

# GrASP: A Library for Extracting and Exploring Human-Interpretable Textual Patterns

Piyawat Lertvittayakumjorn<sup>1</sup>, Leshem Choshen<sup>2</sup>, Eyal Shnarch<sup>2</sup>, Francesca Toni<sup>1</sup>

<sup>1</sup>Department of Computing, Imperial College London, United Kingdom

<sup>2</sup>IBM Research

pl1515@imperial.ac.uk, {leshem.choshen, eyals}@il.ibm.com, ft@imperial.ac.uk

## Abstract

Data exploration is an important step of every data science and machine learning project, including those involving textual data. We provide a novel language tool, in the form of a publicly available Python library for extracting patterns from textual data. The library integrates a first public implementation of the existing GrASP algorithm. It allows users to extract patterns using a number of general-purpose *built-in* linguistic attributes (such as hypernyms, part-of-speech tags, and syntactic dependency tags), as envisaged for the original algorithm, as well as domain-specific *custom* attributes which can be incorporated into the library by implementing two functions. The library is equipped with a web-based interface empowering human users to conveniently explore data via the extracted patterns, using complementary *pattern-centric* and *example-centric* views: the former includes a reading in natural language and statistics of each extracted pattern; the latter shows applications of each extracted pattern to training examples. We demonstrate the usefulness of the library in classification (spam detection and argument mining), model analysis (machine translation), and artifact discovery in datasets (SNLI and 20Newsgroups).

**Keywords:** pattern extraction, textual data exploration, resource tool, grasp, python library, model debugging, model analysis, data analysis

## 1. Introduction

Pattern-based analysis is a promising approach for exploring textual data in many tasks, such as authorship classification (Houvardas and Stamatatos, 2006), relation extraction (Hearst, 1992; Peng et al., 2014) and argument mining (Madnani et al., 2012). Patterns uncover prominent characteristics of the data which lead to interesting insights, helping humans proceed to the next steps (e.g., data pre-processing, setting model hyperparameters). In this paper, we implement and publicly release a Python library named GrASP after a supervised algorithm which learns expressive patterns (in the form of sequences of conjunctions of attributes) characterizing linguistic phenomena (Shnarch et al., 2017). Apart from the usefulness of the web resource, this is the first public implementation of GrASP. To illustrate the richness of GrASP patterns, examples 1–3 in Table 1 are all SMS spam messages from the dataset by Almeida et al. (2011). While there is little word overlap between them, their commonality is apparent, even if hard to name. The GrASP algorithm can reveal an underlying structure which generalizes these three realizations of spams within a single pattern: a positive-sentiment word (represented by `SENTIMENT:pos`), closely followed by a determiner (represented by `POS:det`, where `POS` stands for ‘part of speech’), and then by a proper noun (represented by `POS:propn`).

The input for the algorithm amounts to two sets of texts: in one (the positive set) the target phenomenon appears in all examples (e.g., all are spam messages), and in the other (the negative set) it does not. GrASP looks for commonalities prominent within the texts of one set but not shared across the sets. To be able to rec-

ognize common aspects of texts, beyond their surface form realizations, all input tokens are augmented with a variety of linguistic attributes such as part-of-speech tags, named entity information, or pertinence to a lexicon (e.g., of sentiment words). Attributes are selected to maximize a score (by default their information gain about the label). Then, they are combined by a greedy algorithm to generate patterns which are most indicative either to the positive or the negative set.

Since the patterns are a combination of readable attributes, they are human interpretable. So, they can be used to provide insights about the data and contribute to explainability. For instance, example 1 in Table 1 is identified as a spam message because it contains the phrase *awarded a SiPix*. Patterns can also be used for classification (e.g., if the pattern mentioned above is matched in a message, we can classify it as spam).

