What Can Tweets and Knowledge Graphs Tell Us About Eating Disorders?

José Alberto Benítez-Andrades [a,*], María Teresa García-Ordás [b], Mayra Russo [c], Ahmad Sakor [c], Luis Daniel Fernandes Rotger [d] and María-Esther Vidal [c]

[a] *SALBIS Research Group, Department of Electric, Systems and Automatics Engineering, Universidad de León, León, Spain*
E-mail: jbena@unileon.es

[b] *SECOMUCI Research Group, Escuela de Ingenierías Industrial e Informática, Universidad de León,León, Spain*
E-mail: mgaro@unileon.es

[c] *Leibniz University of Hannover and L3S Research Center and TIB Leibniz Information Centre for Science and Technology, Germany*
E-mails: mrusso@l3s.de, ahmad.sakor@tib.eu, maria.vidal@tib.eu

[d] *Bakken and Baeck GmbH, Germany*
E-mail: luis@bakkenbaeck.no

Abstract. Social networks have become information dissemination channels, where announcements are posted frequently; they also serve as frameworks for debates in various areas (e.g., scientific, political, and social). In particular, in the health area, social networks represent a channel to communicate and disseminate novel treatments' success; they also allow ordinary people to express their concerns about a disease or disorder. As a response, the Artificial Intelligence (AI) community has developed analytical methods to uncover and predict patterns from the posts that enable to explain news about a particular topic, e.g., mental disorders expressed as eating disorders or depression. Albeit potentially rich while expressing an idea or concern, posts are presented as short texts, preventing, thus, AI model from accurately encoding these posts' contextual knowledge. We propose a hybrid approach where knowledge encoded in a community maintained knowledge graphs (e.g., Wikidata) is combined with deep learning to categorise social media posts using existing classification models. The proposed approach resorts to state-of-the-art named entity recognizers and linkers (e.g., FALCON 2.0 and EntityLinker in spaCy Python library) to extract entities in short posts and link them to concepts in knowledge graphs (e.g., Wikidata). Then, knowledge graph embeddings (e.g., RDF2Vec) are utilised to compute latent representations of the extracted entities, which result in a vector representation of the posts that encode these entities' contextual knowledge extracted from the knowledge graphs. These knowledge graph embeddings are combined with contextualized word embeddings (e.g., BERT) to generate a context-based representation of the posts that empower prediction models. We apply our proposed approach in the health domain to detect whether a publication is related to an eating disorder (e.g., anorexia or bulimia) and uncover concepts within the discourse that could help healthcare providers prevent and diagnose this type of mental disorder. We evaluate our approach on a dataset composed of 2,000 short texts related to eating disorders. Our experimental results suggest that using knowledge graph exploitation, the semantic enrichment of these messages increases the reliability of the predictive models generated concerning models that do not use the knowledge collected from Wikidata. The ambition is that the proposed method can support health domain experts in discovering patterns that may forecast a mental disorder, enhancing early detection and more precise diagnosis towards personalised medicine.

Keywords: Name Entity Linking, Deep Learning, Natural Language Processing, Health Data, Knowledge Graphs, Wikidata

*Corresponding author. E-mail: jbena@unileon.es.

1. Introduction

Social media is increasingly used as a dissemination channel to announce novel treatments or conditions to respond to health-related problems and even to discuss natural disasters [1]. Moreover, social media networks are utilised to promote or prevent the administration of certain interventions, e.g., in mental disorders like eating disorders (EDs) [2, 3]. Throughout the COVID-19 pandemic, an increase in the number of people suffering from EDs worldwide has been observed [4–6].

Artificial intelligence is increasingly used in the healthcare field [7, 8], to help prevent disease [9], detect pathologies more quickly [10] and find effective diagnoses more quickly [11]. Analysing the discourse available on social networks such as Twitter can help find answers to relevant problems by applying various artificial intelligence techniques such as machine learning and deep learning and, more specifically, data mining, text mining, and different natural language processing (NLP) techniques. For example, thanks to the use of datasets obtained from social networks, it has been possible to generate predictive models capable of detecting anorexia on social media [12].

In addition to artificial intelligence techniques such as those mentioned above, it is also interesting to work with data related to eating disorders in the field of the semantic web, in studies that make use of knowledge graphs such as Wikidata or DBPedia [13–15]. The information contained in these knowledge graphs is increasingly used in the scientific community to solve different problems [16–18]. Among the most widely used techniques in this field is the obtaining of knowledge graph embeddings to calculate the similarity between terms within a knowledge graph [19–21]. This information is of interest to obtain what is known as semantic enrichment in different textual problems for which only textual information is initially available. Following a strategy for semantic enrichment, more data is enriched with meaningful meta-data to empower the performance of recommender systems. Nevertheless, predictive models based on deep learning may perform poorly over short text. We propose a hybrid approach to combine deep learning techniques with knowledge graphs to improve the classification accuracy on short posts like Tweets.

This research addresses different problems. The first is how we can better understand the content of the tweets available in social media data thanks to the use of knowledge graphs and the second problem is how to classify with greater precision and accuracy these tweets that talk about health issues, understanding that BERT models do not have semantic information and have greater difficulties when classifying short texts. Within the field of health, addressing these problems can help to generate early diagnoses of mental disorders with greater accuracy [22], as well as help to generate interventions that can prevent conditions that can cause, for example, suicides in depressed people [23]. The hypothesis is that making use of the information contained in knowledge graphs in the form of knowledge graph embeddings can help improve predictive deep learning models applied to text classification [24–26] thanks to the semantic information contained in them. There are some studies that make use of semantic content contained in ontologies within some domains to improve some automatic classification methods [24–26]. However, there is a lack of studies that make a recognition of the entities and then a linking with these entities in knowledge graphs to end up obtaining the embeddings of knowledge graphs of these entities as semantic content of these texts.

Let us consider a dataset composed of short texts, e.g. related to eating disorders (EDs), there exist numerous natural language processing techniques that can be used to generate machine learning and deep learning models capable of classifying texts using a set of previously labelled texts. Thanks to this automatic classification, it could help to prevent different diseases through early diagnosis. The predictive models currently offering the best results are those known as Bidirectional Encoder Representations from Transformers (BERT) models [27]. The BiLSTM-based (bidirectional long short-memory) models added the novelty of using the bidirectional approach, i.e., reading sentences from left to right and right to left at the same time. However, BERT models, unlike BiLSTMs, are able to learn from words that are in all positions, i.e., from the whole sentence. In order to achieve predictive models that classify texts, these models usually receive texts and labels as input. The performance problem of these models arises when the dataset is composed of short texts.

To illustrate this problem, Fig. 1a shows two tweets labelled into a category representing tweets promoting having an eating disorder. In current classification modelling solutions (Fig. 1b), these texts would be used as input variables to train the model and, using deep learning techniques, a classification model would be created by which the accuracy results of its classification criteria would be obtained. This approach makes use of natural language processing (NLP) techniques that, in order to train a model, convert text into vectors known as embeddings using techniques such as Word2Vec, GloVe [28] or similar. However, this approach is hampered when the texts to be classified are short or poorly structured.

The approach proposed in this research, which can be seen in Fig. 1c, recognises the entities in each tweet and obtains the identifier of this entity in a knowledge base. Once these entities are obtained, the knowledge graph embeddings corresponding to these entities are obtained. Subsequently, it is necessary to combine these embeddings using some techniques already existing in the scientific literature. Finally, it is proposed that the models trained with the embeddings obtained after applying a traditional NLP approach (EMB) together with the embeddings obtained from the knowledge graphs (EMB-KGE) provide the necessary information to improve the classification models for short texts.

