

# Empowering machine learning models with contextual knowledge for enhancing the detection of eating disorders in social media posts

José Alberto Benítez-Andrades<sup>a,\*</sup>, María Teresa García-Ordás<sup>b</sup>, Mayra Russo<sup>c</sup>, Ahmad Sakor<sup>c</sup>, Luis Daniel Fernandes Rotger<sup>d</sup> and Maria-Esther Vidal<sup>c</sup>

<sup>a</sup> *SALBIS Research Group, Department of Electric, Systems and Automatics Engineering, Universidad de León, León, Spain*

*E-mail: [jbena@unileon.es](mailto:jbena@unileon.es)*

<sup>b</sup> *SECOMUCI Research Group, Escuela de Ingenierías Industrial e Informática, Universidad de León, León, Spain*

*E-mail: [mgaro@unileon.es](mailto:mgaro@unileon.es)*

<sup>c</sup> *Leibniz University of Hannover and L3S Research Center and TIB Leibniz Information Centre for Science and Technology, Germany*

*E-mails: [mrusso@l3s.de](mailto:mrusso@l3s.de), [ahmad.sakor@tib.eu](mailto:ahmad.sakor@tib.eu), [maria.vidal@tib.eu](mailto:maria.vidal@tib.eu)*

<sup>d</sup> *Bakken and Baeck GmbH, Germany*

*E-mail: [luis@bakkenbaeck.no](mailto:luis@bakkenbaeck.no)*

**Editors:** Haridimos Kondylakis, FORTH-ICS, Greece; Praveen Rao, University of Missouri, USA; Kostas Stefanidis, Tampere University, Finland

**Solicited reviews:** João Rafael Almeida, University of Aveiro, Portugal; María Navarro Cáceres, University of Salamanca, Spain; Three anonymous reviewers

**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. The Artificial Intelligence (AI) community has developed analytical methods to uncover and predict patterns from posts that enable it 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 models from accurately encoding these posts' contextual knowledge. We propose a hybrid approach where knowledge encoded in community-maintained knowledge graphs (e.g., Wikidata) is combined with deep learning to categorize 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) to extract entities in short posts and link them to concepts in knowledge graphs. Then, knowledge graph embeddings (KGEs) are utilized to compute latent representations of the extracted entities, which result in vector representations of the posts that encode these entities' contextual knowledge extracted from the knowledge graphs. These KGEs 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 diagnose this type of mental disorder. We evaluate our approach on a dataset of 2,000 tweets about eating disorders.

---

\* Corresponding author. E-mail: [jbena@unileon.es](mailto:jbena@unileon.es).

Our experimental results suggest that combining contextual knowledge encoded in word embeddings with the one built from knowledge graphs increases the reliability of the predictive models. 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 personalized medicine.

Keywords: Name entity linking, deep learning, natural language processing, health data, knowledge graphs, Wikidata

## 1. Introduction

The COVID-19 pandemic has considerably burdened mental diseases all over the world [60], and eating disorders (EDs) are not an exception [16,46,62,81]. EDs are health conditions severely disturbing eating behaviors and related thoughts and emotions. A recent study by Zipfel et al. [81] reveals that EDs increased during the pandemic by 15.3% in 2020 with respect to previous years. Moreover, outcomes from a systematic literature review by McLean et al. [43] uncover that children and adolescents are the most vulnerable groups impacted by EDs during the COVID-19 pandemic. This burden of mental incidences raises awareness of the need for early detection mechanisms to effectively take action in clinical and healthcare services. More importantly, these studies provide evidence of the urgent need for scalable methods for effectively supporting communities increasingly suffering from mental disorders.

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 [64]. Moreover, social media networks are utilized to promote or prevent the administration of certain interventions, e.g., in mental conditions like eating disorders [9,39]. Furthermore, Artificial Intelligence (AI) has gained momentum in the healthcare field [8,80], and AI-based solutions have been developed for disease prevention [17], pathology detection [73], and treatment prescription [57]. Analyzing the discourse on social networks such as Twitter can help to find answers to relevant problems by applying various ML techniques. Specifically, in mental health, predictive models have successfully been applied [12,63]. Exemplary contributions include the detection of depression [29], suicidal mental tendencies [21], bipolar disorders [24] and other mental disorders [20]. Moreover, the models proposed by [38] can uncover patterns of anorexia in datasets obtained from social networks. These results put in perspective the role of automatic classification in the scalable and effective detection of patterns of communities suffering from mental disorders [75].

Knowledge graphs (KGs), and Semantic Web technologies in general, have been accepted as data structures that enable the natural representation and management of the convergence of data and knowledge [22]. The information contained in knowledge graphs is increasingly used in the scientific community to solve different problems [1,3,67]. Specifically, community-maintained knowledge graphs such as Wikidata [71] or DBpedia [33] represent rich sources of structured knowledge not only from general domains but also in biomedicine [10,30,79]. Figure 1(b) depicts the contextual knowledge obtained through structured data in a portion of Wikidata. We define this contextual knowledge as entities related to a given resource within the knowledge graph, e.g., the Wikidata resource Q254327 represents the concept “anorexia”, that is “symptom” and “physiological condition” with “decreased appetite”. This contextual knowledge can be interpretable by both humans and machines. This symbolic depiction of anorexia can be encoded in a subsymbolic representation using embeddings computed by Knowledge Graph Embedding methods [19,50,68]. Obtaining knowledge graph embeddings (KGEs) allows for calculating distance or similarity metrics and discovering relatedness among the entities that compose a knowledge graph (shown at the bottom of Fig. 1(b)). 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. Capturing contextual knowledge via semantic enrichment provides the basis for precise and accurate predictive models. This is of paramount relevance for ensuring the impact and credibility of the ED patterns detected from social media posts using AI methods.

Predictive models can also be built over unstructured data, Bidirectional Encoder Representations from Transformers (BERT) models [15] are exemplary solutions to this problem [2]. Contrary to other models (e.g., BiLSTM-based-Bidirectional Long Short-Memory), BERT models can learn from words in all positions, i.e., from the whole sentence. Figure 1(a) illustrates an example of contextual knowledge extracted from unstructured information for

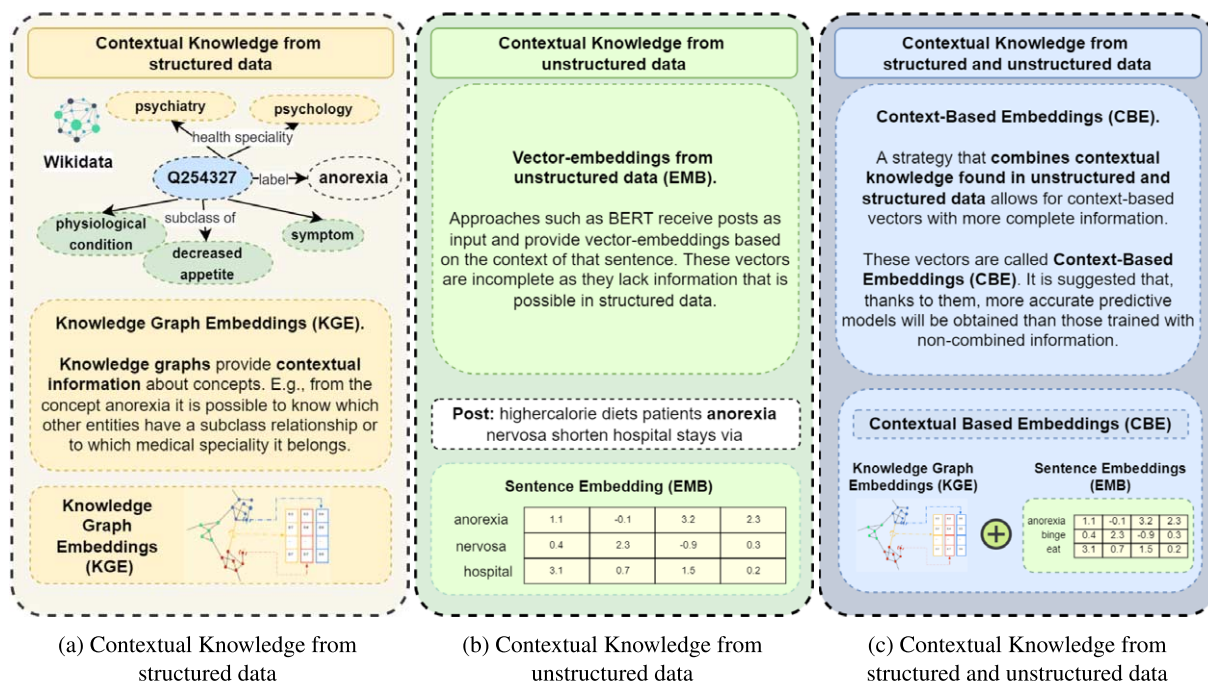


Fig. 1. Contextual knowledge (CK). (a) CK is represented in knowledge graphs and encoded in vector embeddings KGE methods. (b) CK is extracted from unstructured data, e.g., social media posts; it corresponds to the words around a particular concept. (c) The proposed approach, CK extracted from both structured and unstructured data and encoded in context-based embeddings (CBE).

the term “anorexia”. Here, this contextual knowledge is based on the information around the word anorexia within a given sentence. These models usually receive texts and labels as input to achieve predictive models that classify texts. By applying this model, vector embeddings are also obtained, but in this case, the encoded contextual content encodes unstructured data (lower part of Fig. 1(a)). BERT models have exhibited high performance in almost any prediction problem where contextual knowledge is extracted from text. Nevertheless, albeit the large number of covered domains and high-quality predictions, BERT models may perform poorly over short texts [35,78].