The contribution of this paper is threefold:

- We release GrASP, a Python library for extracting interpretable patterns from text. Although the algorithm of GrASP was initially proposed in 2017, we are the first to make its implementation publicly available.<sup>1</sup> The library provides additional extensions, allowing users to customize the algorithm to their specific needs (see §2).
- We provide a novel web-based GUI which displays GrASP patterns and their matches in a dataset. This exploration tool allows users to conveniently explore the patterns and the data with both *pattern-centric* views and *example-centric* views and gain

<sup>1</sup>The library is available at <https://github.com/plkumjorn/GrASP>

Use case	No.	Sentences Matched	Pattern
SMS Spam	1	You are <b>awarded a SiPix</b> Digital Camera...	[[SENTIMENT:pos], [POS:det], [POS:propn]]
	2	...to <b>WIN a FREE Bluetooth</b> Headset...	(A positive-sentiment word, closely followed by a determiner, and then by a proper noun)
	3	...for <b>Free ! Call The Mobile</b> Update...	
Topic-Dependent Argument Mining	4	<b>Evidence suggests that</b> TOPIC is...	[[ARGUMENTATIVE], [POS:verb], [LEMMA:that, POS:adp]]
	5	<b>Poll in 2002 found that</b> 58% of...	
	6	... <b>data indicate that</b> TOPIC reduces...	(An argumentative word, closely followed by a verb, and then by the word "that" which is also a preposition)
Machine Translation	7	<b>choose Windows &gt;</b> Channels	[[HYPERNYM:choose], [DEP:dobj], [LEMMA:>]]
model analysis	8	<b>select Control &gt;</b> Test Movie	(A type of <i>choose</i> (verb), closely followed by a word with the dependency type <i>dobj</i> , and then by >)
	9	<b>select Modify &gt; Transform &gt;</b> Skew	

Table 1: Examples of GrASP patterns capturing the common structure in a variety of surface forms – the sentences matched. Matched words are in bold. A description of each pattern is provided below it, in parentheses.

```

1 from grasp import GrASP
2 # Step 1: Create the GrASP model
3 grasp_model = GrASP(num_patterns = 200,
4                     gaps_allowed = 2, alphabet_size = 200,
5                     include_standard = ['TEXT', 'POS', 'NER',
6                                       'SENTIMENT'])
7 # Step 2: Fit it to the training data
8 grasp_model.fit_transform(pos_exs, neg_exs)
9 # Step 3: Export the results
10 grasp_model.to_csv('results.csv')
11 grasp_model.to_json('results.json')

```

Listing 1: Basic usage of GrASP

insights (see §3).

- We exemplify the usefulness of the provided library in three use cases: (i) classification (see §4), (ii) model analysis (see §5), and (iii) dataset artifacts unveiling (see §6). These use cases cover several applications including spam classification, argument mining, topic classification, natural language inference, and machine translation.

## 2. GrASP Library for Pattern Extraction

Listing 1 shows the basic usage of GrASP for extracting patterns from training data, i.e., two lists of texts containing positive examples (`pos_exs`) and negative examples (`neg_exs`). In the first step, the user specifies hyperparameters such as the desired number of patterns and the set of linguistic attributes for augmenting input tokens. The following general-purpose attributes are built-in: the token text itself (in lower case), its lemma, wordnet hypernyms, part-of-speech tags, named-entity, syntactic dependency, and sentiment tags.<sup>2</sup> After that, the user trains the GrASP model and exports the results to a csv file or a json file. Both file types summarize all the extracted patterns and their relevant

<sup>2</sup>We use spacy for tokenization and tagging, nltk and lesk (Lesk, 1986) for word sense disambiguation and finding hypernyms, and a sentiment lexicon by Hu and Liu (2004). For hypernyms, we consider only three levels above the token of interest in order to avoid terms that are too abstract to comprehend (e.g., *psychological feature*, *group action*, *entity*).

statistics, while the json file also contains the configuration, the alphabet, and the training examples annotated with the patterns matched.

