

# Semi-supervised knowledge extraction for detection of drugs and their effects\*

Fabio Del Vigna<sup>1,2</sup>, Marinella Petrocchi<sup>2</sup>, Alessandro Tommasi<sup>3</sup>,  
Cesare Zavattari<sup>3</sup>, and Maurizio Tesconi<sup>2</sup>

<sup>1</sup> Dept. of Information Engineering, University of Pisa, Italy

<sup>2</sup> Institute of Informatics and Telematics (IIT-CNR), Pisa, Italy

<sup>3</sup> LUCENSE SCaRL, Lucca, Italy

{f.delvigna,m.petrocchi,m.tesconi}@iit.cnr.it  
{alessandro.tommasi,cesare.zavattari}@lucense.it

**Abstract.** New Psychoactive Substances (NPS) are drugs that lay in a grey area of legislation, since they are not internationally and officially banned, possibly leading to their not prosecutable trade. The exacerbation of the phenomenon is that NPS can be easily sold and bought online. Here, we consider large corpora of textual posts, published on online forums specialized on drug discussions, plus a small set of known substances and associated effects, which we call seeds. We propose a semi-supervised approach to knowledge extraction, applied to the detection of drugs (comprising NPS) and effects from the corpora under investigation. Based on the very small set of initial seeds, the work highlights how a contrastive approach and context deduction are effective in detecting substances and effects from the corpora. Our promising results, which feature a F1 score close to 0.9, pave the way for shortening the detection time of new psychoactive substances, once these are discussed and advertised on the Internet.

**Keywords:** Text mining, NPS detection, NPS data mining, drugs forums, social media analysis, machine learning, automatic classification.

## 1 Introduction

US and European countries are facing a raising emergency: the trade of substances that lay in a grey area of legislation, known as New Psychoactive Substances (NPS). The risks connected to this phenomenon are high: every year, hundreds of consumers get overdoses of these chemical substances and hospitals have difficulties to provide effective countermeasures, given the unknown nature of NPS. Government and health departments are struggling to monitor the market to tackle NPS diffusion, forbid NPS trade and sensitise people to the harmful effects of these drugs<sup>4</sup>. Unfortunately, legislation is typically some

---

\* We acknowledge the EU project CASSANDRA, Computer Assisted Solutions for Studying the Availability aNd Distribution of novel psychoActive substances

<sup>4</sup> <http://www.emcdda.europa.eu/start/2016/drug-markets\#pane2/4>; All URLs in the paper have been accessed on July 10, 2016.

steps back and newer NPS quickly replace old generation of substances. Also, the abuse of certain prescription drugs, like opioids, central nervous system depressants, and stimulants, is a widespread as an alarming trend, which can lead to a variety of adverse health effects, including addiction<sup>5</sup>.

The described phenomena are being exacerbating by the fact that online shops and marketplaces convey NPS through the Internet [21]. Moreover, specialised forums offer a fertile stage for questionable organisations to promote NPS, as a replacement of well known drugs. Forums are contact points for people willing to experiment with new substances or looking for alternatives to some chemicals.

In this work, we consider the myriads of posts published on two big drugs forums, namely Bluelight<sup>6</sup> and Drugsforum<sup>7</sup>. Posts consist of natural language, unstructured text, which, generally speaking, can be analysed with text mining techniques to discover meaningful information, useful for some particular purposes [25]. We propose DAGON (DAta Generated jargON), a novel, semi-supervised knowledge extraction methodology, and we apply it to the posts of the drugs forums, with the main goals of: i) detecting substances and their effects; ii) put the basis for linking each substance to its effects. A successful application of our technique is paramount: first, we envisage the possibility to shorten the detection time of NPS; then, it will be possible to group together different names that refer to the same substance, as well as to distinguish between different substances, commonly referred to with the same name (such as “Spice” [20]) and timely detect changes in drug composition over time [8]. Finally, knowing the effects tied to novel substances, first-aid facilities may overcome the current difficulties to provide effective countermeasures.

While traditional supervised techniques usually require large amount of hand-labeled data, our proposal features a semi-supervised learning approach in order to minimize the work required to build an effective detection system. Semi-supervised learning exploits unlabeled data to mitigate the effect of insufficient labeled data on the classifier accuracy. This specific approach attempts to automatically generate high-quality training data from an unlabeled corpus. With very few information, our solution is able to achieve excellent detection results on drugs and their effects, with an FMeasure close to 0.9.