In this research we propose to make use of a hybrid model that combines deep learning techniques and knowledge graph exploitation techniques. In Fig. 1c it is possible to see an example in which two tweets are used as input data to recognise their entities and link them to concepts contained in Wikidata. After this, the knowledge graph embeddings would be obtained using techniques such as RDF2Vec. Finally, the knowledge graph embeddings of the entities contained in each tweet are combined to obtain a single embedding for each tweet, which we have called EMB-KGE. After performing all these steps, and as shown on the right side of Fig. 1c, the EMB-KGE data are combined with the embeddings obtained after applying BERT to the texts, which we have called EMB. Using this combination, we obtain predictive models able to offer a better performance, obtaining correctly labelled tweets using the EMB+(EMB-KGE) approach. In summary, the approach used in the example shown in Fig. 1 is as follows:

– Perform a recognition of the entities in each tweet, apply entity linking techniques that allow obtaining the identifier of that entity in a knowledge base.
– Once entities are obtained, the knowledge graph embeddings are computed, for example, using RDF2Vec [29].
– Subsequently, it is necessary to combine these knowledge graph embeddings using some techniques already existing in the scientific literature [30, 31].
– Finally, it is proposed that the models trained with the embeddings obtained after applying a traditional NLP approach (EMB) together with the embeddings obtained from the knowledge graphs (EMB-KGE) provide the necessary information to improve the classification models for short texts.

To evaluate the hypothesis, ten different predictive models were applied using three different data inputs on the same set. One input consisted of textual information, another input consisted of knowledge graph embeddings and the third data input consisted of the textual information together with the knowledge graph embeddings. This hypothesis was applied to a dataset responding to a health problem related to a mental disorder.

This research proposes three contributions of particular interest in the health sector consisting of:

– A new approach to detecting discourse in social data using knowledge graphs. Thanks to this approach, it would be possible to improve the detection and diagnosis of people who may suffer from some kind of mental disorder.
– A new approach to train automatic prediction models on texts by making use of semantic information obtained through collaborative knowledge graphs. Thus improving the accuracy of disease diagnosis through discourse analysis.
– To evaluate the hypothesis, ten different predictive models were applied using three different data inputs on the same set. One input consisted of textual information, another input consisted of knowledge graph embeddings and the third data input consisted of the textual information together with the knowledge graph embeddings. This hypothesis was applied to a dataset of eating disorders (ED).
– An empirical evaluation was performed on a dataset that respond to health problems. The dataset consists of 2,000 labelled texts categorised into four binary categories from which a total of 1,358 different entities were obtained. For these entities, their embeddings were calculated using the RDF2Vec tool [29] making use of the information contained in the Wikidata knowledge graph. The results indicate that models that make use of texts and combined knowledge graph embeddings perform better in 97.5% of the cases, an improvement of up to 15%. The dataset and the code with which the testbeds are available in a public repository, to ensure, thus, reproducibility.

The remainder of this paper is organized as follows. Section 2 presents an analysis of the related work. Then Section 3 presents the fundamental of RDF, knowledge graphs, knowledge graph embeddings, entity recognition and entity linking. Section 4 formalizes the problem solved by our approach and describes the main componente of the architecture. Section 5 reports on the experimental evaluation and observed outcomes. Finally, we close the paper in Section 6 with conclusions and an overview of future work.

2. Related Work

2.1. Deep Learning in Social Media Data

The use of social media data to train predictive models that make use of machine learning and deep learning techniques is becoming increasingly common in the scientific field. A search for the terms "social media data AND deep learning" in Google Scholar returns approximately 225,000 results since 2017[1]. Some of these studies focus on generating predictive models that are able to classify texts within a given domain, such as mental disorder [12], emotions [32], or fake news [33]. These studies make use of machine learning and deep learning techniques, comparing the results obtained and the computational cost of each. Currently, the techniques that obtain the best results in terms of performance and accuracy of the classification models are long short-term memory (LSTM), bidirectional long short-term memory (Bi-LSTM) neural networks and, since they appeared, bidirectional encoder representations from transformers (BERT) models. Although the predictive models obtained in many studies have a fairly high performance in the form of hit rates, our hypothesis is that most of them could be improved if, in addition, the semantic information of these texts were used.

2.2. Knowledge acquisition from Knowledge Graphs

Knowledge graphs such as DBPedia and Wikidata are increasingly used in research. Moreover, studies have shown that the information contained in them is useful and reliable for use in some fields such as the life sciences [34]. One of the most important problems faced by scientists in the field of knowledge acquisition through these collaborative knowledge networks is what is known as Entity Linking. In order to obtain a resource that represents the concept “help” from Wikidata, it is necessary to know whether we are referring to the concept of help as cooperation between people [2] or whether we are referring to the studio album called "Help!" by The Beatles [3]. Due to the growing use of these knowledge graphs in science, tools have emerged that make it possible to obtain, in a simple way, not only the entities in a sentence, but also the link to these concepts in collaborative knowledge graphs. Some examples are the FALCON 2.0 tool [35] and the EntityLinker [4] function available at spaCy [36] library that is implemented and available for use in Python. Thanks to these tools, it is easier and faster to obtain information associated with the concepts contained in texts and, with it, it is possible to perform queries that allow us to obtain some interesting metrics such as, for example, the embeddings of these concepts using tools such as RDF2Vec [29].

2.3. Combining Knowledge Graphs and Deep Learning

The information contained in the knowledge graphs combined with the use of deep learning techniques are being used to solve different problems raised in the research. For example, thanks to the semantic information obtained through these knowledge graphs, it is possible to obtain a better interpretability and explainability of different predictive models created with [37] deep learning techniques. This increase in the interpretability and explainability of the models is of vital importance in applications related to health or education. Knowledge graphs are also being used to help predict relationships between different concepts, such as, for example, the study [38] predicting the

[1]https://scholar.google.es/scholar?as_ylo=2017&q=social+media+data+AND+deep+learning&hl=es&as_sdt=0,5
[2]https://www.wikidata.org/wiki/Q1643184
[3]https://www.wikidata.org/wiki/Q201816
[4]https://spacy.io/api/entitylinker

relationship between gut microbiota and mental disorder. However, there is a lack of studies that have made use of the information contained in knowledge graphs to classify texts by making use of this information. Our study aims to demonstrate that attaching semantic enrichment to texts can help generate predictive models with better performance.

3. Preliminaries

3.1. RDF and Knowledge Graphs

In this research, we propose the use of collaborative knowledge graphs and, more specifically, Wikidata, to obtain information related to the concepts contained in short texts. The Wikidata knowledge graph is stored internally in JSON format, and can be edited by any user thanks to the web interface of this knowledge graph. The concepts found in this knowledge graph are represented in RDF format. The Wikidata structure consists mainly of data items and data values that are connected by properties. Entities and properties are first class objects that can be used in a global way.

Although Wikidata is primarily based on the RDF format, this knowledge graph does have some differences from RDF. The first difference is that RDF only allows each component to carry a single property or value, whereas, in the Wikidata knowledge graph, objects can be composed. The second difference is that subject-property-object connections can themselves be the subject of auxiliary property-value pairs. These are called qualifiers in Wikidata [39]. It is worth noting that Wikidata is being widely used and provides rich piece of knowledge at encyclopedic and domain specific level, e.g., in the field of life sciences [34].

3.2. Entity recognition and entity linking, making use of knowledge graphs

In our approach, we make use of what are known as named entity recognisers and linkers to extract the entities found in the short texts used as input data and link it to the concepts contained in Wikidata. A named entity linker implements the tasks of named entity recognition (NER) and named entity linking (NEL) allowing the identification of entities in a text and their corresponding resources in a knowledge graph or controlled vocabulary. Although our approach is named entity linker agnostic, as a proof of concept, we make use of the EntityLinker [5] contained in the spaCy [36] Python library and FALCON 2.0 [35]. Both tools resort to background knowledge to perform the linking process. EntityLinker's knowledge base is updated with Wikidata resources from 2019 [6], while FALCON 2.0 has background knowledge updated to 2021. The two systems are used because the recognised entities were not always the same, and it was observed that a union of the entities obtained using both tools are more complete.

3.3. Knowledge Graph Embeddings

In order to understand our approach, it is necessary to define the concept of knowledge graph embedding. Knowledge graph embeddings (KGE) are low-dimensional representations of the entities and relationships of a knowledge graph. KGEs make possible a generalisable representation of an entity based on its context on a global knowledge graph, in the case of our approach, Wikidata. This context allows inferring relationships between concepts. The KGEs provide us with information about the relationships between different terms related with our problem, for example, in an eating disorder context, KGEs could provide us the relationships between 'green tea', 'day', 'good', 'anorexic', 'binge eating' and 'nah'. Thanks to the magnitude of information contained in Wikidata, the KGEs provide information about the interactions between the concepts contained in the short texts we want to classify.