Note that in our version of GrASP, we add several parameters which are not found in the original algorithm. While the original GrASP uses information gain as the criterion to rank and select patterns, Shnarch et al. (2020) preferred  $F_\beta$  (where  $\beta \leq 1$ ) as a criterion, emphasizing precision over recall. Hence, we allow the users to implement their own criteria, tailored to their use cases. Other new parameters for customizing GrASP include the number of gaps allowed between the matched tokens (which overrides the window size parameter) and the minimum coverage for a pattern to be selected.

### 2.1. Extending GrASP Capabilities

Besides the basic usage, we provide the following new useful features.

**Translating patterns into natural language explanations.** It could be difficult for lay users to read and understand linguistic attributes in patterns. Therefore, using templates, we provide a function `pattern2text` for translating a pattern (e.g., `[[HYPERNYM:communication.n.02], [POS:NUM]]` with a window size of 4) into an English explanation (e.g., “A type of communication (n), closely followed by a number”). This feature removes the barrier to an audience without linguistic knowledge to comprehend the resulting patterns. We believe that it is very important in providing explainability for all, not just for NLP or machine learning experts.

**Using custom attributes.** In addition to the built-in attributes, users may require custom attributes, to better fit the patterns to their domain. For instance, given biomedical texts, attributes indicating whether a token refers to a kind of biomedical substance (e.g., protein, drug, enzyme) could be useful for relation extraction from texts. Our library allows users to create custom attributes by using the `CustomAttribute` class and implementing two functions. One extracts the attributes from an input text, and the other explains the attributes

for the translation feature. We provide an example of custom attributes in the README file of our library’s GitHub repository.

**Removing redundant patterns.** We provide a function to remove specific patterns subsumed by more general ones. For instance, all `[[TEXT: .], [POS: NUM]]` matches are also matches of `[[POS: NUM]]`. Hence, this function removes the former (specific) pattern. A flag further allows to condition removal on the fact that the score of the removed pattern must be lower than the score of the remaining one.

**Vectorizing texts using patterns.** Given an input text  $t$  and  $n$  GrASP patterns, we provide a function to create a ternary vector of length  $n$  indicating which of the patterns are matched in  $t$  and whether they are positive or negative. This vector can be used as features of  $t$  for downstream tasks.

### 3. The Web-Based Exploration Tool

We provide a web-based exploration tool implemented using Flask.<sup>3</sup> Taking as input the json file exported by the GrASP library, this tool provides four types of reports: two are pattern-centric and the other two are example-centric. Figures 1–4 show (parts of) these reports for the spam classification example. Additionally, all reports for all the use cases mentioned in this paper are available online for further examination.<sup>4</sup>

**I. Pattern-centric level 1.** (Figure 1) This report lists all the patterns together with the configuration of GrASP used for extracting them and the statistics about the training data. For each pattern (row), its statistics on the training data are reported, including the number of positive and negative examples matched, coverage, the metric score, precision, recall, and F1. Users can click a header to sort the table based on the column values. Moreover, users can click the header of the pattern column to translate the patterns into their natural language meanings (as shown in Figure 1). Once users click a pattern in the table, they will be redirected to that pattern’s level 2 report.

**II. Pattern-centric level 2.** (Figure 2) This report can be generated for each pattern individually. In addition to the pattern’s statistics, it shows the positive and negative examples matched by this pattern, with the corresponding tokens highlighted. Users can click the link icon at the end of each example to go to its example-centric level 2 report (see below).

**III. Example-centric level 1.** (Figure 3) This report lists all positive and negative training examples (with pagination). For each example, words matched by only positive pattern(s) and only negative pattern(s) are highlighted in green and red, respectively, while words matched by both types of patterns are in purple. By clicking a highlighted word, users can see the list of

patterns matching this word. Clicking the link icon of an example directs to its example-centric level 2 report.

**IV. Example-centric level 2.** (Figure 4) This report can be generated for each training example, showing all the positive and negative patterns matching it. Clicking a pattern links to its pattern-centric level 2 report.