The paper is structured as follows. The next section describes our data sources. In Section 3, we introduce our semi-supervised methodology. Section 4 presents a set of experiments and results. Section 5 provides related work on mining drugs over the Internet and it discusses text analysis approaches, highlighting differences and similarities with our proposal. Finally, Section 6 concludes the paper.

---

<sup>5</sup> <https://www.drugabuse.gov/publications/research-reports/prescription-drugs/director>

<sup>6</sup> <http://www.bluelight.org>

<sup>7</sup> <https://drugs-forum.com>

## 2 Datasets

The approach in this work is tested over two different large data sources, in order to consider a variety of contents and information, and to push the automatic detection of drugs. We collected more than a decade of posts from Bluelight and Drugsforum. As shown in Table 1, the available data comprises more than half million users and more than 4.6 million posts. Data was collected through web scraping and stored in a relational database for further querying. These forums were early and partially analysed in [23] and then explored in detail [9]. Here, we present the very same datasets to show how it is possible to extract knowledge from text using few seeds as the starting point for the algorithm introduced in Section 3.

Forum	First post	Last post	Tot posts	Users
Bluelight	22-10-1999	09-02-2016	3,535,378	347,457
Drugsforum	14-01-2003	26-12-2015	1,174,759	220,071

Table 1. Drug forums: Posts and Users

### 2.1 Seeds

We have downloaded a list of 416 drug names of popular psychoactive substances, including the slang which is adopted among consumers to commonly name them, from the website of the project *Talk to Frank*<sup>8</sup> and a dataset containing 8206 pharmaceutical drugs retrieved from Drugbank<sup>9</sup>. This list constitutes a ground truth for known drugs.

Also, we collected a list of 129 symptoms that are typically associated to substance assumption.

## 3 The DAGON methodology (DAta Generated jargON)

In this section, we introduce DAGON, a methodology that will be applied in Section 4 for the task of identifying new “street names” for drugs and their effects. A street name is the name a substance is usually referred to amongst users and pushers.

The task of name identification can be split into two subtasks:

- (a) Identifying text chunks in the forums, which represent candidate drug names (and candidate drug effects);
- (b) Classifying those chunks as drugs, effects, or none of the above.

The first subtask - identification of candidates - could be tackled with different approaches, including a noun-phrase identifier<sup>10</sup>, usually based on a simple part-of-speech-based grammar, or on a technique akin to the identification of named entities, as in [14].

<sup>8</sup> <http://www.talktofrank.com>

<sup>9</sup> <http://www.drugbank.ca>

<sup>10</sup> A noun-phrase is a phrase that plays the role of a noun such as “the kid that Santa Claus forgot”.

In this work, the identification of candidates is based on domain terminology extraction techniques based on a contrastive approach similar to [16]. Essentially, we identify chunks of texts that appear to be especially significant in the context of drug forums. Based on the frequency in which terms appear both in the posts of drugs forums and in contrastive datasets dealing with different topics, we extract the most relevant terms for the forums. We have extracted unigrams, 2-grams, and 3-grams. This approach does not require English specific annotated resources and, thus, it can scale easily to different languages.

The second subtask is a classification problem. Following a supervised approach would have required to have annotated posts and use them as the training set for our classifier. Instead, we have chosen to work on unlabeled data (i.e., the posts on the drugs forums, see Section 2) and to exploit the external list of seeds introduced in Section 2.1.

We represent a candidate by means of the words found along with it when it was used in a post, selecting windows of  $N$  characters surrounding the candidate whenever it was used in the dataset. Hereafter, we call *context* (of a candidate) the text surrounding the term of interest.

Thus, we have shifted the problem: from classifying candidate street names to the classification of their contexts, which are automatically extracted from the unlabeled forum datasets.

It is worth noting that, in the drugs scenario, there would be at least 3 classes, i.e., Substance, Effect, and “none of the above” - the latter to account for the cases where the candidate does not represent substances and effects. However, the seed list at our disposal consists of flat lists of substances/effects names, provided with no additional information (Section 2.1). Therefore, in the following, we will first automatically identify positive examples for the two classes (Substance and Effect), training a classifier on them, and then we will tune the classifier settings to determine when a candidate does not fall in either.