This information, added to the information contained in the tweets themselves, is of vital importance to improve the predictive models that classify these short texts. There are many configurations and methods for calculating these knowledge graph embeddings. In our approach, we have made use of RDF2Vec [29]. Inspired by word2vec [40]

[5] https://spacy.io/api/entitylinker
[6] https://www.kaggle.com/kenshoresearch/kensho-derived-wikimedia-data

which represents words in vector space, RDF2Vec applies this method within a knowledge graph. RDF2Vec allows receiving different ways to create sequences of RDF nodes that are then used as input for the word2vec algorithm. One of the most commonly used strategies is random walks in an RDF graph.

3.4. Strategies for combining embeddings

In our approach, once the knowledge graph embeddings are obtained, it is necessary to combine them, as the number of entities in each tweet may not always be the same and, therefore, the number of KGEs may also be different. As can be seen in the scientific literature, there are many methods of combining embeddings [30, 31, 41, 42], the simplest of which is to find the average of the total number of embeddings obtained in each tweet. However, this does not seem to be the best strategy. In our approach, we have used a technique known as smooth inverse frequency (SIF) [31].This approach also takes into account which words are the most meaningful within a sentence. The algorithm scores from 0 to 5 pairs of sentences, in our case tweets, according to the similarity between them.

4. Problem Definition and Our approach

In this research, we propose a hybrid approach based on the combination of deep learning techniques and the exploitation of knowledge graphs in order to obtain classification models capable of categorizing short texts, such as tweets, with a higher accuracy than deep learning models that make use of textual information such as, for example, BERT models. We define the problem of short text classification through AI models and state the proposed solution throughout this section.

4.1. Problem Statement and Proposed Solution

Problem Statement. The problem that arises in this research is the low accuracy obtained in classification models that make use of advanced AI techniques such as BERT. As shown in Fig. 1, when a dataset of short texts is available, predictive models that classify texts have problems in assigning labels to these texts. This problem needs to be addressed since, at present, information obtained from social networks such as Twitter is being widely used.

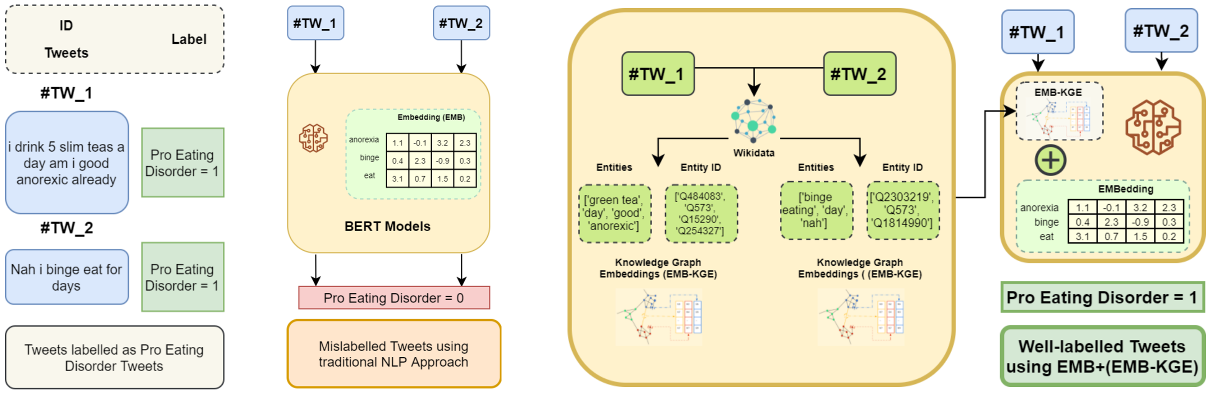

(a) Examplier tweets (b) BERT models approach (c) Deep Learning adding Knowledge Graph Exploitation

Fig. 1. Motivating example. (a) 2 Labelled tweets (b) A text classification model using BERT models approach (c) Short text classification using the novel approach making use of BERT Embeddings (EMB) and Knowledge Graph Emeddings (EMB-KGE).

Proposed Solution. The solution we propose in this research consists of using the information contained in knowledge graphs to improve predictive models for short text classification. To achieve this goal we propose a hybrid model by which the results of applying deep learning to an initial dataset are obtained. After that, entity recognition and entity linking techniques are applied on knowledge graphs. Then, knowledge graphs (e.g. RDF2Vec) are used to compute the latent representations of these resources, which, combined, give rise to a vector representation of the posts. In our solution we propose that making use of the textual information coupled with the information obtained by exploiting knowledge graphs, we obtain short text classification models with higher accuracy that solve the proposed problem.

4.2. Architecture of the proposed approach

Fig. 2 depicts the architecture of the proposed approach. It receives an input data set composed of short texts and outputs the classification generated by a predictive model.

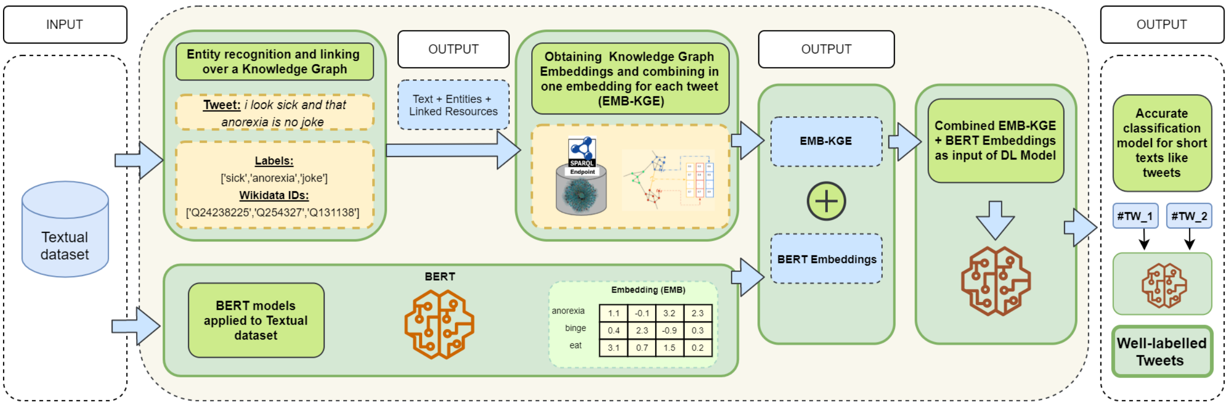

Fig. 2. Architecture of the proposed approach: Starting from a set of text data, (i) entity recognition and linking to concepts within a knowledge graph is performed, (ii) knowledge graph embeddings are extracted from a knowledge base (iii) the KGEs are combined to have a combined KGE called EMB-KGE for each tweet and (iv) the predictive model is trained combining embeddings obtained from texutal dataset using BERT (EMB) with EMB-KGE, EMB+(EMB-KGE).

Entity recognition and linking over a Knowledge Graph. The first step consists of recognizing the entities contained in the short texts used as input and, making use of state-of-the-art named entity linkers (e.g., FALCON 2.0 [35] and EntityLinker in the spaCy [36] Python literature) to link the entities in the short texts with the concepts in a knowledge graph such as Wikidata.

For example, using the FALCON 2.0 API [7], it is possible to obtain all the entities of a sentence and the identifiers that these entities have in Wikidata, for example, of the sentence:

broke my fast at 47hrs cause i felt like i was seriously about to die, i look sick and that anorexia is no joke i'm feeling bad now

The entities recognised are:

['fast', 'death', 'sick', 'anorexia', 'joke']

And the identifiers of these entities in Wikidata are:

['Q44602', 'Q4', 'Q24238225', 'Q254327', 'Q131138']

Obtaining knowledge graph embeddings. After collecting the entities contained in the dataset, their knowledge graph embeddings are obtained using an approach such as RDF2Vec [29]. These knowledge graph embeddings represent the similarity of terms within the Wikidata knowledge graph. Thus, it is possible to add semantic enrichment to the data that has been collected.