Overall, these reports provide a convenient way for quickly exploring the data via the lens of the common patterns found in it. Also, we can use this exploration tool to display results from other pattern extraction algorithms, in addition to GrASP, provided that their results are organized into the specific json format required by this tool (as described in the README file of our GrASP GitHub repository). Next, we will show how insights can be gained from this exploration tool with outputs from the GrASP algorithm.

## 4. Use Case: Classification

Apart from the spam classification example discussed throughout the paper, we present another classification use case of our tools targeting the corpus of topic-dependent argument mining (Shnarch et al., 2018). Generally, argument mining aims to identify and extract the structure of inference and reasoning expressed as arguments in texts (Lawrence and Reed, 2020). The dataset used here in particular focuses on identifying, in text, evidence for claims. It includes 4,065 training and 1,720 test pairs of topics and sentences, each labeled with whether the sentence is a piece of valid evidence for the topic.

We augment each token with the built-in attributes as well as a custom binary attribute indicating whether a token was found in a lexicon of argumentative words. We use information gain to select patterns, set the number of patterns to 100, and allow up to 2 gaps in each pattern. Next, we train and examine the patterns with the pattern-centric reports.

We find that the most indicative pattern is the word *that* in its *preposition* meaning (as opposed to other meanings of that word, e.g., a determiner, as in this sentence). This is a very nice outcome since GrASP automatically revealed what experts have found – “that” is an indicator of argumentative content and can be used as an effective source of weak supervision (Levy et al., 2017). Other informative patterns are about studies: presenting them (e.g., `[[POS: NUM],[LEMMA: study]]`) which matches phrases like *one study*, *Two major studies*, *A 1999 meta study*, and *a 25 year longitudinal study*), declaring their findings (e.g., examples 4–6 in Table 1), and reporting their experimental results (e.g., `[[POS: ADP],[POS: VERB],[HYPERNYM: risk]]`) which covers expressions such as *in reducing the risk*, *more than tripled the risk*, and *as controlling the risks*). These patterns, too, have rediscovered what is known in the domain – texts describing studies are a good source for evidence (Rinott et al., 2015; Ein-Dor et al., 2020). The exploration tool enables having a human in the loop by making it easy to compare patterns (e.g., by their

<sup>3</sup>[flask.palletsprojects.com/en/1.1.x/](https://flask.palletsprojects.com/en/1.1.x/)