Summarising, we have split the task of classifying a candidate into the following sub-tasks:

- (a) Fetch a set of occurrences of the term along with the surrounding text (forming in such a way the so called contexts).
- (b) Classify each context along the 2 known classes (Section 3.3).
- (c) Determine a classification for the term given the classification result for the context related to that term (obtained at step (b)).

The single context classification task [1] falls within the realm of standard text categorization, for which there is a rich literature.

Hereafter, we detail the training phase for our classifier (3.1), we give detail on the choice of seeds (3.2), we specify the procedure for classifying a new candidate (3.3), and we illustrate a simple approach to link substances to their effects (3.4).

### 3.1 Training phase

We are equipped with a list of examples for both the drugs and the effects, as described in Section 2.1. This list of entry terms is the training set for the classification task and we call it *list of seeds*.

Each post in the target drug forums was indexed by a full-text indexer (Apache Lucene<sup>11</sup>) as a single document.

The training phase is as follows:

- (i) Let  $T_S$  and  $T_E$  be the set of example contexts, for the Substance and Effects classes respectively, initialized empty.
- (ii) From the lists of seeds, we pick a new seed (a drug name) for the Substance class and one (an effect name) for the Effects class. A seed is therefore an example of the corresponding class taken from the seed list (Section 2.1). See Section 3.2 for the heuristic to select a seed out of the list.
- (iii) We use the full-text index to retrieve  $M$  posts containing the seed  $s$ ; we only use the bit of text surrounding the seed. In Section 4, we will show how results change by varying  $M$ . We pick a window of 50 characters surrounding the searched seed.
- (iv) We strip  $s$  from the text, replacing it always with the same unlikely string (such as “CTHULHUFHTAGN”), in order to avoid the bias carried by the term itself. We call the texts thus obtained  $ctx_s$  (context of seed  $s$ ).
- (v) We add the texts thus generated to the set of training examples for the category  $C$  the seed belongs to (either  $T_S$  or  $T_E$ ).
- (vi) We use the training examples to train a multiclass classification model  $M^{ctx}$ , which can be any multiclass model, as long as it features a measure (e.g., a probability) interpretable as a confidence score of the classification. In section 4 we will show results when using SVM with linear kernel [5].

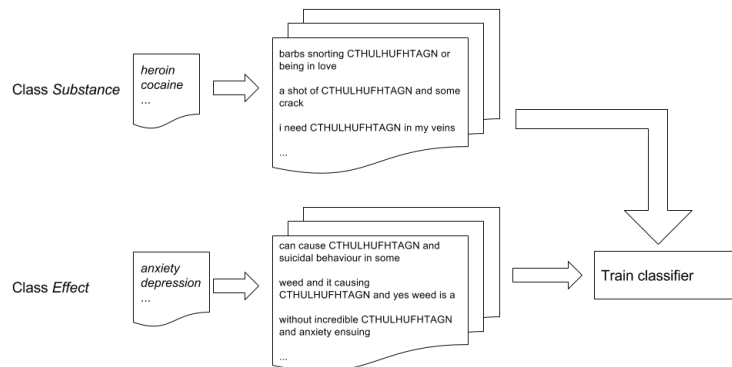
At the end of these steps, we have obtained a classifier of contexts ( $M^{ctx}$ ), but as seeds (not contexts) are labeled, we are unable to assess its performance directly. We therefore define a classifier of candidate terms ( $M^{trm}$ ) using the method described later in Section 3.3, the performance of which we can assess against the seed list. This allows us to optionally iterate back to step (ii), in order to provide additional seeds to extend the training sets, and improve performances.

The rationale behind this process is that drug (and effects) mentions will likely share at least part of their immediate contexts. Clearly, when a very small number of seeds is provided (e.g., 1 per class) there will be a strong bias in the examples ultimately used for training, which means that the resulting model will be overly specific to the type of drug used in the training. By providing more seeds, and with enough variety, the model will eventually become more generic to encompass the various drug types, and the relative differences in the contexts in which they are mentioned in the dataset.

### 3.2 Choosing a seed

Choosing an effective seed is paramount, and, in doing so, there are various aspects to consider:

<sup>11</sup> <http://lucene.apache.org/>



**Fig. 1.** Training phase

- (a) Is the seed mentioned verbatim enough times in the data collection? Failing this, the seed will only serve to collect a small number of additional training elements, and it will not impact the model enough;
- (b) Is the seed adding new information? The most effective seeds are those whose contexts are misclassified by the current iteration of the classification model. In order to pick the most useful one, we could select, from the list of available unused seeds, those whose contexts are frequently misclassified. Using these seeds, the model is modified to address a larger number of potential errors.