**Problem Statement and Objectives.** This paper addresses the problem of effectively classifying short posts referring to EDs. Our main research objective is to generate vector-embedding representations that encode contextual knowledge reported in structured data structures from knowledge graphs and unstructured corpora (e.g., scientific publications or social media posts). Concretely, we aim at encoding richer contextual knowledge, as depicted in Fig. 1(c). As a result, our goal is to generate vector-embeddings obtained from structured and unstructured data. The result is called contextual-based embeddings (CBEs). We hypothesize that CBEs will empower machine learning models in a way that the accuracy of pattern detection from social media posts is enhanced.

**Proposed Solution.** We propose a hybrid framework that combines vector embeddings generated from methods that extract contextual knowledge from various types of data sources (Fig. 1(c)). As a proof of concept, we show the benefits of the proposed approach while combining word embeddings from pre-trained BERT models and knowledge graph embedding methods learned from Wikidata. Although our approach is domain-agnostic, we will illustrate its performance in the problem of classifying Tweets according to mentions of eating disorders.

The resulting embeddings are used in various state-of-the-art predictive models. The methodology followed to implement the proposed approach is composed of the following steps: **1) Contextual-based knowledge from unstructured data:** Vector-embeddings are obtained from unstructured data by applying pre-trained BERT models. **2) Contextual-based knowledge from structured data:** RDF2Vec [56] is utilized to traverse the KGs, e.g., Wikidata. The resources to be traversed are extracted from social posts using Named Entity Recognition (NER) and Named Entity Linking (NEL). Thus, an entity embedding encodes the contextual knowledge represented in the traversed neighborhood of the entity which is recognized in the post. **3) Combined contextual knowledge:** Contextual-based

embeddings (CBEs) are computed by combining BERT and RDF2Vec embeddings. **4) Creation of Predictive Models:** CBEs are utilized to train and test the predictive models in a given classification task. The corpus of social media posts, annotated with labels, is also provided as input. In this paper, we have evaluated the proposed framework in a corpus of 2,000 Tweets related to eating disorders. This corpus has been created, curated, and annotated by the authors. The quality of the combined embeddings has been empirically evaluated in 12 state-of-the-art predictive models. The observed results indicate that considering the CBE increases the performance of all the models in accuracy and F-measure. Albeit limited to one corpus, these experimental outcomes provide evidence of the impact that contextual knowledge has on the accuracy of classification tasks conducted by existing methods.

**Contributions.** In summary, the scientific contributions of this work are as follows:

- **A corpus of Tweets related to Eating Disorders.** The dataset was collected using Twitter’s API and tagged in four categories. The corpus is composed of 2,000 tweets from which a total of 1,358 different entities are recognized. For these entities, their sentence embedding vectors from text data are calculated using Word2Vec, Tokenisers, and BERT and knowledge graph embeddings are generated using RDF2Vec [56]. The dataset quality is analyzed in terms of diversity and potential bias.
- **A Contextual-based Method to Empower Predictive Models.** The proposed method computes vector embeddings from structured and unstructured data and generates embeddings that integrate the two perspectives. This method is applied to the corpus of Tweets to solve four classification tasks related to eating disorders.
- **Experimental Assessment of the Proposed Methods.** An empirical evaluation is performed on the corpus over 12 different predictive models. Four binary classification tasks are solved. The results indicate that models that use contextual knowledge from structured and unstructured data in combination perform better in 97.5% of the cases, an improvement of up to 15%. The dataset and the code with the testbeds are available in a public repository, thus, ensuring reproducibility [7].

The remainder of this paper is organized as follows. Section 2 presents the fundamentals of RDF, knowledge graphs, knowledge graph embeddings, entity recognition, and entity linking. Section 3 analyzes related approaches from the state of the art. Then Section 4 defines our problem statement and states our proposed solution. This section also describes the main components of the architecture that implements our proposed method. Section 5 reports on the experimental evaluation and discusses the observed outcomes. Finally, we close the paper in Section 6 with conclusions and an overview of future work.

## 2. Preliminaries

### 2.1. Wikidata – Community-maintained knowledge graph

In this research, we use community-maintained 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. Wikidata can be downloaded in RDF, however, a subject-property-object triple is annotated with qualifiers [41], representing metadata about the triple. Lastly, although Wikidata is a source of encyclopedic knowledge, Waagmeester et al. [72] report the main characteristics of Wikidata as a relevant provider of knowledge from Life Sciences.

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

In our approach, we use named entity recognition and linking tools to extract the entities found in the short texts used as input data and link them 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<sup>1</sup> contained in the spaCy

---

<sup>1</sup><https://spacy.io/api/entitylinker>

[26] Python library and Falcon 2.0 [58]. Both tools resort to background knowledge to perform the linking process. EntityLinker's knowledge base is updated with Wikidata resources from 2019,<sup>2</sup> while Falcon 2.0 has a background knowledge updated to 2021. The two systems are used because the recognized entities were not always the same, and it was observed that a union of the entities obtained using both tools is more complete.

### 2.3. Knowledge graph embeddings

Knowledge graph embeddings (KGE) are low-dimensional representations of the entities and relationships that make up a knowledge graph. KGEs make possible a generalizable representation of an entity based on its context on a global knowledge graph, in our case, Wikidata. This context allows inferring relationships between concepts. The KGEs provide us with information about the relationships between different terms related to our problem. For example, in an ED context, KGEs provide us with 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 found 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 use of RDF2Vec [56]. Inspired by Word2vec [44] 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.

### 2.4. Extracting knowledge graph embeddings through RDF2Vec

RDF2Vec, as proposed by Ristoski et al. [56], adapts the language modeling approach in Word2vec for latent representations of entities in RDF graphs. Ristoski et al. [56] demonstrate that projecting such latent representations of entities into a lower dimensional feature space shows that semantically similar entities appear closer to each other in the knowledge graph where these entities are represented. The following steps are proposed to generate RDF2Vec embeddings from a corpus of short texts:

1. The entities of each of the texts in the dataset are obtained, as explained in Section 2.2.
2. The set of unique entities is used to generate a vocabulary.
3. An extension of RDF2Vec (implemented by the authors) is used to generate knowledge graph embeddings of a set of entities. This procedure calculates the embeddings based on the entities that are in the neighborhood of each entity within Wikidata. This enables us to obtain word vectors based on the contextual knowledge within the entities found in the collected dataset.

### 2.5. Strategies for combining embeddings

In the proposed 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 [4,34,52,54], 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 smooth inverse frequency (SIF) [4]. SIF also considers the most meaningful words within a sentence; it scores from 0 to 5 pairs of sentences, in our case tweets, according to the similarity between them. SIF is used as it is one of the most widely used strategies [49,76].

---

<sup>2</sup><https://www.kaggle.com/kenshoresearch/kensho-derived-wikimedia-data>



### 3. Related work

#### 3.1. Deep learning in social media data

The use of social media data to train predictive models that use 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.<sup>3</sup> Some of these studies focus on generating predictive models that are able to classify texts within a given domain, such as mental disorder [38], emotions [51], or fake news [77]. These studies resort to 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.