Combining knowledge graph embeddings in one embedding for each tweet (EMB-KGE). Having obtained these knowledge graph embeddings, a combined knowledge graph embedding for each tweet is calculated, in order

[7]https://labs.tib.eu/falcon/falcon2/api-use

to train deep learning models using a single vector of the same dimension for each tweet. In order to perform this calculation, an embedding combination system known as smooth inverse frequency (SIF) [31]. The method used by SIF to calculate the combined embedding in a sentence, or, in the case of this research, in a tweet, does not rely solely on finding the average of the knowledge graph embeddings of the entities in each tweet, but uses an approach based on discourse vectors and the random walk model. These combined embeddings are called embeddings obtained from the knowledge graph explotation (EMB-KGE).

BERT models applied to textual dataset. The last step consists of combining the embeddings obtained from the initial dataset composed of short texts after applying deep learning models.

DL Model using combined (EMB-KGE)+(Textual embeddings) as input. Through this process, text embeddings (EMB) are obtained, which together with the EMB-KGE provide a dataset with information through which, by training predictive models based on deep learning techniques, provide a better classification of short texts in different categories.

5. Experiments and results

The empirical study aims at answering the following research questions:

RQ1) What are the tweets and knowledge graphs telling us about eating disorders? What terminology is being used?

RQ2) Is it possible to obtain better results in a deep learning short text classification model using only the information contained in knowledge graphs?

RQ3) Can existing knowledge in knowledge graphs be used to improve deep learning models applied to text?

Data availability of the experiments and code of empirical evaluation of the proposal approach [8] are publicly available; this facilitates the reproducibility of the reported results.

5.1. Experimental setup

Dataset used: The social data used in this study were an own dataset on eating disorders (ED dataset).

5.1.1. Collecting and labelling ED Dataset

The own dataset (ED dataset) was retrieved through information-filtering based methods of social media posts using a set of keywords as search queries. While the method of collection used allows for faster and less costly access to a bigger volume of data and with a specialized level of detail as to the desired output when compared to what could be obtained through traditional collection techniques such as clinical screening surveys [43, 44]. The datasets produced come with their limitations due in great part to the pervasive noise associated with their source, which may impact on the quality of the data itself, as well as on the reliability and representativeness of the results obtained [43, 45].

In this particular case, the social media platform used is Twitter through the T-Hoarder tool [46], and the non-comprehensive list of hashtags used to filter tweets is: anorexia, anorexic, dietary disorders, inappetence, feeding disorder, food problem, binge-eating, eating disorders, bulimia, food issues, loss of appetite. Regarding this, we would like to acknowledge that this selection might bias the sampling and results, as we might be making assumptions about the behaviour of the users generating the content or overlooking relevant information outside these limits [43]. The collection yielded the capture of 494,025 tweets, and in order to mitigate issues associated with lexical and semantic redudancy of the content [43], a process of cleaning (removal of duplicates, re-tweets, re-shared content), manual curation and annotation was performed, thus, resulting in the creation of a subset of 2,000 tweets. This subset was labelled in four categories:

- **ED I:** tweets written by people with ED (1) or not (0).
- **ED II:** tweets that promote having ED (1) or not (0).
- **ED III:** tweets of an informative nature (1) or not (0).
- **ED IV:** tweets of a scientific nature (1) or not (0).

[8]https://github.com/knowledgeb/Combining-Knowledge-Graphs-and-Deep-Learning-techniques-for-Categorizing-Tweets

This dataset is available for download from github [9].
After data collection and pre-processing, a dataset consisting of 2,000 tweets messages, and 4 binary columns representing the 4 categories mentioned above is collected. On account of the characteristics of the data and the research itself, it is not our objective to profile the sampled users, thus we avoid making assumptions based on 'gender', 'nationality', 'socio-economic background', 'age', to name a few features, that could lead us to fall back on perpetuating damaging stereotypes associated with persons who suffer eating disorders, as recent research has shown that they are global illnesses that do not discriminate [47, 48]. However, all the collected tweets belong to users that express themselves in English, a result of the predetermined search criteria. Additionally, as part of the 2,000 tweets, 1,567 belong to unique Twitter accounts, meaning that 433 tweets are concentrated in 183 Twitter accounts, a figure that could indicate there is no presence of over-representation of 'noisy' social media users, notwithstanding the possibility that this distribution includes accounts that belong to nonhumans (i.e., bots, corporate accounts), multiple users posting from the same account, or the same entity posting from different accounts [43, 49]
Through an exploration of the frequency of terms associated with the hashtags collected, we managed to identify a predominant mention of anorexia and other derivative terms such as "anorexia", "anorexic", "proana", "ana", "anatwt", "anorexiatips", (1,034 associated hashtags), while an underrepresentation of hashtags associated with other eating disorders, such as bulimia (182 associated hashtags) is identified. This pattern is also replicated in the analysis of the content of the tweets themselves, as the most frequent terms were also associated with anorexia (more than 428 unigram tokens) over bulimia (109 tokens), a small caveat regarding this will be that binge (282 tokens) eating, is a behaviour that could both be identified as a symptom of bulimia or a disorder in itself. While this detangling has not been the focus of this study, the results leave the door open to perform further captures that could cover a wider variety of eating disorders or for a fine-grained analysis that could help semantically differentiate these disorders in short text analysis.

Figure 3 shows a summary of the analysis of the data. Most frequent terms with unigram tokens (see Fig. 3b and table 3a) and a breakdown of the annotated categories and the corresponding class count is showed in Fig. 3c. The top 35 hashtags show a predominant mention of 'anorexia' and other derivative terms such as "anorexia", "anorexic", "proana", "ana", "anatwt", "anorexiatips".

word	frequency	word	frequency
eat	1132	research	226
disorder	830	study	202
food	410	problem	199
recovery	382	anorexia	184
edtwt	301	help	151
binge	282	...	...
people	245	bulimia	109
anorexic	244		

(a) Most frequent terms in the dataset.

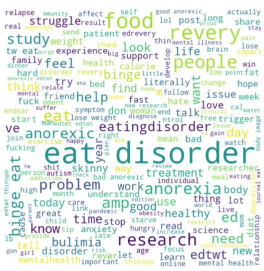

(b) Frequent unigram wordcloud.

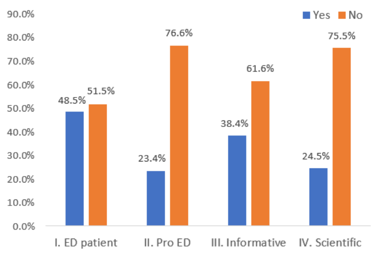

(c) Distribution of data by category.

Fig. 3. Results of the analysis of the dataset showing the most frequent terms and the distribution of the categories.

5.1.2. *Knowledge Acquisition through Knowledge Graph*

The texts of the dataset were processed with FALCON 2.0 [35] and EntityLinker [36], resulting in a total number of entities in the Wikipedia knowledge graph, as shown in table 1. From the total number of entities, the total number of unique entities was calculated and a dataset was obtained consisting of those entities that appeared at least twice in each of the tweets. By manually reviewing the entities obtained in the datasets, some disambiguation errors were detected and manually corrected. Some entity linking errors that were manually modified are shown in table 2a.
In addition, several concepts were added to Wikidata showed in Table 2b. After obtaining all the entities, the knowledge graph embeddings were collected using the pyRDF2Vec tool [50] (setting *RandomWalker* with *max_depth=4*

[9] https://github.com/knowledgeb/Combining-Knowledge-Graphs-and-Deep-Learning-techniques-for-Categorizing-Tweets/blob/main/kge_ml_dl_models/ed-dataset-falcon_spacy2-embeddings-sentence-md4.csv

Dataset	Tweets	Entities	Unique entities	Unique entities >2
Eating disorders	2,000	11,680	1,743	1,358

Table 1

Number of entities, unique entities and unique entities appearing two or more times obtained from Wikidata for each of the datasets used.