In information retrieval, Inverse Document Frequency [19] (*idf*) is often used along with term frequency (*tf*) as a measure of relevance of a term, capturing the fact that a term is frequent, but not so frequent to be essentially meaningless (non-meaning words, such as articles and conjunctions, are normally the most frequent ones). A common way to address point (a) would therefore be using a standard *tf-idf* metric. However, because our seeds list is guaranteed to only contain meaningful entries, we can safely select the terms occurring in more documents first (i.e., with an increasing *idf*). We leave point (b) for future work.

### 3.3 Classification of a new candidate

At the end of the training phase, the classifier  $M^{ctx}$  has been trained - on contexts of the selected seeds - to classify as either pertaining to substances or effects. Here, we describe the procedure by which, given a new candidate  $c$ , we establish what class (Substance or Effect) it belongs to. The new candidates are chosen from the terms which are more relevant for the forums. Such terms are extracted according to the contrastive approach described in Section 3, subtask (a).

The training phase produces a model  $M^{ctx}$  by which contexts in which the term appears are classified - we define here a model  $M^{trm}$  by which the term itself is classified into either Substance, Effect, or “none-of-the-above”.  $M^{trm}$  is defined as a function of a candidate  $c$  and the existing model  $M^{ctx}$  as such:

1. We apply steps (iii) and (iv) of the algorithm described in 3.1 to obtain the contexts for  $c$  ( $ctx_c$ ).
2. We classify the elements of  $ctx_c$  using  $M^{ctx}$ . We discard all categorizations whose confidence, according to the model, falls below a threshold  $\theta_p$ , which we have experimentally set to 0.8 as a reference value.
3. We consider the remaining categorizations thus obtained. If a sizeable portion of them ( $\theta_c$ , initially set to 0.6, we will show how results vary along with its value) belongs to the same class  $C$ , then  $c$  belongs to  $C$ ; otherwise it is left unassigned.

In Figure 2 we give a high level graphical description of this process.

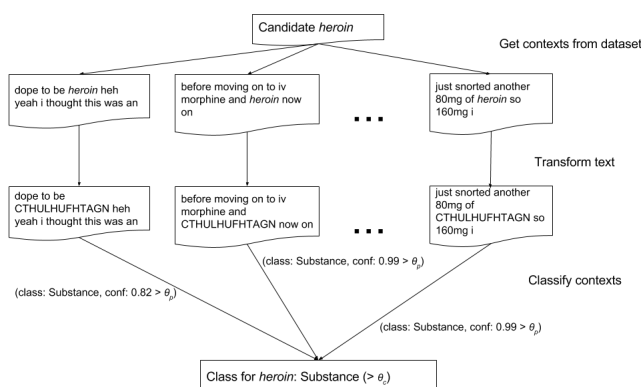


Fig. 2. Classification of a new term

### 3.4 Linking substances to effects

We outline here a simple procedure by which we can associate the substances mentioned in the drugs forums to the effects they produce.

When indexing a post, the significant terminology elements found in the post are linked to it as metadata. As introduced, the terminology elements have been extracted following a contrastive approach, as in [16].

We assume to have already tagged the terminology elements found in each post as referring to substances or effects, using the method described in Section 3.3. Thus, when searching for mentions of a particular substance, we can correspondingly fetch, for each post the substance mention is found in, the relative metadata. We can therefore retrieve a list of effects mentioned along with the substance – by sorting this list of effects by frequency, it is very likely that those effects are related to the substance.

As a simple example, let's suppose to have a single post, with Text: *heroin gave me a terrible headache*; Substances: [heroin]; Effects: [headache].

Intuitively, we can assume that [headache] is an effect of [heroin]. If we consider all the posts in our datasets where the substance [heroin] is among the metadata, and we count the most frequent metadata effects associated to [heroin], we can have an indication of the links between substances and effects. However many substances may appear in the same text. Thus, it is necessary to filter out the rarest links substance-effect since they are often due by chance. Section 4 will report on some findings we were able to achieve for our datasets about drugs and their effects.