#### 3.2. Knowledge acquisition from knowledge graphs

Knowledge graphs such as Wikidata [71] and DBpedia [33] are increasingly used in research. Moreover, studies have shown that the information they contain, is useful and reliable for use in some fields like the Life Sciences [72]. One of the most important problems scientists face in acquiring knowledge 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<sup>4</sup> or whether we are referring to the studio album called “Help!” by The Beatles.<sup>5</sup> Due to the growing use of these knowledge graphs for research, 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 [58] and the EntityLinker<sup>6</sup> function from the spaCy [26] library that is implemented and available for use in Python. Thanks to these tools, obtaining information associated with the concepts contained in texts is easier and faster. The extracted linked entities can be used to perform different operations, for example, to calculate the corresponding embeddings for these entities. These calculations can be done using methods such as RDF2Vec [56].

#### 3.3. Combining knowledge graphs and deep learning

The contextual knowledge found in knowledge graphs combined with the use of deep learning techniques is being used to solve different research problems. For example, thanks to the semantic information obtained through these knowledge graphs, it is possible to improve the interpretability and explainability of different predictive models powered by deep learning techniques [36]. Improving the interpretability and explainability of predictive models is important, especially when deployed in domains related to health or education. Moreover, knowledge graphs are also being used to help predict relationships between different concepts, such as, the gut microbiota and mental disorders [18]. Notwithstanding, there is a lack of studies that have used the contextual knowledge extracted from knowledge graphs combined with the contextual knowledge extracted from unstructured data to classify texts by making use of this information. Our study aims to demonstrate that incorporating semantic enrichment to texts can contribute to improve the performance of predictive models.

---

<sup>3</sup>[https://scholar.google.es/scholar?as\\_ylo=2017&q=social+media+data+AND+deep+learning&hl=es&as\\_sdt=0,5](https://scholar.google.es/scholar?as_ylo=2017&q=social+media+data+AND+deep+learning&hl=es&as_sdt=0,5)

<sup>4</sup><https://www.wikidata.org/wiki/Q1643184>

<sup>5</sup><https://www.wikidata.org/wiki/Q201816>

<sup>6</sup><https://spacy.io/api/entitylinker>

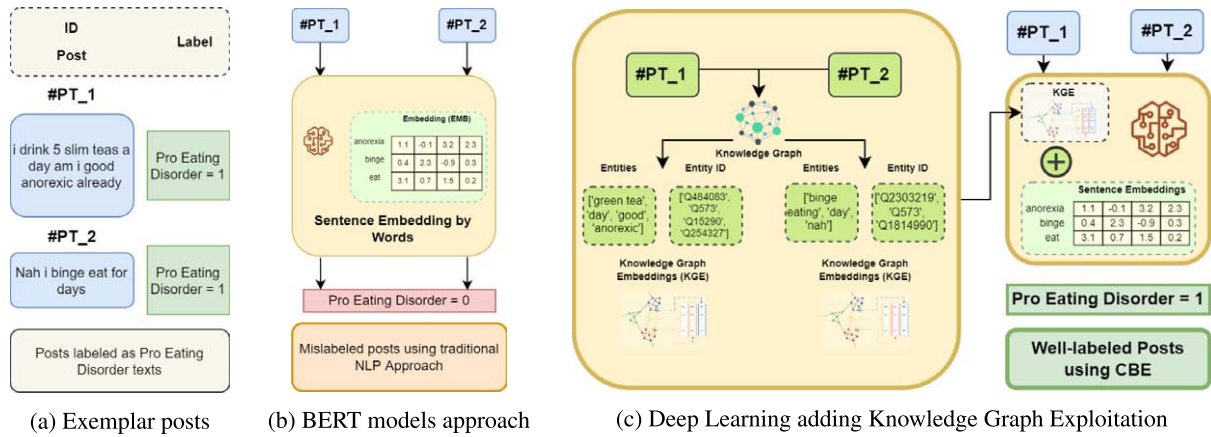


Fig. 2. Motivating example. (a) 2 labeled posts, (b) a text classification model using BERT models approach, (c) short text classification using the novel approach making use of context-based embeddings (CBE).

#### 4. Combining contextual knowledge from unstructured and structured data

In this research, we address the problem of effectively classifying short text. We propose a hybrid method based on the combination of contextual knowledge from structured data and unstructured data, resulting in contextual based embeddings. We elaborate below on the problem statement and the proposed solution.

**Problem Statement.** Machine learning models have shown to underperform when applied to short text classification problems. As depicted in Fig. 2, NLP-based approaches, for example, those that incorporate BERT models, continue to demonstrate accuracy problems, and are capable of mislabeling short text posts in binary classification tasks. This problem needs to be addressed as datasets elaborated from short posts obtained from social media, e.g., Twitter, are widely used.

**Proposed Solution.** The approach proposed in this work resorts to contextual knowledge extracted from unstructured data combined with contextual knowledge extracted from structured data to improve the performance of predictive models for short text classification. Vector-embeddings are obtained based on the context of a sentence (unstructured data); secondly, knowledge graph embeddings are obtained (e.g., using RDF2Vec), which represent the contextual knowledge of structured data. Figure 3 depicts the proposed architecture of the empowered model; concisely, it receives as input a corpus composed of short texts and the corresponding entities extracted from knowledge graphs. Both of these are transformed into embedding vectors to be combined. The output is the classification generated by a predictive model after a process of context-based embeddings encoding and decoding.

##### Component 1: Context-Based Embeddings Encoding.

This component has three main objectives: (i) to obtain contextual knowledge through unstructured data (sentence embeddings), (ii) to obtain contextual knowledge through structured data (knowledge graph embeddings), and (iii) to combine this textual knowledge into a single vector structure through the combination of both (context-based embeddings). Objectives (i) and (ii) are achieved through modules that are executed in parallel. Objective (iii) is attained once the two previous modules have been executed. The properties of the three modules that make up this component are detailed below.

**Unstructured based data.** This module generates a vector-embedding for each word within a post. This vector represents the similarity of the terms within the corpus of posts. For example, the term “anorexia” within the sentence “higher-calorie diets patients anorexia nervosa shorten hospital stays via” is contextually related to the terms around it. A word-embedding encodes the similarity values of the terms within the corpus of posts. Once all the word-embeddings of a sentence are computed, a vector-embedding is generated; it represents the complete sentence, combining all the vectors of a sentence. This vector-embedding contains contextual knowledge of the posts. Thus, the closeness between terms is encoded, e.g., “anorexia” and “diets”. These vectors can be obtained using, e.g.,

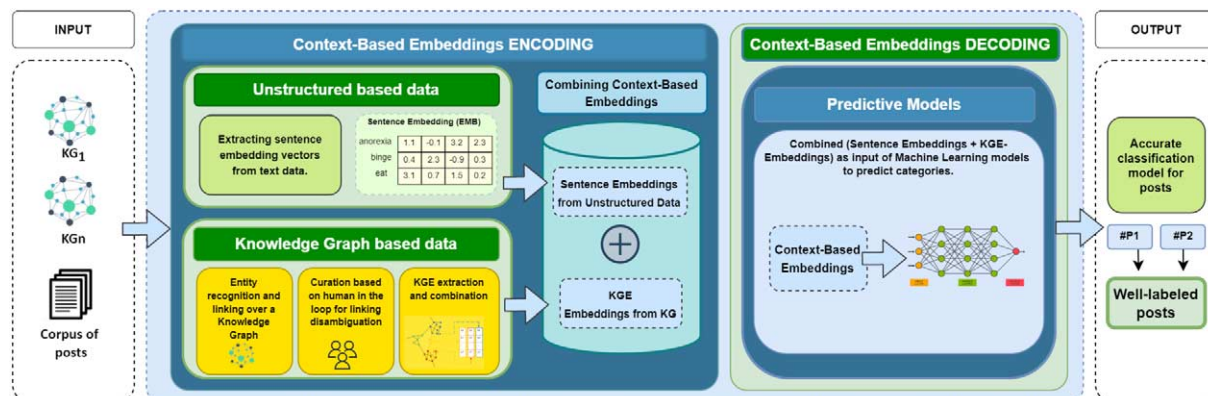


Fig. 3. Architecture of the proposed approach. It receives a corpus of posts and knowledge graphs. The first component, called context-based embeddings encoding, represents how unstructured based data is combined with the information obtained through knowledge graphs applying entity recognition and linking over concepts within a knowledge graph. The knowledge graph-based data module shows the process of knowledge graph embeddings (KGEs) extraction from a knowledge base. Then, sentence embeddings from unstructured data are combined with KGE. The second component, called context-based embeddings, is computed, and predictive machine learning models are executed with these embeddings in decoding, resulting in accurate classification models for posts.