Term	Wrong link	Correct link
recovery	Recovery, music album (Q274533)	Recovery approach (Q2135807)
anorexia	Anorexia, music album (Q4770169)	anorexia, medical syntom (Q254327)
ed	Ed, tv serie (Q930797)	eating disorder (Q373822)
binger	binger, town (Q544455)	binge eating (Q2303219)
help	The Help, a film (Q204374)	help, cooperation (Q1643184)

(a) Terms tagged with misconceptions on Wikidata.

Concept	Wikidata ID
Food avoidance emotional disorder	Q108760799
fatspo	Q111780867
Addictive Eaters Anonymous	Q111781180
Meanspo	Q111781194
Ultra-processed food	Q111781198

(b) New terms added to Wikidata.

Table 2

Experimental tables. (a) Terms tagged with misconceptions on Wikidata by FALCON [35] and EntityLinker [36]. (b) News terms added about eating disorders added to Wikidata.

and *max_walks=50*) and the SIF algorithm is applied to combine the knowledge graph embeddings of each tweet into a single embedding. In this way, a final dataset is obtained containing: the texts of the tweets, the four binary categories, and the information obtained after exploiting the knowledge graphs, i.e. the combined knowledge graph embeddings.

5.1.3. *Deep Learning techniques for categorizing tweets*

After obtaining the combined knowledge graph embeddings of each tweet, it was decided to train and validate different classification models using machine learning and deep learning techniques applied to the four different categorisations. To validate the hypothesis, the same models were trained using three different datasets for each of the datasets used:

- **Texts dataset (Text):** dataset with texts and the labels.
- **Knowledge Graph Embeddings dataset (KGE):** dataset containing the combined knowledge graph embeddings of each tweet, and the labels.
- **Texts + KGE Dataset (Text+KGE):** dataset containing texts, combined knowledge graph embeddings and the labels.

The predictive models used in the validation of the approach proposed in this research are described below:

- **Random Forests (RF):** a Random Forest is an ensemble of decision trees combined with bagging [51].
- **Recurrent Neural Networks (RNN):** A recurrent neural network does not have a defined layered structure, but allows arbitrary connections between neurons, and can even create cycles, thus creating temporality, allowing the network to have memory [52].
- **Bidirectional Long Short-Term Memory (Bi-LSTM):** LSTMs and their bidirectional variants are popular because they have tried to learn how and when to forget and when not to use gates in their architecture. In the textual information domain, they tend to get better results [53].
- **Bidirectional Encoder Representations from Transformers (BERT):** BERT main technical innovation is the application of bidirectional Transformer training, a popular model of attention, to language modelling [27]. In this research, we have decided to test the following seven pre-trained models: TweetBERT [54], BERT [55], RoBERTa [56], DistilBERT [57], CamemBERT [58], Albert [59], and FlauBERT [60].

A k-fold cross validation has been carried out with 5 folds for all of the experiments in order to avoid randomness. Every model has been evaluated through an extensive hyperparameter grid search. In the results, the score for the hyperparameter configuration with the best mean value of the 10 folds is shown.

Metrics. To evaluate the results obtained in each of the models developed, 2 different metrics have been used: F_1-score (F_1) and accuracy (acc) (Table 3). F_1-score is a metric used to calculate the effectiveness of a classifier by taking into account its accuracy and recall values. F1 assumes that the two metrics used are of equal importance in calculating effectiveness. If one of them is more important than the other, a different formula F_β would have to be used. The formula used to calculate this metric is as follows, where P equals the precision value and R equals the recall value. Accuracy refers to how close the result of a measurement is to the true value. In statistical terms, accuracy is related to the bias of an estimate. It is represented as the proportion of true results (both true positives (TP) and true negatives (TN)) divided by the total number of cases examined (true positives, false positives, true negatives, false negatives).

Metric	Formula
Precision	$P(c) = \frac{TP}{TP+FP}$
Recall	$R(c) = \frac{TP}{TP+FN}$
F_1-score	$F_1 = \frac{2*P*R}{P+R}$
Accuracy	$Acc = \frac{TP+TN}{(TP+TN+FP+FN)}$

Table 3

Formulas of the evaluation metrics used

5.1.4. Hardware and software configuration

Each of the ten predictive models applied to classify the data into four different categories, having as input data a dataset of eating disorders, is trained 5 times to ensure validation of the results obtained. All experiments are run on the same computer. The experiments are executed on a Windows 10 Pro x64 machine with an Intel® Core® i7-9700K CPU @ 3.60GHz (eight physical cores, eight threads) and 32 GiB DDR4 RAM with a graphic card NVIDIA GeForce RTX 2080 SUPER (8 GB GDDR6). All experiments are run in an environment with Python 3.6 and CUDA 11.1 installed.

5.2. Discussion of Observed Results

Ten different classification models to classify tweets in four binary categories are compared using three datasets composed by labels and (i) texts, (ii) combined knowledge graph embeddings and (iii) the combination of texts and knowledge graph embeddings.

Table 4 shows all the results obtained after train ten different machine learning and deep learning models to categorize tweets in different categories using three different datasets for:

– **ED dataset:** 2,000 labelled tweets, 2,000 combined knowledge graph embeddings and the combination of both datasets (texts with combined knowledge graph embeddings).

All the results obtained using the dataset formed by the combined knowledge graph embeddings are the ones that have obtained the best score for the health problem of our study.

The models trained with the Text+KGE dataset are the ones that obtain the best results in 97.5% of the cases. The highest percentage improvement in the F_1 metric occurred in ED I in the RF model, and in the *acc* metric occurred in the same category and model. The biggest difference in performance between the model trained with the texts and the model trained with the combination of texts and KGE in the F_1 metric is in the RF model in ED I, with an improvement of 11.11% in the model with texts with KGE information and in the *acc* metric in the RF model in ED III with an improvement of 15.01%. There are only two cases in which the F_1 metric applied to the texts with KGE information data is the same or worsens with respect to that obtained in the text data, in ED III in the CamemBERT and TweetBERT models.

	ED I		ED II		ED III		ED IV	
Model (Data)	F_1	*acc*	F_1	*acc*	F_1	*acc*	F_1	*acc*
RF (KGE)	0.350	0.510	0.701	0.787	0.461	0.591	0.631	0.739
RF (Text)	0.774	0.799	0.910	0.860	0.808	0.733	0.921	0.875
RF (Text+KGE)	**0.860**	**0.860**	**0.927**	**0.878**	**0.874**	**0.843**	**0.933**	**0.902**
RNN (KGE)	0.360	0.499	0.631	0.787	0.504	0.602	0.622	0.739
RNN (Text)	0.801	0.794	0.831	0.812	0.792	0.770	0.839	0.801
RNN (Text+KGE)	**0.851**	**0.858**	**0.848**	**0.823**	**0.829**	**0.802**	**0.863**	**0.842**
Bi-LSTM (KGE)	0.211	0.499	0.656	0.787	0.215	0.602	0.432	0.739 2
Bi-LSTM (Text)	0.791	0.785	0.841	0.823	0.777	0.762	0.864	0.847
Bi-LSTM (Text+KGE)	**0.835**	**0.842**	**0.853**	**0.841**	**0.835**	**0.811**	**0.893**	**0.873**
Albert (KGE)	0.411	0.552	0.254	0.782	0.398	0.580	0.621	0.739
Albert (Text)	0.852	0.848	0.930	0.890	0.872	0.846	0.944	0.919
Albert (Text+KGE)	**0.894**	**0.890**	**0.959**	**0.936**	**0.886**	**0.865**	**0.948**	**0.924**
BERT (KGE)	0.275	0.526	0.479	0.796	0.171	0.586	0.212	0.648
BERT (Text)	**0.876**	0.870	0.933	0.895	0.887	0.863	0.951	0.927
BERT (Text+KGE)	0.875	**0.873**	**0.963**	**0.942**	**0.898**	**0.878**	**0.960**	**0.941**
CamemBERT (KGE)	0.453	0.499	0.166	0.792	0.302	0.602	0.234	0.739
CamemBERT (Text)	0.854	0.841	0.937	0.900	**0.870**	0.846	0.947	0.924
CamemBERT (Text+KGE)	**0.885**	**0.876**	**0.961**	**0.939**	**0.870**	**0.854**	**0.950**	**0.927**
DistilBERT (KGE)	0.515	0.577	0.477	0.786	0.380	0.551	0.229	0.670
DistilBERT (Text)	0.868	0.868	0.932	0.895	0.888	0.865	0.948	0.924
DistilBERT (Text+KGE)	**0.890**	**0.888**	**0.961**	**0.939**	**0.900**	**0.878**	**0.952**	**0.929**
FlauBERT (KGE)	0.041	0.509	0.137	0.797	0.211	0.602	0.672	0.739
FlauBERT (Text)	0.852	0.843	0.928	0.885	0.848	0.827	0.949	0.924
FlauBERT (Text+KGE)	**0.892**	**0.888**	**0.953**	**0.927**	**0.888**	**0.868**	**0.954**	**0.932**
RoBERTa (KGE)	0.499	0.694	0.350	0.768	0.211	0.602	0.234	0.739
RoBERTa (Text)	0.882	0.883	0.941	0.909	0.899	0.875	0.945	0.919
RoBERTa (Text+KGE)	**0.897**	**0.897**	**0.964**	**0.944**	**0.906**	**0.888**	**0.949**	**0.926**
TweetBERT (KGE)	0.354	0.521	0.487	0.785	0.277	0.558	0.234	0.739
TweetBERT (Text)	0.888	0.887	0.933	0.898	**0.890**	0.868	0.953	0.931
TweetBERT (Text+KGE)	**0.909**	**0.909**	**0.964**	**0.944**	0.888	**0.868**	**0.957**	**0.937**