## 4 Experiments

We show a set of experiments on the data described in Section 2. First, from all the posts, we need to identify a list of candidates (unless we want to try and classify every term – a possible, but undesirable strategy, to pinpoint substances or effects out of which. Candidates are selected using a contrastive terminology extraction [16], to identify terms and phrases common within the community and yet specific to it; this is the first subtask outlined in section 3. Then, we apply the  $M^{trm}$  classifier, described in Section 3.3, to assign to candidates either the class Substance or Effect or none of the above, and evaluate the performance of the classification. The intermediate  $M^{ctx}$  classifier was trained using SVM with linear kernel [5].

We report experiments and results for the Bluelight forum. The lists used to select seeds and to validate results have been described in Section 2.1. These lists represent 2 classes: Substance and Effect.

It is worth noting that, for our experiments, we consider the intersection between the lists of seeds and the extracted terminology. This is necessary because: i) items that are present in the lists may not be present in the downloaded dataset; ii) many terminological entries might be neither drug names nor drug effects. The intersection contains 226 substances and 89 effects. Some of these will be used as seeds, the rest of the entries to validate the results.

The results are given in terms of three standard metrics in text categorization, based on true positives (TP - items classified in category  $C$ , actually belonging to  $C$ ), false positives (FP - items classified in  $C$ , actually not belonging to  $C$ ) and false negatives (FN - items not classified in  $C$ , actually belonging to  $C$ ), computed over the decisions taken by the classifier: precision<sup>12</sup>, recall<sup>13</sup> and F1-micro averaged<sup>14</sup>.

The first results are in Table 2 and Figure 3. Even though the training set is limited to a small number of entries, the results are interesting: with only 6 seeds, the proposed methodology achieves a F1 score close to 0.88 (on the 2 classes - Substance and Effect). With the aim of monitoring the diffusion of new substances, the result is quite promising, since it is able to detect unknown substances without human supervision.

<sup>12</sup>  $precision = \frac{TP}{TP+FP}$


<sup>13</sup>  $recall = \frac{TP}{TP+FN}$

<sup>14</sup> harmonic mean of *precision* and *recall*:  $F1 = 2 \cdot \frac{precision \cdot recall}{precision+recall}$



# of seeds	Recall	Precision	F1
1	0.502	0.649	0.566
2	0.576	0.734	0.645
3	0.65	0.827	0.728
4	0.769	0.891	0.826
5	0.823	0.909	0.864
6	0.832	0.926	0.876

**Table 2.** Classification results for substances and effects, varying the number of seeds




**Fig. 3.** Recall, precision and F1 varying the number of seeds

**Dealing with “the rest”.** Finding mentions of new substances or effects means classifying candidate terms in either one class. Playing with thresholds, we can discard some candidates, as belonging to none of the two classes (see Section 3.3).

Thus, within the extracted terminology, we have manually labeled about 100 entries as neither drugs nor effects, and we have used them as candidates. This has been done to evaluate the effectiveness of using the parameter  $\theta_c$  to avoid classifying these terms as either substances or effects. Performance-wise, this resulted in few more false positives given by terms erroneously assigned to the substance and effect classes, when instead these 100 candidates should ideally all be discarded. The results are in Table 3 and Figure 4. We can observe that, when we include in the evaluation also those data that are neither substances nor effects, with no training data other than the original seeds, and operating only on the thresholds, the precision drops significantly.

To achieve comparable performances, we have conducted experiments changing the number of seeds and  $\theta_c$  used to keep relevant terms. The results are shown in Table 4 and Figure 5. The higher the threshold, the higher the precision, while increasing the number of seeds improves the recall, which is to be

# of seeds	Recall	Precision	F1
1	0.502	0.502	0.502
2	0.576	0.563	0.569
3	0.650	0.628	0.639
4	0.769	0.694	0.730
5	0.823	0.723	0.770
6	0.832	0.733	0.779


**Table 3.** Classification results for substances and effects, including the “rest” category**Fig. 4.** Recall, precision and F1 including the “rest” category

expected: adding seeds “teaches” the system more about the variety of the data. Moreover, recall augments when we increase the number of contexts per seed used to train the system (Table 5 and Figure 6).