Tokenisers, Word2Vec or BERT. Although this combined vector represents the words of a sentence, it does not include knowledge modeled in structured data.

**Knowledge Graph-based data.** The main objective of this module is to obtain knowledge graph embeddings for each post; these vectors will enrich the ones generated from the unstructured data. Three steps are performed:

(I) *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 linking the recognized entities to a knowledge graph; these engines are known as named entity recognizers and linkers (e.g., Falcon 2.0 [58] and EntityLinker in the spaCy [26] Python literature). In the current version of this module, the recognized entities are linked to Wikidata. The Falcon 2.0 API<sup>7</sup> extracts the entities of a sentence, as well as the resources that correspond to these entities in Wikidata and DBpedia. The following example illustrates the tasks of entity recognition and linking over the following sentence:

*higher-calorie diets patients anorexia nervosa shorten hospital stays via*

The recognized entities are:

['kilocalorie', 'diet', 'patient', 'anorexia', 'nervosa']

And the identifiers of these entities in Wikidata are:

['Q26708069', 'Q474191', 'Q181600', 'Q254327', 'Q131749']

(II) *Curation based on human in the loop for linking disambiguation.* Entity recognizers and linkers may be inaccurate. For example, the word “Help” can represent the action of help or support to someone, or it could be a music album called “Help!”. In order to assist this type of disambiguation, we resort to a list of *tabu types* to help determine when the linked resource needs to be manually validated and curated. This list includes the following types: *Album, Book, Streets, Organization, Song, and Movie*. This process enhances the quality of the description of a post in terms of entities in a knowledge graph.

(III) *KGE – Embeddings Extraction and Combination.* Knowledge graphs contain contextual knowledge of entities about concepts. In this step, knowledge graph embeddings are obtained for each entity of each post present. For example, the concept “anorexia” is represented in Wikidata as resource Q254327, which is a “symptom” and “physical condition” with “decreased appetite”. The symbolic representation of anorexia in Wikidata is encoded in a sub-symbolic representation or knowledge graph embedding. RDF2Vec [56] or SDM-RDF2Vec [61] can generate

<sup>7</sup><https://labs.tib.eu/falcon/falcon2/api-use>



these vectors. After obtaining each knowledge graph embedding for a post's entities, a combination of these vectors, which have the same size, is generated. As a result, a knowledge graph embedding for each post is generated.

**Combining Context-Based Embeddings.** Once the contextual knowledge has been obtained from the unstructured based data and knowledge graphs data modules, this third module is in charge of combining both vectors. Combining these two vectors allows the complete contextual knowledge of each concept to be represented by a single vector. For example, in the initial sentence, we used “highercalorie diets patients anorexia nervosa shorten hospital stays via” we would have a single vector called context-based embedding. This vector would contain the information of the post obtained from its context as unstructured data, but it would also contain knowledge of the neighborhood of the words that compose it. The vector would represent the similarity between each post in the corpus from two different contextual knowledge. This combined vector contains the complete contextual information of each post.

### **Component 2: Context-Based Embeddings Decoding.**

This component uses the information contained in the context-based embeddings generated by the first component by decoding them. According to the hypothesis put forward in this study, the complete contextual knowledge contained in these vectors should provide predictive models with a higher degree of accuracy. This knowledge allows the vectors to contain a more accurate inter-post similarity function than vectors generated from unstructured data alone or from unstructured data. In the running example, “anorexia” was close to “nervosa” in the vector obtained from the unstructured data, but not close to “symptom”. The combination of both types of embeddings encodes different contextual knowledge; it offers a contextual-based enhanced vector.

**Predictive models.** In this module, different machine learning models are trained. The result is empowered models that outperform traditional models, i.e., have higher accuracy, when classifying a corpus of text. This decoding methodology can be applied not only to text classification problems, but also to other supervised and unsupervised learning problems.

## **5. Experiments and results**

The experiments conducted in this study demonstrate that machine learning models enhanced with contextual knowledge can improve accuracy in predictive short text classification models. The objective of our particular case study is to improve the classification of eating disorders in social media posts. The empirical study sets out to answer the following research questions:

- RQ1)** What is the impact of contextual knowledge extracted from structured data on the performance of the model?
- RQ2)** Do BERT pre-trained models based on short texts perform better than other pre-trained models?
- RQ3)** What type of contextual information provides the most accurate predictive models – structured, unstructured or a combination of both?

Data availability of the experiments and code of the empirical evaluation of the proposed approach<sup>8</sup> are publicly available; this facilitates the reproducibility of the reported results.

### *5.1. Experimental setup*

The following settings are configured to answer our research questions.

**Benchmarks.** The experiments are executed using a proprietary eating disorders dataset. The dataset was elaborated by collecting tweets from Twitter and contains 2,000 posts labeled as detailed in the Section 5.3. The dataset is divided into training and test sets to train the different predictive models. This separation is performed using a stratified separation function, in a proportion of 70% for training and 30% for testing. This way, it is possible to ensure that the categories are balanced in both sets and to avoid possible overfitting or under fitting of the models.

---

<sup>8</sup><https://github.com/knowledgeb/Combining-Knowledge-Graphs-and-Deep-Learning-techniques-for-Categorizing-Tweets>

Table 1  
Formulas of the evaluation metrics used

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	$Acc = \frac{TP+TN}{(TP+TN+FP+FN)}$

*Cross-validation.* A k-fold cross-validation has been carried out with 10 folds for all experiments to avoid randomness. Every model has been evaluated through an extensive hyperparameter grid search. The results show the score for the hyperparameter configuration with the best mean value of the ten folds.

**Metrics.** To evaluate the results obtained by each of the models, two different metrics have been used:  $F_1$ -score ( $F_1$ ) and accuracy (acc) (Table 1).  $F_1$ -score is a metric used to calculate the effectiveness of a classifier by taking into account its accuracy and recall values.  $F_1$  assumes that the two metrics used are equally important in calculating effectiveness. If one of them is more important, a different formula  $F_\beta$  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).

**Implementations.** The proposed approach is implemented in Python 3.6. having CUDA 11.1 installed. The following libraries are required in the indicated versions: numpy v1.19.5, params v0.9.0, pandas v1.2.2, pyrdf2vec v0.2.3, qwikidata v0.4.0, spacy v3.4.1, tensorflow v2.4.0, torch v1.10.2, tweet-preprocessor v0.6.0, scikit\_learn v0.24.2 and simpletransformers v0.60.9. The experiments are executed on a Windows 10 Pro x64 machine with an Intel® Core® i7-9700K CPU @ 3.60 GHz (eight physical cores, eight threads) and 32 GiB DDR4 RAM with a graphic card NVIDIA GeForce RTX 2080 SUPER (8 GB GDDR6).

**Engines.** Falcon 2.0 [58] and EntityLinker [26] are used to perform entity recognition and linking. Moreover, knowledge graph embeddings are computed with pyRDF2vec [69]; it is configured as follows: *RandomWalker* with *max\_depth=4* and *max\_walks=50*.

**Models.** 12 different machine learning models are executed to classify the four binary categories. The classification models used are: Random Forest Classifier [13], LSTM [27], Bi-LSTM [28], CNN, CNN + LSTM [80] and 7 pre-trained BERT models [15]. The pre-trained BERT models are as follows: TweetBERT [45], BERT [14], RoBERTa [37], DistilBERT [59], CamemBERT [42], Albert [31], and FlauBERT [32]. These models receive three inputs: (i) the texts of the posts, (ii) the structured information obtained from knowledge graphs, and (iii) the combination of structured and unstructured information in the form of context-based embeddings (CBE). These models are evaluated using the metrics explained above.

## 5.2. Experimental methodology

The different steps taken to validate the model proposed in this research using a dataset of our own are detailed below. The working pipeline is divided into 4 phases:

1. **Collection and labeling of ED data:** in this phase, texts are collected through web scraping on the social network Twitter. After this collection, the tweets are labeled in the four binary categories of study.
2. **Analysis of the data collected:** once the data has been tagged, an analysis of the texts obtained is carried out. This analysis provides information on the context of the texts obtained, and also on the importance of emojis in the dataset. This analysis allows us to choose pre-processing strategies to avoid possible biases in the experiments.
3. **Knowledge acquisition through knowledge graphs:** once the labeled dataset is curated, information is obtained through knowledge graphs, in this case, Wikidata. In this step, the knowledge graph embeddings of the texts are obtained using RDF2Vec [56]. Then the work of combining these embeddings with the embeddings obtained in the texts is carried out, giving rise to context-based embeddings.