Table 4

Results after applying 10 machine learning models using 3 different input variables: knowledge graph embeddings, texts and the combination of both. The best results are highlighted in bold

Answer to RQ1.

By using the information contained in knowledge graphs, it has been possible to detect concepts mentioned in the discourse contained in the collected social media data. This action can help in the early detection and diagnosis of different mental disorders through discourse analysis.

In the particular case where this research has been applied, mental disorders and, more specifically, eating disorders, it has been possible to carry out an analysis of the content of the tweets that has helped to identify new terminology or jargon found in the social networks, such as, for example, the concepts shown in T¡table 2b (Food avoidance emotional disorder, fatspo, meanspo, etc.).

Answer to RQ2.

From the analysis of the results, it is clear that the use of semantic information obtained through knowledge graphs has a positive effect on the performance of the generated predictive models. Of all the experiments carried out over a dataset related to eating disorders, in 97.5% of the cases, the experiments carried out using as input data those with knowledge graph exploitation obtained better results. In the remaining 2.5% of cases, the model is equally effective or slightly less effective, not exceeding 1% worse than the model using data without KGE.

Answer to RQ3. It is possible to highlight that knowledge graphs can help to improve predictive models by making use of the information contained in them through knowledge graph embeddings between concepts within the graph. Nevertheless, the high percentage of models that provide better performance on the dataset with KGE may be an indication that the performance of any model applied to texts to which this information is added will improve.

6. Conclusions and Future work

We address the problem of improving the performance of predictive models generated through machine learning and deep learning techniques for short text classification by adding semantic information through the exploitation of knowledge graphs. Given the importance of these text classification models in different scientific and industrial contexts, any contribution that improves the performance of these predictive models is of worldwide interest and, more specifically, in the health field. Thanks to the improvement of these text classifiers, it is possible to detect and diagnose diseases earlier, so that interventions can be planned to help reduce the number of patients who may have serious consequences as a result of these diseases.

This research presents a methodology using a combination of semantic text enrichment that exploits the information contained in knowledge graphs in combination with deep learning techniques to improve classification models in short texts. It has been demonstrated, on a real health-related problem, that the approach presented in this research combining the exploitation of knowledge graphs with the use of deep learning techniques to generate predictive models on short texts, obtains a higher success rate than the one obtained with models trained only with textual information. As a result, 97.5% of the trained models performed better with our approach than with the usual approaches. Thus, it is possible to highlight that the use of semantic information contained in knowledge graphs can improve the performance of predictive models in the field of natural language processing and text mining, with the corresponding importance in the health sector. We hope that the results we present will encourage the various communities to add this approach to their text classification research, where these results can be reproduced and generalised to real-world scenarios and healthcare settings. With this new approach, the health and well-being of society can be improved.

Given that the approach presented in this research has been tested on a dataset within a particular health domain, eating disorders, further testing is needed to bring more robustness to the validation of the approach. In the future, we will test on a dataset with a larger amount of data to provide further validity, and we will test this approach using. Finally, the creation of a framework that applies our approach to a given dataset is on our future agenda.

References

[1] G. Song and D. Huang, A Sentiment-Aware Contextual Model for Real-Time Disaster Prediction Using Twitter Data, *Future Internet* **13**(7) (2021). doi:10.3390/fi13070163. https://www.mdpi.com/1999-5903/13/7/163.

[2] A. Budenz, A. Klassen, J. Purtle, E. Yom Tov, M. Yudell and P. Massey, Mental illness and bipolar disorder on Twitter: implications for stigma and social support, *Journal of Mental Health* **29**(2) (2020), 191–199, Publisher: Routledge _eprint: https://doi.org/10.1080/09638237.2019.1677878. doi:10.1080/09638237.2019.1677878.

[3] M. Makita, A. Mas-Bleda, S. Morris and M. Thelwall, Mental Health Discourses on Twitter during Mental Health Awareness Week, *Issues in Mental Health Nursing* **42**(5) (2021), 437–450, Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/01612840.2020.1814914. doi:10.1080/01612840.2020.1814914.

[4] M. Shah, M. Sachdeva and H. Johnston, Eating disorders in the age of COVID-19, *Psychiatry Research* **290** (2020), 113122. doi:10.1016/j.psychres.2020.113122. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7259905/.

[5] S.K. Nutley, A.M. Falise, R. Henderson, V. Apostolou, C.A. Mathews and C.W. Striley, Impact of the COVID-19 Pandemic on Disordered Eating Behavior: Qualitative Analysis of Social Media Posts, *JMIR Mental Health* **8**(1) (2021), e26011, Company: JMIR Mental Health Distributor: JMIR Mental Health Institution: JMIR Mental Health Label: JMIR Mental Health Publisher: JMIR Publications Inc., Toronto, Canada. doi:10.2196/26011. https://mental.jmir.org/2021/1/e26011.

[6] J. Feinmann, Eating disorders during the covid-19 pandemic, *BMJ* **374** (2021), n1787, Publisher: British Medical Journal Publishing Group Section: Feature. doi:10.1136/bmj.n1787. https://www.bmj.com/content/374/bmj.n1787.

[7] J.A. Benítez-Andrades, J.M. Alija-Pérez, I. García-Rodríguez, C. Benavides, H. Alaiz-Moretón, R.P. Vargas and M.T. García-Ordás, BERT Model-Based Approach For Detecting Categories of Tweets in the Field of Eating Disorders (ED), in: *2021 IEEE 34th International Symposium on Computer-Based Medical Systems (CBMS)*, 2021, pp. 586–590, ISSN: 2372-9198. doi:10.1109/CBMS52027.2021.00105.

[8] S. Zhou, Y. Zhao, J. Bian, A.F. Haynos and R. Zhang, Exploring Eating Disorder Topics on Twitter: Machine Learning Approach, *JMIR Medical Informatics* **8**(10) (2020), e18273, Company: JMIR Medical Informatics Distributor: JMIR Medical Informatics Institution: JMIR Medical Informatics Label: JMIR Medical Informatics Publisher: JMIR Publications Inc., Toronto, Canada. doi:10.2196/18273. https://medinform.jmir.org/2020/10/e18273.

[9] M.T. García-ordás, C. Benavides, J.A. Benítez-andrades, H. Alaiz-moretón and I. García-rodríguez, Diabetes detection using deep learning techniques with oversampling and feature augmentation, *Computer Methods and Programs in Biomedicine* **202** (2021). doi:10.1016/j.cmpb.2021.105968.