<sup>4</sup><https://plkumjorn.github.io/GrASP>

# A Summary of Extracted Patterns

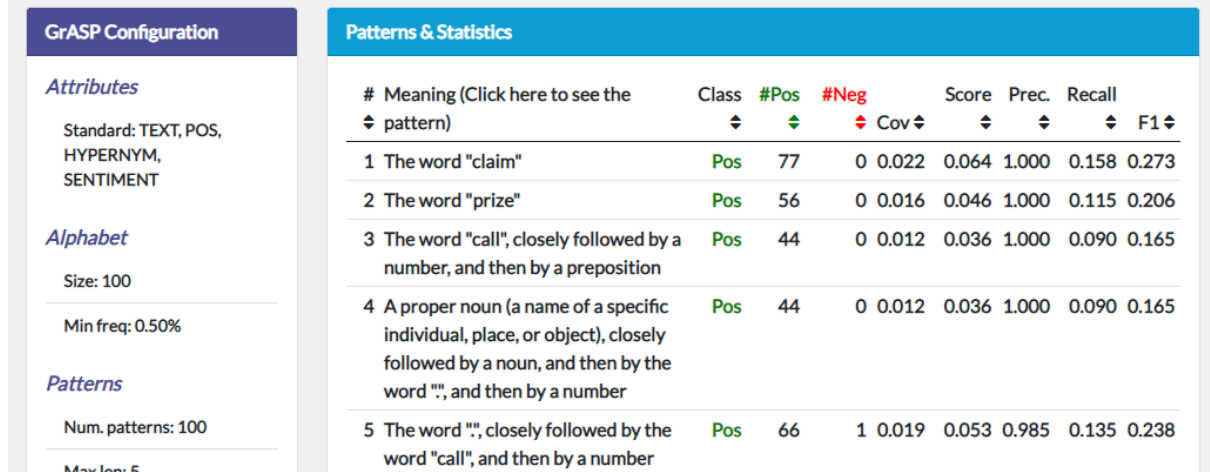


Figure 1: A part of the pattern-centric level 1 report listing the patterns and their metrics with respect to the training data, the configuration of GrASP, and the training set statistics.

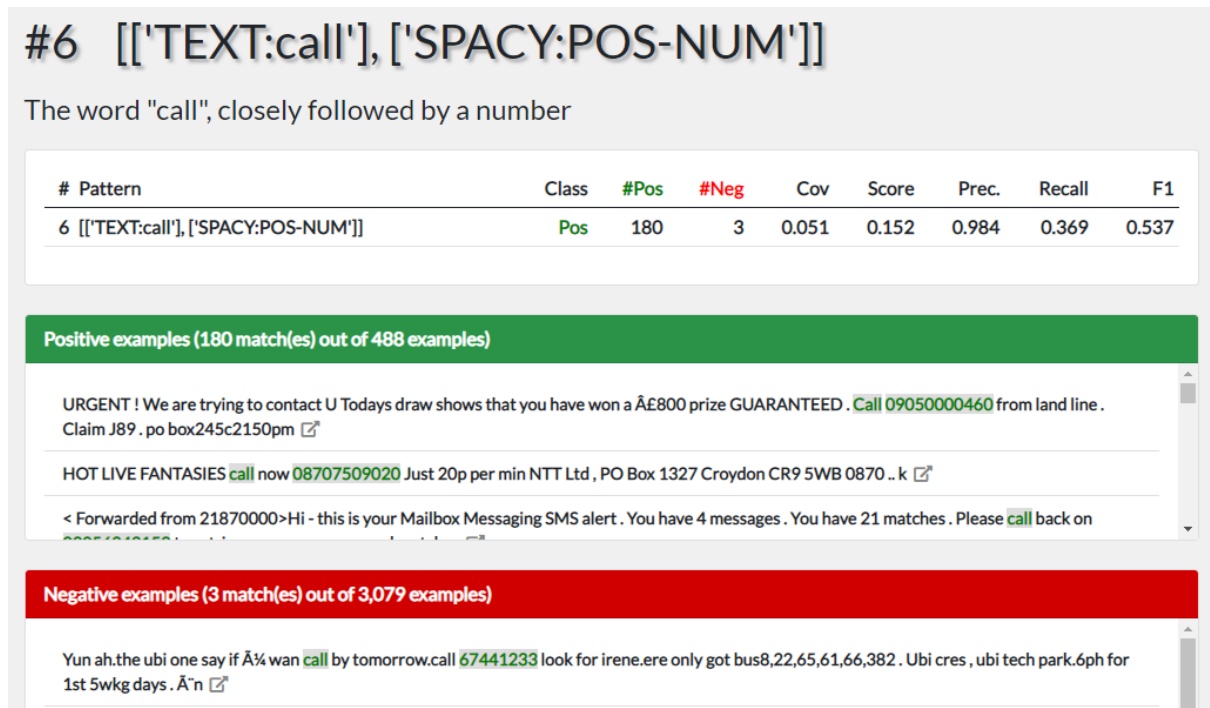


Figure 2: A part of the pattern-centric level 2 report listing positive and negative matches by a specific pattern.

recall-precision trade off), examine their realizations in the positive and negative sets, check pattern combination (e.g., click on a purple word in the example-centric report to see all patterns that matched it), and evaluate the contribution of adding custom attributes (such as the lexicon of argumentative words). After that, users can improve on their own the list of patterns by filtering, correcting, and merging patterns. The selected patterns can then contribute to the downstream pipeline (e.g., by us-

ing the most precise/high-coverage patterns to improve precision/recall). Alternatively, the patterns can be used to obtain weak supervision, combined with labeled data, to improve performance as in (Shnarch et al., 2018). Finally, similar to (Sen et al., 2020), the human-readable patterns matching an example can provide explainability for the model decision for it.