4. **Application of machine learning models:** in this phase, different machine learning models are applied using the texts, the knowledge graph embeddings, and the context-based embeddings as input. In this way, it is possible to compare the accuracy results of these models using traditional approaches and the approach we propose in this research. In addition, to verify that our model is not dependent on the dataset used, an analysis has been carried out on the input dataset and on the output obtained after applying the different models.

The phases mentioned above are explained in detail in the following subsections.

5.3. *Collecting and labeling ED dataset*

Figure 4 depicts the process followed to collect and label the tweets that compose our corpus.

We employ information-filtering-based methods of social media posts using a set of keywords as search queries. This method of collection 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 [48,74]. Nonetheless, the datasets produced through this way come with their limitations due in great part to the pervasive noise associated with their source, which may impact the quality of the data itself, as well as the reliability and representativeness of the results obtained [5,48].

Our chosen social media platform is Twitter, and we perform the capture of tweets using the T-Hoarder tool [11]. The list of hashtags used to filter tweets is anorexia, anorexic, dietary disorders, inappetence, feeding disorder, food problem, binge-eating, eating disorders, bulimia, food issues, and loss of appetite. These hashtags were selected taking into account the most relevant keywords that yielded tweets about eating disorders in other research [23,25,40,65]. However, this sub-selection could still bias the sampling and results, as we might be making assumptions about the behavior of the users generating the content or overlooking relevant information outside these limits [48]. The collection yielded the capture of 494,025 tweets, and to mitigate issues associated with lexical and semantic redundancy of the content [48], a process of cleaning (removal of duplicates, re-tweets, re-shared content), manual curation and annotation were performed, thus, resulting in the creation of a subset of 2,000 tweets.

To annotate the tweets, we first determine four binary classification problems of interest guided by preexisting ED research [25,40,65]. The resulting categories for each problem are detailed below, followed by the labeling criteria of each tweet:

- **ED I:** Tweets written by people with an ED (1) or not (0). For each tweet, the author’s profile was accessed. All those who indicated in their biography to suffer from an ED and/or wrote tweets explicitly indicating that they suffered from an ED were put into the positive class (1).
- **ED II:** Tweets that promote having ED (1) or not (0). In the community, there are terms that expose a person to be in favor of having an ED, for example: proana (proanorexia) or promia (probulimia) [70]. Tweets labeled as promoting EDs were written by authors expressing these terms in their biographies or timelines.
- **ED III:** Tweets of an informative nature (1) or not (0). Tweets that contained information on an ED-related subject were put into the positive class (1). Content that reflected opinions was put in the negative class (0).

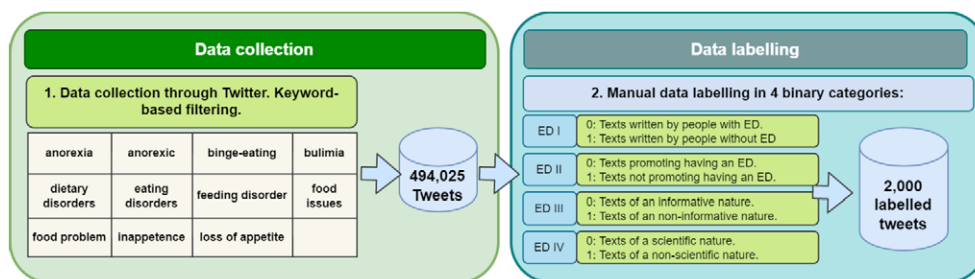


Fig. 4. Data collection and data labeling steps. First, a set of tweets were collected through a keyword-based search. Subsequently, a subset of the tweets was manually labeled into four binary categories.

- **ED IV:** Tweets of scientific nature (1) or not (0). Tweets written by ED health professionals and researchers disseminating relevant information, data, and results were categorized as scientific. Tweets linked to an article in a scientific journal were also categorized as scientific.

This dataset is available for download from github [7].

#### 5.4. Dataset analysis

As per 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 [47,53]. 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 [6,48].

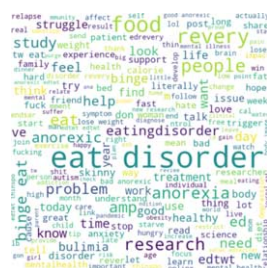
Through an exploration of the frequency of terms associated with the hashtags collected, we managed to identify a predominant mention of anorexia and other derivative terms such as “anorexia”, “anorexic”, “proana”, “ana”, “anatwt”, “anorexiatips”, (1,034 associated hashtags), while an underrepresentation 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 behavior 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 of this exploration 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 5 presents a summary of the data analysis. Most frequent terms with unigram tokens (see Fig. 5(b) and Figure 5(a)) and a breakdown of the annotated categories and the corresponding class count is showed in Fig. 5(c). The top 35 hashtags show a predominant mention of ‘anorexia’ and other derivative terms such as “anorexia”, “anorexic”, “proana”, “ana”, “anatwt”, “anorexiatips”.

Emojis were eliminated from the tweets in the corpus. The following analysis provides evidence of the lack of effect of the emojis’ removal.

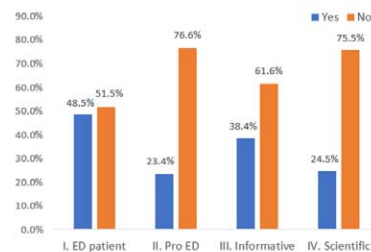
- The frequencies of the different emojis in the texts were obtained, and it was determined that only 17.9% of texts contained some emoji, i.e., 359 of the 2,000 total texts.
- Statistical calculations were carried out on the distributions of emojis according to the four categorizations made. To statistically analyze these data, two statistical techniques were used: (i) overlap analysis and (ii) Spearman’s rho correlation.

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

(a) Most frequent terms in the dataset.



(b) Frequent unigram Word Cloud.



(c) Distribution of data by category.

Fig. 5. Results of the analysis of the dataset, showing the most frequent terms and the distribution of the categories.

Table 2

Frequency of emojis in each of the four categories, and Spearman correlation between the distributions for each category

Categorization	Frequency (label 0)	Frequency (label 1)	Spearman results
ED I	111	248	$r(163) = 0.85, p < 0.001$
ED II	231	128	$r(163) = 0.96, p < 0.001$
ED III	285	74	$r(163) = 0.57, p < 0.001$
ED IV	326	33	$r(163) = 0.82, p < 0.001$

Table 3

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

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

- \* **Overlap analysis:** An overlap analysis compares the values between columns in one table or across tables. This analysis allows us to identify overlapping or redundant data in columns. In our particular case, we compared whether or not the number of emojis found in tweets that have been tagged with 0 in each of the binary categories overlapped with the emojis found that have been tagged with the value 1 in each category.
- \* **Spearman's rho correlation:** This analysis makes it possible to determine whether two data distributions are correlated. The analysis of emojis allows us to determine whether the distribution of emojis found in the tweets categorized in each binary class of the four categorizations performed is significantly similar or not. That is, it helps us determine whether the number of emojis found in tweets labeled as "tweets written by ED patients" is similar, statistically speaking, to the number of emojis found in tweets labeled as "tweets not written by ED patients".

It was observed that, in all cases, the similarity between the two distributions was statistically significant. The correlations are shown in Table 2. Note that in category ED III, the Spearman's rho value is lower because 8 emojis were used more than 90% of the time in tweets categorized as non-informative, compared to those categorized as informative. This difference makes sense, since, when an opinion is expressed, there are emojis that relate more to the feelings of that opinion, and informative tweets are more objective.

- Therefore, in line with what has been done in other similar research [55,66], it was decided to eliminate emojis in the pre-processing of the texts.

### 5.5. Knowledge acquisition through knowledge graph

The texts of the dataset were processed with Falcon 2.0 [58] and Entity Linker [26], resulting in a total number of entities in the Wikipedia knowledge graph, as shown in Table 3.

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 (Table 4(a)).

In addition, several concepts were added to Wikidata, shown in Table 4(b). After obtaining all the entities, the knowledge graph embeddings were collected using the pyRDF2Vec tool [69] (setting *RandomWalker* with *max\_depth=4* and *max\_walks=50*) and the SIF algorithm is applied to combine the knowledge graph embeddings of each tweet into a single embedding. 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 (knowledge graph embeddings).

