proceedings/papers/kaleidoskop/camera_ready_barbaresi.pdf).
- [20] R. Mihalcea, A. Csomai, Wikify!: linking documents to encyclopedic knowledge, in: M. J. Silva, A. H. F. Laender, R. A. Baeza-Yates, D. L. McGuinness, B. Olstad, Ø. H. Olsen, A. O. Falcão (Eds.), Proceedings of the Sixteenth ACM Conference on Information and Knowledge Management, CIKM 2007, Lisbon, Portugal, November 6-10, 2007, ACM, 2007, pp. 233–242. URL: <https://doi.org/10.1145/1321440.1321475>. doi:10.1145/1321440.1321475.
- [21] L. Page, S. Brin, R. Motwani, T. Winograd, The PageRank citation ranking: Bringing order to the web., Technical Report, Stanford InfoLab, 1999.
- [22] A. Moro, F. Cecconi, R. Navigli, Multilingual word sense disambiguation and entity linking for everybody, in: M. Horridge, M. Rospocher, J. van Ossenbruggen (Eds.), Proceedings of the ISWC 2014 Posters & Demonstrations Track a track within the 13th International Semantic Web Conference, ISWC 2014, Riva del Garda, Italy, October 21, 2014, volume 1272 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2014, pp. 25–28. URL: [http://ceur-ws.org/Vol-1272/paper\\_30.pdf](http://ceur-ws.org/Vol-1272/paper_30.pdf).
- [23] R. Campos, V. Mangaravite, A. Pasquali, A. Jorge, C. Nunes, A. Jatowt, Yake! keyword extraction from single documents using multiple local features, *Inf. Sci.* 509 (2020) 257–289. URL: <https://doi.org/10.1016/j.ins.2019.09.013>. doi:10.1016/j.ins.2019.09.013.
- [24] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: B. Krishnapuram, M. Shah, A. J. Smola, C. C. Aggarwal, D. Shen, R. Rastogi (Eds.), Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016, ACM, 2016, pp. 785–794. URL: <https://doi.org/10.1145/2939672.2939785>. doi:10.1145/2939672.2939785.
- [25] E. Frank, M. A. Hall, A simple approach to ordinal classification, in: L. D. Raedt, P. A. Flach (Eds.), Machine Learning: EMCL 2001, 12th European Conference on Machine Learning, Freiburg, Germany, September 5-7, 2001, Proceedings, volume 2167 of *Lecture Notes in Computer Science*, Springer, 2001, pp. 145–156. URL: [https://doi.org/10.1007/3-540-44795-4\\_13](https://doi.org/10.1007/3-540-44795-4_13). doi:10.1007/3-540-44795-4\_13.
- [26] E. Terreau, A. Gourru, J. Velcin, Writing style author embedding evaluation, in: Proceedings of the 2nd Workshop on Evaluation and Comparison of NLP Systems, 2021, pp. 84–93.
- [27] F. Petroni, T. Rocktäschel, P. S. H. Lewis, A. Bakhtin, Y. Wu, A. H. Miller, S. Riedel, Language models as knowledge bases?, *CoRR abs/1909.01066* (2019). URL: <http://arxiv.org/abs/1909.01066>. arXiv:1909.01066.
- [28] C. K. Pereira, B. P. Nunes, S. W. M. Siqueira, R. Manrique, J. F. Medeiros, 'a little knowledge

is a dangerous thing': A method to automatically detect knowledge compartmentalization and oversimplification, in: 20th IEEE International Conference on Advanced Learning Technologies, ICALT 2020, Tartu, Estonia, July 6-9, 2020, IEEE, 2020, pp. 140–144. URL: <https://doi.org/10.1109/ICALT49669.2020.00048>. doi:10.1109/ICALT49669.2020.00048.

- [29] R. Manrique, B. P. Nunes, O. Mariño, Exploring knowledge graphs for the identification of concept prerequisites, *Smart Learn. Environ.* 6 (2019) 21. URL: <https://doi.org/10.1186/s40561-019-0104-3>. doi:10.1186/s40561-019-0104-3.