# of seeds	Recall 0.75	Precision 0.75	F1 0.75	Recall 0.8	Precision 0.8	F1 0.8
5	0.607	0.755	0.673	0.508	0.787	0.618
10	0.759	0.852	0.803	0.654	0.889	0.754
15	0.811	0.837	0.824	0.705	0.874	0.781
20	0.833	0.854	0.843	0.753	0.866	0.805

**Table 4.** Precision, Recall and F1 with  $\theta_c$  set to 0.75 and 0.8 (incl. “rest” category)

It is worth noting that increasing the number of contexts used to classify a new term seems to have no effect after few contexts, as shown in Table 6 and Figure 7). This indirectly conveys an information on the variety of contexts present on the investigated datasets.




**Fig. 5.** Precision and Recall with  $\theta_c$  set to 0.75 and 0.8 (incl. “rest” category)



**Fig. 6.** Recall and precision varying the number of contexts (snippets) per seed, 10 seeds used

Interestingly, the automated drug detection reported 1846 drugs in Bluelight and 1857 in DrugsForum, with 1520 drugs in common between the two forums. Moreover, some drugs appear exclusively in one of the two forums, like the *triptorelin*, *candesartan* and *thiorphan* in Bluelight and the *lymecycline*, *boceprevir* and *imipenem* in Drugsforum, although the majority is shared.

Finally, upon training the system with the seeds, for every post it is possible to link the drugs to their effects. An example of links is in Table 7.



**Fig. 7.** Recall and precision varying the number of contexts (snippets) per new term, 10 seeds used

# of contexts	Recall	Precision	F1
100	0.437	0.709	0.541
1000	0.675	0.802	0.733
2000	0.742	0.830	0.784
3000	0.759	0.852	0.803
4000	0.769	0.838	0.802
5000	0.817	0.867	0.841
6000	0.831	0.851	0.840

**Table 5.** Results varying the number of contexts per seed

# of contexts	Recall	Precision	F1
10	0.763	0.758	0.760
50	0.746	0.815	0.779
100	0.759	0.852	0.803
150	0.763	0.852	0.805
200	0.753	0.854	0.800
300	0.759	0.852	0.803

**Table 6.** Results varying the number of contexts per new term

## 5 Related work

Recently, Academia has started mining online communities, to seek for comments on drugs and drugs reactions [27]. Indeed, websites, forums and social networks are a source of spontaneous information from users, with abundance of data about experiences, doses, assumption methods [7, 9]. Authors in [15] realized ADRMine, a tool for adverse drugs reaction detection. The tool relies on advanced machine learning algorithms and semantic features based on word clusters - generated from pre-trained word representation vectors using deep learning techniques. Also, intelligence analysis has been applied to social media to detect new outbreaking trends in drug markets, as in [24]. One of the raising phenomenon connected to the consumption of psychoactive substances is the increasing adoption of nonmedical use of prescription drugs [13], such as sedatives, opioids, and stimulants, which could be very harmful for the individual. Even

Drug	Effects
heroin	anxiety, euphoria
cocaine	euphoria, anxiety, comedown, paranoia, psychosis
ketamine	euphoria, anxiety, visuals, comedown, hallucinations, nausea
methadone	anxiety, euphoria
codeine	euphoria, anxiety, nausea
morphine	euphoria, anxiety, analgesic, nausea
amphetamine	euphoria, anxiety, comedown, psychosis, visuals
oxycodone	euphoria, anxiety
methamphetamine	euphoria, anxiety, psychosis, comedown, paranoia
dopamine	euphoria, anxiety, comedown, psychosis

**Table 7.** Main effects of the most discussed drugs on Bluelight

these drugs are often traded online by fake pharmacies [12, 11] and advertised through social networks.

Automated text analysis enjoys a rich literature, due to the relevance of the task being ever more poignant in the communication age, and the collaborative web especially. The amount of data available nowadays has made the discipline veer towards more machine learning-based approaches. Because complex tasks might require many training examples, however, there is a vivid study on unsupervised and semi-supervised approaches.