### 5.6. Experiments

To validate the proposed approach, 144 tests are carried out, from which the evaluation metrics indicated above, accuracy and f1-score, are obtained. These 144 tests are divided as follows:



Table 4

Experimental tables. (a) Terms tagged with misconceptions on Wikidata by Falcon [58] and EntityLinker [26]. (b) News terms added about eating disorders added to Wikidata

(a) Terms tagged with misconceptions on Wikidata			(b) New terms added to Wikidata	
Term	Wrong link	Correct link	Concept	Wikidata ID
recovery	Recovery, music album (Q274533)	Recovery approach (Q2135807)	Food avoidance emotional disorder	Q108760799
anorexia	Anorexia, music album (Q4770169)	anorexia, medical symptom (Q254327)	fatspo	Q111780867
ed	Ed, tv serie (Q930797)	eating disorder (Q373822)	Addictive Eaters Anonymous	Q111781180
binger	binger, town (Q544455)	binge eating (Q2303219)	Meanspo	Q111781194
help	The Help, a film (Q204374)	help, cooperation (Q1643184)	Ultra-processed food	Q111781198

- Three inputs, one for each approach:** Three input datasets are applied, one for each of the approaches mentioned in this study: (i) texts, which allow obtaining the results after analyzing contextual knowledge using unstructured information; (ii) knowledge graph embeddings, which allow obtaining classifications based on contextual knowledge based on structured data; and (iii) using context-based embeddings obtained through the proposed approach.
- Four binary classifications:** Each machine learning model is trained four different times with each of the three inputs in order to classify the four categories mentioned in this study: ED I (tweets written by ED patient, or not), ED II (tweets promoting having an ED or not), ED III (informative tweets or not) and ED IV (scientific tweets or not).
- Twelve machine learning models:** The 12 models indicated in the experimental setup for the classification of the four categories indicated are run with the three different data inputs. This gives a total of 144 different outputs, 36 for each binary category to be classified, 12 with each data input in each category.

### 5.7. Dataset properties in the classification outcome

To verify that our model is not dependent on the dataset used, an analysis has been carried out on the original dataset and on the outputs obtained after applying the different models. This analysis consists of obtaining the Spearman correlation between the distributions of the most frequent terms found in tweets labeled as “0” or “1” in each category in each of the tests performed. A correlation close to 1.0 with a p-value of less than 0.05 would indicate that the distributions of the most frequent terms analyzed are statistically similar, thus being a valid and unbiased dataset for training the models. Likewise, keeping the same distribution at the output may suggest that the models do not increase or decrease the possible bias that could be present in the dataset.

After analyzing the output of the 12 models, applied on the 3 different input datasets, in the 4 tested categorizations, it has been observed that the distributions of the most frequent terms in the dataset are similar in the 144 outputs of the analyzed algorithms with a p-value < 0.001 in all cases, the maximum correlation was 1.0 and the minimum 0.996. This means that the models do not amplify the possible bias of the dataset.

### 5.8. Discussion of observed results

Twelve different classification models to classify tweets in four binary categories are compared using three different approaches: (i) contextual knowledge extracted from unstructured data, (ii) contextual knowledge extracted from structured data (KGE), and (iii) the approach proposed based in contextual knowledge extracted from structured data combined with contextual knowledge extracted from unstructured data (context-based embeddings, CBE). The results are shown in Table 5.

These results are obtained after training twelve different machine learning and deep learning models to categorize tweets in different categories using three different inputs for our dataset about eating disorders which consist of: (i) 2,000 labeled tweets, (ii) 2,000 knowledge graph embeddings and (iii) context-based embeddings obtained through our proposed approach.

Table 5

Results after applying 12 machine learning models using 3 different input variables: texts, knowledge graph embeddings, and the context-based embeddings obtained using our approach. The best results are highlighted in bold

Model (Data)	ED I		ED II		ED III		ED IV	
	$F_1$	$acc$	$F_1$	$acc$	$F_1$	$acc$	$F_1$	$acc$
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 (CBE)	<b>0.860</b>	<b>0.860</b>	<b>0.927</b>	<b>0.878</b>	<b>0.874</b>	<b>0.843</b>	<b>0.933</b>	<b>0.902</b>
LSTM (KGE)	0.360	0.499	0.631	0.787	0.504	0.602	0.622	0.739
LSTM (Text)	0.801	0.794	0.831	0.812	0.792	0.770	0.839	0.801
LSTM (CBE)	<b>0.846</b>	<b>0.835</b>	<b>0.848</b>	<b>0.823</b>	<b>0.829</b>	<b>0.802</b>	<b>0.863</b>	<b>0.842</b>
Bi-LSTM (KGE)	0.211	0.499	0.656	0.787	0.215	0.602	0.432	0.739
Bi-LSTM (Text)	0.791	0.785	0.841	0.823	0.777	0.762	0.864	0.847
Bi-LSTM (CBE)	<b>0.835</b>	<b>0.842</b>	<b>0.853</b>	<b>0.841</b>	<b>0.835</b>	<b>0.811</b>	<b>0.893</b>	<b>0.873</b>
CNN (KGE)	0.421	0.435	0.601	0.611	0.545	0.581	0.495	0.505
CNN (Text)	0.720	0.745	0.803	0.811	0.712	0.732	0.768	0.771
CNN (CBE)	<b>0.821</b>	<b>0.832</b>	<b>0.821</b>	<b>0.833</b>	<b>0.805</b>	<b>0.817</b>	<b>0.834</b>	<b>0.842</b>
CNN + LSTM (KGE)	0.490	0.499	0.490	0.498	0.314	0.398	0.421	0.501
CNN + LSTM (Text)	0.824	0.831	0.748	0.797	0.715	0.770	0.656	0.750
CNN + LSTM (CBE)	<b>0.849</b>	<b>0.841</b>	<b>0.853</b>	<b>0.857</b>	<b>0.817</b>	<b>0.871</b>	<b>0.862</b>	<b>0.854</b>
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 (CBE)	<b>0.894</b>	<b>0.890</b>	<b>0.959</b>	<b>0.936</b>	<b>0.886</b>	<b>0.865</b>	<b>0.948</b>	<b>0.924</b>
BERT (KGE)	0.275	0.526	0.479	0.796	0.171	0.586	0.212	0.648
BERT (Text)	<b>0.876</b>	0.870	0.933	0.895	0.887	0.863	0.951	0.927
BERT (CBE)	0.875	<b>0.873</b>	<b>0.963</b>	<b>0.942</b>	<b>0.898</b>	<b>0.878</b>	<b>0.960</b>	<b>0.941</b>
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	<b>0.870</b>	0.846	0.947	0.924
CamemBERT (CBE)	<b>0.885</b>	<b>0.876</b>	<b>0.961</b>	<b>0.939</b>	<b>0.870</b>	<b>0.854</b>	<b>0.950</b>	<b>0.927</b>
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 (CBE)	<b>0.890</b>	<b>0.888</b>	<b>0.961</b>	<b>0.939</b>	<b>0.900</b>	<b>0.878</b>	<b>0.952</b>	<b>0.929</b>
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 (CBE)	<b>0.892</b>	<b>0.888</b>	<b>0.953</b>	<b>0.927</b>	<b>0.888</b>	<b>0.868</b>	<b>0.954</b>	<b>0.932</b>
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 (CBE)	<b>0.897</b>	<b>0.897</b>	<b>0.964</b>	<b>0.944</b>	<b>0.906</b>	<b>0.888</b>	<b>0.949</b>	<b>0.926</b>
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	<b>0.890</b>	0.868	0.953	0.931
TweetBERT (CBE)	<b>0.909</b>	<b>0.909</b>	<b>0.964</b>	<b>0.944</b>	0.888	<b>0.868</b>	<b>0.957</b>	<b>0.937</b>

The models trained with the context-based embeddings (CBE) 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.

**Answer to RQ1.**

From the analysis of the results, it is clear that the use of contextual knowledge extracted from structured 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 those with knowledge graph exploitation as input data 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 RQ2.**

Comparing the performance results obtained by each of the seven pre-trained BERT models, the results suggest that the pre-trained model with short texts, TweetBERT, is the one that offers the best results in the metrics evaluated against its opponents in all but one occasion. These results suggest that pre-trained models with short texts may perform better than others.