# Training Examples

## Positive examples (#1 - #100 out of 488 examples)

Dear Voucher Holder 2 claim your 1st class airport lounge passes when using Your holiday voucher call 08704439680 . When booking quote 1st class x 2

Congratulations ! Thanks to a good friend U have WON the Â£2,000 Xmas prize , 2 claim is easy , just call 08718726971 NOW ! Only 10p per minute . BT - national - rate .

Guess who am I?This is the first time I created a web page WWW.ASJESUS.COM read all I wrote . I'm waiting for your opinions . I want to be your

Prev

Page 1 out of 5

Next

## Negative examples (#1 - #100 out of 3,079 examples)

Great . I was getting worried about you . Just know that a wonderful and caring person like you will have only the best in life . Know that u r wonderful and God 's love is yours .

Watching cartoon , listening music & amp ; at eve had to go temple & amp ; church .. What about u ?

Huh so slow i tot u reach long ago liao ... U 2 more days only i 4 more leh ...

Prev

Page 1 out of 31

Next

Figure 3: The example-centric level 1 report listing all positive and negative examples in the training data. Words matched by at least one pattern are highlighted.

## Example ID: pos.1

Do you want a New Nokia 3510i colour phone DeliveredTomorrow ? With 300 free minutes to any mobile + 100 free texts + Free Camcorder reply or call 08000930705 (Label: pos)

### 14 positive pattern(s)

Pattern 6: The word "call" (call) , closely followed by a number (08000930705)

Pattern 9: The word "mobile" (mobile)

Pattern 12: A proper noun (a name of a specific individual, place, or object) (DeliveredTomorrow) , closely followed by a preposition (With) , and then by a number (300)

### 4 negative pattern(s)

Pattern 34: A preposition (With)

Pattern 39: A noun (colour) . closelv followed by a noun (ophone)

Figure 4: A part of the example-centric level 2 report showing positive and negative patterns matching an example.

## 5. Use Case: Model Analysis

In this use case, we show how GrASP library can assist the process of analyzing an NLP model, and specifically, finding its strength and weaknesses. While for some tasks, such as Parsing (Nilsson and Nivre, 2008) or Grammatical Error Correction (Choshen et al., 2020), exploring model's decisions naturally supports its anal-

ysis (e.g., which types of grammatical errors the model does not cover well), in many tasks this is not the case. We take Machine Translation (MT) as an example where characterizing model performance poses a challenge (Stanovsky et al., 2019; Renduchintala and Williams, 2021). We demonstrate how GrASP library can be used to identify both challenging and easy inputs for a given

MT model. An example of the practical usage of such analysis is found in translation systems which combine automatic MT and human translation. If we can identify that some types of inputs are translated well by the model, then we can refer them to the MT and manually translate only the rest. This will lower the system costs and aid in focusing human effort in the difficult translation cases and revision (Hutchins, 2001; Taravella and Villeneuve, 2013).