Our task encompasses identifying names in text, something often associated with named-entity extraction. Unsupervised methods such as [22] use unlabeled data contrasted with other data assumed irrelevant - to use as negative examples - in order to build a classification model. Instead, we use seeds, a small set of examples, because the writers on forums often attempt not to mention drugs explicitly, and resorting to paraphrases or nicknames, making a purely contrastive approach difficult to apply. Also multi-level bootstrapping proved to be a valid improvement in information extraction [17]; this techniques feature an iterative process to gradually enlarge and refine a dictionary of common terms. Our approach, instead, splits the problem of finding candidate terms and classifying them in two separate subproblems, the second of which is fed with a small number of annotated examples, i.e., the seeds. Co-training is a common technique [3] to evaluate whether to use an unlabeled piece of data as a training example: the idea is building different classifiers, and use the label assigned by one as a training example for another. In our case, we instead leverage the redundancy among the data, to ensure candidate examples are selected with a high degree of confidence. Relation extraction is an even more complex task which seeks for the relationships among the entities. This is of relevance to us, because substances can only be identified basing on the role that they have in the sentence, since common names are often used to refer to them. Work in [18] proposes a method based on corpus statistics that requires no human supervision and no additional corpus resources beyond the corpus used for relation extraction. Our approach does not explicitly address relation extraction, but it exploits the redundancy of a substance (or effect) being often associated with other entities to identify

them. KnowItAll [10] is a tool for unsupervised named entity extraction with improved recall, thanks to the pattern learning, the subclass extraction and the list extraction features that still includes bootstrapping to learn domain independent extraction patterns. There is a similarity with our approach due to the fact that, for us, common mention patterns are also strong indicators of the substance or effect class; however, we do not use patterns to extract, but only, implicitly, for classification purposes. Furthermore, [4] pursues the thesis that much greater accuracy can be achieved by further constraining the learning task, by coupling the semi-supervised training of many extractors for different categories and relations; we use a single multiclass classifier to achieve the same goal. Under the assumption that the number of labeled data points is extremely small and the two classes are highly unbalanced, the authors of [26] realized a stochastic semi-supervised learning approach that was used in the 2009-2010 Active Learning Challenge. While the task is similar, our approach is different, because we do not need to use unlabeled data as negative examples. The framework proposed in [6] suggests to use domain knowledge, such as dictionaries and ontologies, as a way to guide semi-supervised learning, so as to inject knowledge into the learning process. We have not relied on rare expert knowledge for our task, preferring to stay with the notion that a few labeled seeds were easier to produce than dictionaries or other forms of expert knowledge representations. A mixed case of learning extraction patterns, relation extraction and injecting expert knowledge is in [2], which also shows the challenge of evaluating a technique when few labeled examples are available. As shown above, the problem of building a model with a limited set of information, but with a large enough amount of data, has been tackled by various angles. Our main staples were: a) the availability of a large set of unlabeled data, and b) the availability of a small set of labeled substance and effect names.

## 6 Conclusions

We have shown the capability to automatically identify and classify substances and effects from posts of drugs forums, making use of a semi-supervised text mining approach. Human intervention is required for the creation of a small training set, but the algorithm is able to automatically discover substances and effects with such a very few initial information. We believe our proposal will help contrasting the diffusion of novel psychoactive substances, which spreads on the online market at an impressive high rate. Furthermore, shortening the time of drugs detection will help saving lives and sensitize drug consumers about the risks of their choices.