**Answer to RQ3.**

It is possible to highlight that the approach proposed in this research, based on the combination of contextual knowledge obtained through structured and unstructured information, offers better results than the other two approaches studied. These results suggest that the contextual content that can be obtained from structured information, combined with the textual content of unstructured information, can help us generate more accurate predictive models. This improvement in the accuracy of models may contribute to assist some tasks of early disease detection.

## 6. Conclusions and future work

We address the problem of improving the performance of predictive models generated through machine learning techniques for short texts classification proposing a hybrid framework based on the combination of contextual knowledge from unstructured and structured data sources. 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 due to these diseases.

This research presents a hybrid approach using a combination of contextual knowledge from unstructured data (e.g., BERT) and contextual knowledge from structured data (e.g., knowledge graphs), obtaining contextual knowledge embeddings. One of the objectives achieved with these embeddings is to improve classification models in short texts. It has been demonstrated on a real health-related problem over a dataset about eating disorders. The tests to validate the proposed approach show that the models trained with contextual-based embeddings have a higher success rate than those obtained with models trained only with contextual knowledge from structured or unstructured data. As a result, 97.5% of the trained models perform better with our approach than the usual ones. It is possible to highlight that using semantic information in knowledge graphs can improve the performance of predictive models in natural language processing and text mining, with the corresponding importance in the health sector. Thus, the proposed method contributes to the portfolio of tools to support understanding the disorders that social media users may suffer. Given the urgent need to help these communities, we hope these results will motivate using our proposed methods in other health-related classification problems.

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 by adding the information of emojis by combining other Knowledge Bases such as, e.g., Emojinet or Emojipedia. Finally, creating a framework that applies our approach to a given dataset is on our future agenda.

## Acknowledgements

Part of this research was funded by the European Union’s Horizon 2020 research and innovation programme under Marie Skłodowska-Curie Actions (grant agreement number 860630) for the project “NoBIAS – Artificial Intelligence without Bias”. This work reflects only the authors’ views, and the European Research Executive Agency (REA) is not responsible for any use that may be made of the information it contains. Furthermore, Maria-Esther Vidal is partially supported by Leibniz Association in the program “Leibniz Best Minds: Programme for Women Professors”, project TrustKG-Transforming Data in Trustable Insights with grant P99/2020.

## References

- [1] K. Abhishek, V. Pratihari, 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* **44**(9) (2022), 838–846, Publisher: Taylor & Francis \_eprint. doi:[10.1080/1206212X.2021.1957551](https://doi.org/10.1080/1206212X.2021.1957551).
- [2] F.A. Acheampong, H. Nunoo-Mensah and W. Chen, Transformer models for text-based emotion detection: A review of BERT-based approaches, *Artificial Intelligence Review* **54**(8) (2021), 5789–5829. doi:[10.1007/s10462-021-09958-2](https://doi.org/10.1007/s10462-021-09958-2).
- [3] S. Angioni, A.A. Salatino, F. Osborne, D.R. Recupero and E. Motta, Integrating knowledge graphs for analysing academia and industry dynamics, in: *Communications in Computer and Information Science* 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, Springer International Publishing, Cham, 2020, pp. 219–225. ISBN: 978-3-030-55814-7. doi:[10.1007/978-3-030-55814-7\\_18](https://doi.org/10.1007/978-3-030-55814-7_18).
- [4] S. Arora, Y. Liang and T. Ma, A simple but tough-to-beat baseline for sentence embeddings, in: *ICLR*, 2017.
- [5] 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](https://doi.org/10.1145/1526993.1526995).
- [6] R. Baeza-Yates, Bias on the web, *Commun. ACM* **61**(6) (2018), 54–61. doi:[10.1145/3209581](https://doi.org/10.1145/3209581).
- [7] J.A. Benítez-Andrades, jabenitez88/Empowering-Machine-Learning-Models- with-Contextual-Knowledge: r1.0 Semantic Web Journal, Zenodo, 2022. doi:[10.5281/zenodo.7101682](https://doi.org/10.5281/zenodo.7101682).
- [8] 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](https://doi.org/10.1109/CBMS52027.2021.00105).
- [9] 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. doi:[10.1080/09638237.2019.1677878](https://doi.org/10.1080/09638237.2019.1677878).
- [10] E. Cabrio, Artificial intelligence to extract, analyze and generate knowledge and arguments from texts to support informed interaction and decision making (2020), 253.
- [11] 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](https://doi.org/10.1016/j.jnca.2017.01.029).
- [12] M. Conway and D. O’Connor, Social media, big data, and mental health: Current advances and ethical implications, *Current Opinion in Psychology* **9** (2016), 77–82. doi:[10.1016/j.copsyc.2016.01.004](https://doi.org/10.1016/j.copsyc.2016.01.004).
- [13] 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](https://doi.org/10.1007/978-1-4419-9326-7_5).
- [14] 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>.
- [15] 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. <https://www.aclweb.org/anthology/N19-1423>. doi:[10.18653/v1/N19-1423](https://doi.org/10.18653/v1/N19-1423).
- [16] J. Feinmann, Eating disorders during the COVID-19 pandemic, *BMJ* **374** (2021), n1787. <https://www.bmj.com/content/374/bmj.n1787>. doi:[10.1136/bmj.n1787](https://doi.org/10.1136/bmj.n1787).
- [17] 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](https://doi.org/10.1016/j.cmpb.2021.105968).
- [18] 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](https://doi.org/10.1109/MIC.2020.3031769).
- [19] G.A. Gesese, M. Alam and H. Sack, Semantic entity enrichment by leveraging multilingual descriptions for link prediction (2020).

- [20] S. Graham, C. Depp, E.E. Lee, C. Nebeker, X. Tu, H.-C. Kim and D.V. Jeste, Artificial intelligence for mental health and mental illnesses: An overview, *Current Psychiatry Reports* **21**(11) (2019), 116. doi:10.1007/s11920-019-1094-0.
- [21] V.K. Gunjan, Y. Vijayalata, S. Valli, S. Kumar, M.O. Mohamed and V. Saravanan, Machine learning and cloud-based knowledge graphs to recognize suicidal mental tendencies, *Computational Intelligence and Neuroscience* **2022** (2022), e3604113, Publisher: Hindawi. <https://www.hindawi.com/journals/cin/2022/3604113/>. doi:10.1155/2022/3604113.
- [22] C. Gutiérrez and J.F. Sequeda, Knowledge graphs, *Communications of the ACM* **64**(3) (2021), 96–104. doi:10.1145/3418294.
- [23] J.K. Harris, A. Duncan, V. Men, N. Shevick, M.J. Krauss and P.A. Cavazos-Rehg, Messengers and messages for tweets that used #thin-spo and #fitspo hashtags in 2016, preventing, *Chronic Disease* **15** (2018), E01. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5757384/>. doi:10.5888/pcd15.170309.
- [24] D. Harvey, F. Lobban, P. Rayson, A. Warner and S. Jones, Natural language processing methods and bipolar disorder: Scoping review, *JMIR Mental Health* **9**(4) (2022), e35928, Company: JMIR Mental Health Distributor: JMIR Mental Health Institution: JMIR Mental Health Label: JMIR Mental Health Publisher: JMIR Publications Inc., Toronto, Canada. <https://mental.jmir.org/2022/4/e35928>. doi:10.2196/35928.
- [25] L. He and J. Luo, “What makes a pro eating disorder hashtag”: Using hashtags to identify pro eating disorder tumblr posts and Twitter users, in: *2016 IEEE International Conference on Big Data (Big Data)*, 2016, pp. 3977–3979. doi:10.1109/BigData.2016.7841081.
- [26] M. Honnibal and I. Montani, spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing, 2017, To appear.
- [27] L.C. Jain and L.R. Medsker, *Recurrent Neural Networks: Design and Applications*, 1st edn, CRC Press, Inc., USA, 1999. ISBN: 0849371813.
- [28] 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. <https://www.mdpi.com/2076-3417/10/17/5841>. doi:10.3390/app10175841.
- [29] H. Kour and M.K. Gupta, Predicting the language of depression from multivariate Twitter data using a feature-rich hybrid deep learning model, *Concurrency and Computation: Practice and Experience* **34**(24) (2022), e7224, \_eprint. <https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpe.7224>. doi:10.1002/cpe.7224.
- [30] T. Kuhn, A. Meroño-Peñuela, A. Malic, J.H. Poelen, A.H. Hurlbert, E. Centeno Ortiz, L.I. Furlong, N. Queral-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.
- [31] 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>.
- [32] 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>.
- [33] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P.N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer and C. Bizer, DBpedia – A large-scale, multilingual knowledge base extracted from Wikipedia, *Semantic Web* **6**(2) (2015), 167–195. doi:10.3233/SW-140134.
- [34] B. Li, H. Zhou, J. He, M. Wang, Y. Yang and L. Li, On the sentence embeddings from pre-trained language models (2020).
- [35] X. Li, Q. Sun and J. Shi, Research on automatic categorization of product reviews based on short text extension and BERT model, in: *WHICEB 2022 Proceedings*, 2022. <https://aisel.aisnet.org/whiceb2022/30>. doi:10.1145/3532577.
- [36] 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.
- [37] 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>.
- [38] 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. <https://www.mdpi.com/2076-3417/11/4/1838>. doi:10.3390/app11041838.
- [39] 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. doi:10.1080/01612840.2020.1814914.
- [40] C. Malighetti, A. Chirico, S. Sciara and G. Riva, #Eating disorders and Instagram: What emotions do you express?, *Annual Review of Cybertherapy and Telemedicine* (2020), 8.
- [41] 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. doi:10.1007/978-3-030-00668-6\_23.
- [42] 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.
- [43] C.P. McLean, R. Utpala and G. Sharp, The impacts of COVID-19 on eating disorders and disordered eating: A mixed studies systematic review and implications, *Frontiers in Psychology* (2022). doi:10.3389/fpsyg.2022.926709.
- [44] T. Mikolov, K. Chen, G. Corrado and J. Dean, Efficient estimation of word representations in vector space (2013).