We use the WMT19 English-German Quality Estimation data (Fonseca et al., 2019) which consists of the automatic translation of 13K sentences by a specific model. Each translated sentence was judged by humans and was assigned with a translation quality score. We sort all sentences by this quality score and take the top 25% (inputs on which the model performed best) as the positive set for GrASP, and the bottom 25% as the negative set. We train GrASP with the default hyperparameters and output 100 patterns.<sup>5</sup>

First, we focus on indicative patterns in the positive set, i.e., the easiest inputs for the model. There are 60 highly precise patterns (precision  $\geq 80\%$ ), each with recall ranging from 5% up to 35%. Among those, we find various patterns which indicate that instruction texts are easy to translate. These patterns include the sign  $>$ , the lemmas *click* and *select*, the *hypernym of choose* (which matches both *choose* and *select*). For an example of such pattern and its matches, see lines 7–9 in Table 1. Furthermore, by examining these instruction patterns, we did notice that the dataset is composed of two specific domains (Reviews and IT), unlike several other machine translation datasets that consist of general domain data, such as news and Wikipedia (Barrault et al., 2020; Specia et al., 2020). This accidentally exemplifies the power of the exploration tool to teach us about the data we use.

Revealing which inputs are easy for the translation model is important, but so is the opposite. Various studies have analyzed what challenges translation models to aid their improvement (Macketanz et al., 2018; Choshen and Abend, 2019). Hence, we observe the negative-set patterns, i.e., the most challenging inputs. GrASP finds much less such patterns, suggesting that the challenging inputs are more diverse and harder to classify. Still, 16 patterns with precision larger than 60% are found.

Of specific interest to us are patterns that match inputs that are known in the literature to be challenging. We report two such examples. Gralinski et al. (2019) highlight the lemma *be* as an especially hard case. Indeed, GrASP dedicates a special pattern for this lemma. Another known difficult input is that of structural ambiguity (Avramidis et al., 2019). For example, *new blog entry*, where the translation model needs to decide which is new, the *blog entry* or the *blog* itself. GrASP finds the pattern `[[POS:ADJ], [POS:NOUN], [POS:NOUN]]`

<sup>5</sup>The full list of GrASP hyperparameters for each use case in this paper can be found in the pattern-centric level 1 report of the use case on our live demo website.

which captures these cases.

Overall, this use case shows that the library is beneficial for analyzing models by extracting known hurdles and identifying easy cases. Other uses for analysis may include different tasks or comparing models to each other.

## 6. Use Case: Dataset Artifacts

Artifacts in text classification datasets are tokens or phrases which are irrelevant, but frequently appear in examples of some of the classes. Hence, the trained models could rely on spurious correlations caused by these artifacts to make predictions. This is usually undesirable since it prevents the models from generalizing well out of the training distribution. In this use case, we leverage GrASP to help humans detect such artifacts in training datasets.

**SNLI.** Natural language inference is a task aiming to predict whether a hypothesis statement is true (entailment), false (contradiction), or undetermined (neither) given a premise statement. Using the pointwise mutual information (PMI), Gururangan et al. (2018) discovered annotation artifacts in many hypotheses of the Stanford Natural Language Inference (SNLI) dataset (Bowman et al., 2015). These artifacts make the fasttext classifier (Joulin et al., 2017) classify examples correctly for 67.0% without even looking at the premises (while the majority class baseline is only 34.3%).

We applied GrASP (with the number of patterns and the alphabet size of 200, using information gain as the selection criteria) to a subset of hypotheses in SNLI, i.e., 10K sentences from each of the entailment (positive) and contradiction (negative) classes. Then we kept only patterns with precision  $\geq 75\%$ . Among the final set of patterns, we found artifacts reported by Gururangan et al. (2018) yet with additional contexts and beyond. While PMI found that *outside* and *outdoors* are important words for the entailment class, GrASP discovered the pattern `[[POS:NOUN], [LEMMA:be], [HYPERNYM:outside]]` with the precision of 97% for the entailment class. This pattern covers both *outside* and *outdoors* (thanks to the *HYPERNYM* attribute) and further reveals their semantic role in the sentences. Another similar pattern, this time for the contradiction class, is `[[HYPERNYM:person], [HYPERNYM:be], [POS:VERB], [LEMMA:in, POS:ADP]]` (matching, e.g., *man/women/girl/hiker/Firefighters is/are lying/bored/swimming/reading in* as shown in the pattern-centric level 2 report). It shows that many contradicting hypotheses were crafted by indicating the actions (verbs) a human took and their locations (following the word *in*). Besides, the pattern `[[HYPERNYM:nutriment, POS:NOUN]]`, grouping *lunch, pizza, dessert, picnic*, etc. into one category, helps us quickly grasp the artifact theme. Interestingly, we also found from the pattern `[[SENTIMENT:neg, POS:VERB]]` that verbs with negative sentiments such as *attack, stole, lost* and *fell*