## References

1. Attardi, G., Gull, A., Sebastiani, F.: Theseus: Categorization by context. Universal Computer Science (1998)

2. Bellandi, A., Nasoni, S., Tommasi, A., Zavattari, C.: Ontology-driven relation extraction by pattern discovery. In: Information, Process, and Knowledge Management. pp. 1–6. IEEE Computer Society (2010)
3. Blum, A., Mitchell, T.: Combining labeled and unlabeled data with co-training. In: Computational Learning Theory. pp. 92–100. ACM (1998)
4. Carlson, A., Betteridge, J., Wang, R.C., Hruschka Jr, E.R., Mitchell, T.M.: Coupled semi-supervised learning for information extraction. In: Web Search and Data Mining. pp. 101–110. ACM (2010)
5. Chang, C.C., Lin, C.J.: LIBSVM: A library for support vector machines. ACM Trans. Intell. Syst. Technol. 2(3), 27:1–27:27 (May 2011)
6. Chang, M.W., Ratinov, L., Roth, D.: Guiding semi-supervision with constraint-driven learning. In: Annual Meeting - Association for Computational Linguistics. pp. 280–287 (2007)
7. Davey, Z., Schifano, F., Corazza, O., Deluca, P.: e-Psychonauts: Conducting research in online drug forum communities. Journal of Mental Health 21(4), 386–394 (2012)
8. Davies, S., et al.: Purchasing legal highs on the Internet - is there consistency in what you get? QJM 103(7), 489–493 (2010)
9. Del Vigna, F., Avvenuti, M., Bacciu, C., Deluca, P., Marchetti, A., Petrocchi, M., Tesconi, M.: Spotting the diffusion of new psychoactive substances over the internet. arXiv preprint arXiv:1605.03817 (2016)
10. Etzioni, O., Cafarella, M., Downey, D., Popescu, A.M., Shaked, T., Soderland, S., Weld, D.S., Yates, A.: Unsupervised named-entity extraction from the web: An experimental study. Artificial intelligence 165(1), 91–134 (2005)
11. Freifeld, C.C., Brownstein, J.S., Menone, C.M., Bao, W., Filice, R., Kass-Hout, T., Dasgupta, N.: Digital drug safety surveillance: monitoring pharmaceutical products in Twitter. Drug Safety 37(5), 343–350 (2014)
12. Katsuki, T., Mackey, T.K., Cuomo, R.: Establishing a link between prescription drug abuse and illicit online pharmacies: Analysis of Twitter data. Journal of Medical Internet Research 17(12) (2015)
13. Mackey, T.K., Liang, B.A., Strathdee, S.A.: Digital social media, youth, and non-medical use of prescription drugs: the need for reform. Journal of Medical Internet Research 15(7), e143 (2013)
14. Marsh, E., Perzanowski, D.: MUC-7 evaluation of IE technology: Overview of results. In: Seventh Message Understanding Conference (MUC-7) (1998)
15. Nikfarjam, A., Sarker, A., OConnor, K., Ginn, R., Gonzalez, G.: Pharmacovigilance from social media: mining adverse drug reaction mentions using sequence labeling with word embedding cluster features. Journal of the American Medical Informatics Association 22(3), 671–681 (2015)
16. Penas, A., Verdejo, F., Gonzalo, J.: Corpus-based terminology extraction applied to information access. In: Corpus Linguistics. pp. 458–465 (2001)
17. Riloff, E., Jones, R., et al.: Learning dictionaries for information extraction by multi-level bootstrapping. In: AAAI/IAAI. pp. 474–479 (1999)
18. Rosenfeld, B., Feldman, R.: Using corpus statistics on entities to improve semi-supervised relation extraction from the web. In: Annual Meeting - Association for Computational Linguistics. pp. 600–607 (2007)
19. Salton, G., Buckley, C.: Term-weighting approaches in automatic text retrieval. Inf. Process. Management 24(5), 513–523 (1988)
20. Schifano, F., Corazza, O., Deluca, P., Davey, Z., Furia, L.D., Farre', M., Flesland, L., Mannonen, M., Pagani, S., Peltoniemi, T., Pezzolesi, C., Scherbaum, N., Siemann, H., Skutle, A., Torrens, M., Kreeft, P.V.D.: Psychoactive drug or mystical

- incense? Overview of the online available information on Spice products. *International Journal of Culture and Mental Health* 2(2), 137–144 (2009)
21. Schmidt, M.M., Sharma, A., Schifano, F., Feinmann, C.: Legal highs on the net: Evaluation of UK-based websites, products and product information. *Forensic Science International* 206(1), 92–97 (2011)
  22. Smith, N.A., Eisner, J.: Contrastive estimation: Training log-linear models on unlabeled data. In: *Annual Meeting - Association for Computational Linguistics*. pp. 354–362 (2005)
  23. Soussan, C., Kjellgren, A.: Harm reduction and knowledge exchange—a qualitative analysis of drug-related Internet discussion forums. *Harm Reduction Journal* 11(1), 1–9 (2014)
  24. Watters, P.A., Phair, N.: Detecting illicit drugs on social media using automated social media intelligence analysis (ASMIA). In: *Cyberspace Safety and Security*, pp. 66–76. Springer (2012)
  25. Witten, H.I., Don, J.K., Dewsnip, M., Tablan, V.: Text mining in a digital library. *International Journal on Digital Libraries* 4(1), 56–59 (2004)
  26. Xie, J., Xiong, T.: Stochastic semi-supervised learning on partially labeled imbalanced data. *Active Learning Challenge Challenges in Machine Learning* (2011)
  27. Yang, C.C., Yang, H., Jiang, L.: Postmarketing drug safety surveillance using publicly available health-consumer-contributed content in social media. *ACM Trans. Manage. Inf. Syst.* 5(1), 2:1–2:21 (Apr 2014)