[10] S. Wang, D.M. Yang, R. Rong, X. Zhan, J. Fujimoto, H. Liu, J. Minna, I.I. Wistuba, Y. Xie and G. Xiao, Artificial Intelligence in Lung Cancer Pathology Image Analysis, *Cancers* **11**(11) (2019). doi:10.3390/cancers11111673. https://www.mdpi.com/2072-6694/11/11/1673.

[11] A. Rodriguez-Gonzalez, G. Hernandez-Chan, R. Colomo-Palacios, J. Miguel Gomez-Berbis, A. Garcia-Crespo, G. Alor-Hernandez and R. Valencia-Garcia, Towards an Ontology to Support Semantics Enabled Diagnostic Decision Support Systems, *Current Bioinformatics* **7**(3) (2012), 234–245. ISBN 1574-8936/2212-392X. doi:http://dx.doi.org/10.2174/157489312802460721. http://www.eurekaselect.com/article/44965.

[12] P. López-Úbeda, F.M. Plaza-del-Arco, M.C. Díaz-Galiano and M.-T. Martín-Valdivia, How Successful Is Transfer Learning for Detecting Anorexia on Social Media?, *Applied Sciences* **11**(4) (2021), 1838, Number: 4 Publisher: Multidisciplinary Digital Publishing Institute. doi:10.3390/app11041838. https://www.mdpi.com/2076-3417/11/4/1838.

[13] E. Cabrio, Artificial Intelligence to Extract, Analyze and Generate Knowledge and Arguments from Texts to Support Informed Interaction and Decision Making (2020), 253.

[14] H. Yu, Health Causal Probability Knowledge Graph: Another Intelligent Health Knowledge Discovery Approach, in: *2020 7th International Conference on Bioinformatics Research and Applications*, ICBRA 2020, Association for Computing Machinery, New York, NY, USA, 2020, pp. 49–58. ISBN 978-1-4503-8813-9. doi:10.1145/3440067.3440077.

[15] T. Kuhn, A. Meroño-Peñuela, A. Malic, J.H. Poelen, A.H. Hurlbert, E. Centeno Ortiz, L.I. Furlong, N. Queralt-Rosinach, C. Chichester, J.M. Banda, E. Willighagen, F. Ehrhart, C. Evelo, T.B. Malas and M. Dumontier, Nanopublications: A Growing Resource of Provenance-Centric Scientific Linked Data, in: *2018 IEEE 14th International Conference on e-Science (e-Science)*, 2018, pp. 83–92. doi:10.1109/eScience.2018.00024.

[16] S. Angioni, A.A. Salatino, F. Osborne, D.R. Recupero and E. Motta, Integrating Knowledge Graphs for Analysing Academia and Industry Dynamics, in: *ADBIS, TPDL and EDA 2020 Common Workshops and Doctoral Consortium*, L. Bellatreche, M. Bieliková, O. Boussaïd, B. Catania, J. Darmont, E. Demidova, F. Duchateau, M. Hall, T. Merčun, B. Novikov, C. Papatheodorou, T. Risse, O. Romero, L. Sautot, G. Talens, R. Wrembel and M. Žumer, eds, Communications in Computer and Information Science, Springer International Publishing, Cham, 2020, pp. 219–225. ISBN 978-3-030-55814-7. doi:10.1007/978-3-030-55814-7_18.

[17] H. Turki, M.A. Hadj Taieb, T. Shafee, T. Lubiana, D. Jemielniak, M.B. Aouicha, J.E. Labra Gayo, E.A. Youngstrom, M. Banat, D. Das, D. Mietchen and on behalf of WikiProject COVID-19, Representing COVID-19 information in collaborative knowledge graphs: The case of Wikidata, *Semantic Web* **Preprint**(Preprint) (2021), 1–32, Publisher: IOS Press. doi:10.3233/SW-210444. https://content.iospress.com/articles/semantic-web/sw210444.

[18] K. Abhishek, V. Pratihar, S.K. Shandilya, S. Tiwari, V.K. Ranjan and S. Tripathi, An intelligent approach for mining knowledge graphs of online news, *International Journal of Computers and Applications* **0**(0) (2021), 1–9, Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/1206212X.2021.1957551. doi:10.1080/1206212X.2021.1957551.

[19] S. Vallurupalli, J. Sleeman and T. Finin, Fine and ultra-fine entity type embeddings for question answering, *CEUR Workshop Proceedings* **2774** (2020), 57–70.

[20] F. Özcan, C. Lei, A. Quamar and V. Efthymiou, Semantic Enrichment of Data for AI Applications, in: *Proceedings of the Fifth Workshop on Data Management for End-To-End Machine Learning*, DEEM '21, Association for Computing Machinery, New York, NY, USA, 2021. ISBN 9781450384865. doi:10.1145/3462462.3468881.

[21] G.A. Gesese, M. Alam and H. Sack, Semantic Entity Enrichment by Leveraging Multilingual Descriptions for Link Prediction, 2020.

[22] K. Zeberga, M. Attique, B. Shah, F. Ali, Y.Z. Jembre and T.-S. Chung, A Novel Text Mining Approach for Mental Health Prediction Using Bi-LSTM and BERT Model, *Computational Intelligence and Neuroscience* **2022** (2022), 7893775. doi:10.1155/2022/7893775.

[23] M. Matero, A. Idnani, Y. Son, S. Giorgi, H. Vu, M. Zamani, P. Limbachiya, S.C. Guntuku and H.A. Schwartz, Suicide Risk Assessment with Multi-level Dual-Context Language and BERT, in: *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 39–44. doi:10.18653/v1/W19-3005. https://aclanthology.org/W19-3005.

[24] R. Alfrjani, T. Osman and G. Cosma, A Hybrid Semantic Knowledgebase-Machine Learning Approach for Opinion Mining, *Data & Knowledge Engineering* **121** (2019), 88–108. doi:https://doi.org/10.1016/j.datak.2019.05.002. https://www.sciencedirect.com/science/article/pii/S0169023X18300399.

[25] M. Kulmanov, F.Z. Smaili, X. Gao and R. Hoehndorf, Semantic similarity and machine learning with ontologies, *Briefings in Bioinformatics* **22**(4) (2021). doi:10.1093/bib/bbaa199.

[26] Y. Abgaz, R. Rocha Souza, J. Methuku, G. Koch and A. Dorn, A Methodology for Semantic Enrichment of Cultural Heritage Images Using Artificial Intelligence Technologies, *Journal of Imaging* **7**(8) (2021), 121, Number: 8 Publisher: Multidisciplinary Digital Publishing Institute. doi:10.3390/jimaging7080121. https://www.mdpi.com/2313-433X/7/8/121.

[27] J. Devlin, M.-W. Chang, K. Lee and K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 4171–4186. doi:10.18653/v1/N19-1423. https://www.aclweb.org/anthology/N19-1423.

[28] T. Pickard, Comparing word2vec and GloVe for Automatic Measurement of MWE Compositionality, in: *Proceedings of the Joint Workshop on Multiword Expressions and Electronic Lexicons*, Association for Computational Linguistics, online, 2020, pp. 95–100. https://aclanthology.org/2020.mwe-1.12.

[29] P. Ristoski and H. Paulheim, RDF2Vec: RDF Graph Embeddings for Data Mining, in: *SEMWEB*, 2016.

[30] Z. Parcheta, G. Sanchis-Trilles, F. Casacuberta and R. Rendahl, Combining Embeddings of Input Data for Text Classification, *Neural Processing Letters* **53** (2021), 3123–3151. doi:10.1007/s11063-020-10312-w.

[31] S. Arora, Y. Liang and T. Ma, A Simple but Tough-to-Beat Baseline for Sentence Embeddings, in: *ICLR*, 2017.

[32] W.S. Paka, R. Bansal, A. Kaushik, S. Sengupta and T. Chakraborty, Cross-SEAN: A cross-stitch semi-supervised neural attention model for COVID-19 fake news detection, *Applied Soft Computing* **107** (2021), 107393. doi:https://doi.org/10.1016/j.asoc.2021.107393. https://www.sciencedirect.com/science/article/pii/S1568494621003161.