- [45] 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. doi:10.18653/v1/2020.emnlp-demos.2.
- [46] 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. <https://mental.jmir.org/2021/1/e26011>. doi:10.2196/26011.
- [47] U.S.D. of Health and H. Services, NIMH, Let's talk about eating disorders. <https://www.nimh.nih.gov/health/publications/lets-talk-about-eating-disorders>. Accessed: 2021-10-17.
- [48] A. Olteanu, C. Castillo, F. Diaz and E. Kiciman, Social data: Biases, methodological pitfalls, and ethical boundaries, *Frontiers in Big Data* **2** (2019), 13. doi:10.3389/fdata.2019.00013.
- [49] A. Origlia and S. Di Martino, On the impact of location-related terms in neural embeddings for content similarity measures in cultural heritage recommender systems, in: *Web and Wireless Geographical Information Systems*, F. Karimipour and S. Storandt, eds, Lecture Notes in Computer Science, Springer International Publishing, Cham, 2022, pp. 108–120. ISBN: 978-3-031-06245-2. doi:10.1007/978-3-031-06245-2\_10.
- [50] 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.
- [51] 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. <https://www.sciencedirect.com/science/article/pii/S1568494621003161>. doi:10.1016/j.asoc.2021.107393.
- [52] 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.
- [53] 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.
- [54] C.S. Perone, R. Silveira and T.S. Paula, Evaluation of sentence embeddings in downstream and linguistic probing tasks (2018).
- [55] V. Pinheiro, *Computational Processing of the Portuguese Language: 15th International Conference, PROPOR 2022*, Fortaleza, Brazil, March 21–23, 2022, Proceedings, Springer Nature, 2022, Google-Books-ID: Df9kEAAQBAJ. ISBN: 978-3-030-98305-5. doi:10.1007/978-3-030-98305-5.
- [56] P. Ristoski and H. Paulheim, RDF2Vec: RDF graph embeddings for data mining, in: *SEMWEB*, 2016.
- [57] 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. ISSN: 1574-8936/2212-392X. <http://www.eurekaselect.com/article/44965>. doi:10.2174/157489312802460721.
- [58] 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.
- [59] 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).
- [60] D. Santomauro et al., Global prevalence and burden of depressive and anxiety disorders in 204 countries and territories in 2020 due to the COVID-19 pandemic, *The Lancet* **398** (2021), 700–1712.
- [61] S.D.M.G.- TIB, SDM-RDF2vec, GitHub, 2022.
- [62] M. Shah, M. Sachdeva and H. Johnston, Eating disorders in the age of COVID-19, *Psychiatry Research* **290** (2020), 113122. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7259905/>. doi:10.1016/j.psychres.2020.113122.
- [63] R. Skaik and D. Inkpen, Using social media for mental health surveillance: A review, *ACM Computing Surveys* **53**(6) (2020), 129:1–129:31. doi:10.1145/3422824.
- [64] G. Song and D. Huang, A sentiment-aware contextual model for real-time disaster prediction using Twitter data, *Future Internet* **13**(7) (2021). <https://www.mdpi.com/1999-5903/13/7/163>. doi:10.3390/fi13070163.
- [65] S. Sukunesan, M. Huynh and G. Sharp, Examining the pro-eating disorders community on Twitter via the hashtag #proana: Statistical modeling approach, *JMIR Mental Health* **8**(7) (2021), e24340, Company: JMIR Mental Health Distributor: JMIR Mental Health Institution: JMIR Mental Health Label: JMIR Mental Health Publisher: JMIR Publications Inc., Toronto, Canada. <https://mental.jmir.org/2021/7/e24340>. doi:10.2196/24340.
- [66] S. Talebi, K. Manoj and G. Hemantha Kumar, Building knowledge graph based on user tweets, in: *Data Analytics and Learning*, P. Nagabhushan, D.S. Guru, B.H. Shekar and Y.H.S. Kumar, eds, Lecture Notes in Networks and Systems, Springer, Singapore, 2019, pp. 433–443. ISBN: 9789811325144. doi:10.1007/978-981-13-2514-4\_36.
- [67] 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* **14**(2) (2022), 233–264, Publisher: IOS Press. <https://content.iospress.com/articles/semantic-web/sw210444>. doi:10.3233/SW-210444.
- [68] S. Vallurupalli, J. Sleeman and T. Finin, Fine and ultra-fine entity type embeddings for question answering, *CEUR Workshop Proceedings* **2774** (2020), 57–70.

- [69] G. Vandewiele, B. Steenwinckel, T. Agozzino, M. Weyns, P. Bonte, F. Ongenaes and F.D. Turck, pyRDF2Vec: Python implementation and extension of RDF2Vec, IDLab, 2020. <https://github.com/IBCNServices/pyRDF2Vec>.
- [70] I. Viguria, M.A. Alvarez-Mon, M. Llaveró-Valero, A. Asunsolo del Barco, F. Ortuño and M. Alvarez-Mon, Eating disorder awareness campaigns: Thematic and quantitative analysis using Twitter, *J Med Internet Res* **22**(7) (2020), e17626. <http://www.jmir.org/2020/7/e17626/>. doi:10.2196/17626.
- [71] D. Vrandečić and M. Krötzsch, Wikidata: A free collaborative knowledgebase, *Commun. ACM* **57**(10) (2014), 78–85. doi:10.1145/2629489.
- [72] 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. doi:10.7554/eLife.52614.
- [73] 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). <https://www.mdpi.com/2072-6694/11/11/1673>. doi:10.3390/cancers11111673.
- [74] 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.
- [75] A. Wongkoblap, M.A. Vadillo and V. Curcin, 6 – Social media big data analysis for mental health research, in: *Mental Health in a Digital World*, D.J. Stein, N.A. Fineberg and S.R. Chamberlain, eds, Global Mental Health in Practice, Academic Press, 2022, pp. 109–143. ISBN: 978-0-12-822201-0. doi:10.1016/B978-0-12-822201-0.00018-6.
- [76] S. Xu, C. Xiong and P. Shi, Research on argument text clustering method based on autoencoder, in: *2021 6th International Symposium on Computer and Information Processing Technology (ISCIPT)*, 2021, pp. 783–788. doi:10.1109/ISCIPT53667.2021.00164.
- [77] 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. <http://www.jmir.org/2020/11/e20550/>. doi:10.2196/20550.
- [78] Z. Ye, G. Jiang, Y. Liu, Z. Li and J. Yuan, Document and word representations generated by graph convolutional network and BERT for short text classification, in: *ECAI 2020*, IOS Press, 2020, pp. 2275–2281. doi:10.3233/FAIA200355.
- [79] 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.
- [80] 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. <https://medinform.jmir.org/2020/10/e18273>. doi:10.2196/18273.
- [81] S. Zipfel et al., The hidden burden of eating disorders during the COVID-19 pandemic, *The Lancet Psychiatry* **9** (2022), 9–11. doi:10.1016/S2215-0366(21)00435-1.