[33] J. Xue, J. Chen, R. Hu, C. Chen, C. Zheng, Y. Su and T. Zhu, Twitter Discussions and Emotions About the COVID-19 Pandemic: Machine Learning Approach, *Journal of Medical Internet Research* **22**(11) (2020), e20550. doi:10.2196/20550. http://www.jmir.org/2020/11/e20550/.

[34] A. Waagmeester, G. Stupp, S. Burgstaller-Muehlbacher, B.M. Good, M. Griffith, O.L. Griffith, K. Hanspers, H. Hermjakob, T.S. Hudson, K. Hybiske, S.M. Keating, M. Manske, M. Mayers, D. Mietchen, E. Mitraka, A.R. Pico, T. Putman, A. Riutta, N. Queralt-Rosinach, L.M. Schriml, T. Shafee, D. Slenter, R. Stephan, K. Thornton, G. Tsueng, R. Tu, S. Ul-Hasan, E. Willighagen, C. Wu and A.I. Su, Wikidata as a knowledge graph for the life sciences, *eLife* **9** (2020), e52614, Publisher: eLife Sciences Publications, Ltd. doi:10.7554/eLife.52614.

[35] A. Sakor, K. Singh, A. Patel and M.-E. Vidal, Falcon 2.0: An Entity and Relation Linking Tool over Wikidata, in: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, CIKM '20, Association for Computing Machinery, New York, NY, USA, 2020, pp. 3141–3148. ISBN 978-1-4503-6859-9. doi:10.1145/3340531.3412777.

[36] M. Honnibal and I. Montani, spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing, 2017, To appear.

[37] T. Liu, X. Pan, X. Wang, K.A. Feenstra, J. Heringa and Z. Huang, Predicting the relationships between gut microbiota and mental disorders with knowledge graphs, *Health Information Science and Systems* **9**(1) (2020), 3. doi:10.1007/s13755-020-00128-2.

[38] M. Gaur, K. Faldu and A. Sheth, Semantics of the Black-Box: Can Knowledge Graphs Help Make Deep Learning Systems More Interpretable and Explainable?, *IEEE Internet Computing* **25**(1) (2021), 51–59, Conference Name: IEEE Internet Computing. doi:10.1109/MIC.2020.3031769.

[39] S. Malyshev, M. Krötzsch, L. González, J. Gonsior and A. Bielefeldt, Getting the Most Out of Wikidata: Semantic Technology Usage in Wikipedia's Knowledge Graph, in: *The Semantic Web – ISWC 2018*, D. Vrandečić, K. Bontcheva, M.C. Suárez-Figueroa, V. Presutti, I. Celino, M. Sabou, L.-A. Kaffee and E. Simperl, eds, Springer International Publishing, Cham, 2018, pp. 376–394. ISBN 978-3-030-00668-6.

[40] T. Mikolov, K. Chen, G. Corrado and J. Dean, Efficient Estimation of Word Representations in Vector Space, 2013.

[41] B. Li, H. Zhou, J. He, M. Wang, Y. Yang and L. Li, On the Sentence Embeddings from Pre-trained Language Models, 2020.

[42] C.S. Perone, R. Silveira and T.S. Paula, Evaluation of sentence embeddings in downstream and linguistic probing tasks, 2018.

[43] A. Olteanu, C. Castillo, F. Diaz and E. Kıcıman, Social Data: Biases, Methodological Pitfalls, and Ethical Boundaries, *Frontiers in Big Data* **2** (2019), 13. doi:10.3389/fdata.2019.00013.

[44] T. Wang, M. Brede, A. Ianni and E. Mentzakis, Detecting and Characterizing Eating-Disorder Communities on Social Media, in: *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, WSDM '17, Association for Computing Machinery, New York, NY, USA, 2017, pp. 91–100–. ISBN 9781450346757. doi:10.1145/3018661.3018706.

[45] R. Baeza-Yates, User Generated Content: How Good is It?, in: *Proceedings of the 3rd Workshop on Information Credibility on the Web*, WICOW '09, Association for Computing Machinery, New York, NY, USA, 2009, pp. 1–2–. ISBN 9781605584881. doi:10.1145/1526993.1526995.

[46] M. Congosto, P. Basanta-Val and L. Sanchez-Fernandez, T-Hoarder: A framework to process Twitter data streams, *Journal of Network and Computer Applications* **83** (2017), 28–39. doi:10.1016/j.jnca.2017.01.029.

[47] J.A. Pater, L.E. Reining, A.D. Miller, T. Toscos and E.D. Mynatt, "Notjustgirls": Exploring Male-Related Eating Disordered Content across Social Media Platforms, in: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, Association for Computing Machinery, New York, NY, USA, 2019, pp. 1–13–. ISBN 9781450359702. doi:10.1145/3290605.3300881.

[48] U.S.D. of Health and H. Services, NIMH » Let's Talk About Eating Disorders, Accesed: 2021-10-17. https://www.nimh.nih.gov/health/publications/lets-talk-about-eating-disorders.

[49] R. Baeza-Yates, Bias on the Web, *Commun. ACM* **61**(6) (2018), 54–61–. doi:10.1145/3209581.

[50] G. Vandewiele, B. Steenwinckel, T. Agozzino, M. Weyns, P. Bonte, F. Ongenae and F.D. Turck, pyRDF2Vec: Python Implementation and Extension of RDF2Vec, IDLab, 2020. https://github.com/IBCNServices/pyRDF2Vec.

[51] A. Cutler, D.R. Cutler and J.R. Stevens, Random Forests, in: *Ensemble Machine Learning: Methods and Applications*, C. Zhang and Y. Ma, eds, Springer US, Boston, MA, 2012, pp. 157–175. ISBN 978-1-4419-9326-7. doi:10.1007/978-1-4419-9326-7_5.

[52] L.C. Jain and L.R. Medsker, *Recurrent Neural Networks: Design and Applications*, 1st edn, CRC Press, Inc., USA, 1999. ISBN 0849371813.

[53] B. Jang, M. Kim, G. Harerimana, S.-u. Kang and J.W. Kim, Bi-LSTM Model to Increase Accuracy in Text Classification: Combining Word2vec CNN and Attention Mechanism, *Applied Sciences* **10**(17) (2020), 5841, Number: 17 Publisher: Multidisciplinary Digital Publishing Institute. doi:10.3390/app10175841. https://www.mdpi.com/2076-3417/10/17/5841.

[54] D.Q. Nguyen, T. Vu and A.T. Nguyen, BERTweet: A pre-trained language model for English Tweets, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2020, pp. 9–14.

[55] J. Devlin, M. Chang, K. Lee and K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, *CoRR* **abs/1810.04805** (2018). http://arxiv.org/abs/1810.04805.
[56] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer and V. Stoyanov, RoBERTa: A Robustly Optimized BERT Pretraining Approach, *CoRR* **abs/1907.11692** (2019). http://arxiv.org/abs/1907.11692.
[57] V. Sanh, L. Debut, J. Chaumond and T. Wolf, DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter, *ArXiv* **abs/1910.01108** (2019).
[58] L. Martin, B. Muller, P.J.O. Suárez, Y. Dupont, L. Romary, É.V. de la Clergerie, D. Seddah and B. Sagot, CamemBERT: a Tasty French Language Model, in: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020.
[59] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma and R. Soricut, ALBERT: A Lite BERT for Self-supervised Learning of Language Representations, *CoRR* **abs/1909.11942** (2019). http://arxiv.org/abs/1909.11942.
[60] H. Le, L. Vial, J. Frej, V. Segonne, M. Coavoux, B. Lecouteux, A. Allauzen, B. Crabbé, L. Besacier and D. Schwab, FlauBERT: Unsupervised Language Model Pre-training for French, in: *Proceedings of The 12th Language Resources and Evaluation Conference*, European Language Resources Association, Marseille, France, 2020, pp. 2479–2490. https://www.aclweb.org/anthology/2020.lrec-1.302.