fairly correlate with the contradiction class (with 75.7% precision). All in all, a strength of GrASP is the ability to group words based on their semantics or functions. Together with the ordered representation of patterns, it provides better understanding of the artifacts.

**20Newsgroups** (Lang, 1995) is another dataset containing artifacts. Here, we aim to not only identify the artifacts but also partially remove them to make the classifier generalize better. Particularly, we focus on distinguishing the Christianity class from the two related classes, Atheism and Religion (miscellaneous). The goal is to improve generalization of a model trained on the 20Newsgroups dataset and tested on the Religion dataset (Ribeiro et al., 2016).

To detect the artifacts, we applied GrASP using Christianity as the positive class and the other two classes as the negative class. Generally, if the positive and the negative examples are not balanced, GrASP’s precision might be biased towards the majority class. Hence, we downsampled the larger class (non-Christianity) to reach the same size as the smaller class. We then kept only patterns with at least 75% precision, resulting in 133 patterns.

After that, as the patterns are human-interpretable, we manually checked them using the web exploration tool and flagged patterns that are semantically irrelevant to the classification task (i.e., likely to be artifacts). The pattern-centric view level 1 shows all the patterns to be annotated. When being unsure what a pattern covers, we used the pattern-centric view level 2 to see matched examples before making a decision. For example, one can see that the pattern `[[HYPERNYM:sacrament]]` covers *liturgy*, *baptism*, *confession*, and *Eucharist*, and accepting this pattern helps us annotate words or phrases in bulk. Overall, 40 out of 133 patterns were marked irrelevant. These include patterns for irrelevant names and surnames (e.g., `[[LEMMA:keith]]`, `[[LEMMA:mozumder]]`), for generic words (e.g., `[[LEMMA:cheers]]`, `[[LEMMA:recent]]`), and for punctuation (e.g., `[[LEMMA:>]]`, `[[LEMMA:#]]`).

With the artifact patterns identified, there could be several ways to address this problem and enhance model generalizability. Since the main objective of this paper is to demonstrate the applications of GrASP, we opted for simplicity and removed all the training examples that matched at least one artifact pattern. Then we used this filtered dataset to train a 1D Convolutional Neural Network (Kim, 2014) with 300-dimensional GloVe embeddings (Pennington et al., 2014). We compared this model with another model trained on a sample of the original dataset matching the filtered dataset in the number of positive and negative examples (to disregard the effect of the dataset size). For each model, we ran the train and test process five times and averaged the results, reported in Table 2.

We can see that, both models perform almost equally well on the in-domain test set (from 20Newsgroups). However, the model trained on the filtered dataset per-

Setting	Accuracy	Macro F1
Test dataset: 20Newsgroups		
Sampled	<b>0.810 ± 0.03</b>	<b>0.818 ± 0.03</b>
Filtered	0.800 ± 0.01	0.804 ± 0.01
Test dataset: Religion (Out-of-distribution)		
Sampled	0.674 ± 0.03	0.682 ± 0.03
Filtered	<b>0.725 ± 0.06</b>	<b>0.739 ± 0.05</b>

Table 2: The results of the data artifacts (20Newsgroups) use case (average of five repetitions ± SD).

forms better on the out-of-distribution test set (about 5% and 6% absolute difference for the accuracy and macro F1, respectively). This demonstrates that GrASP can help humans identify artifacts in training data which is a prerequisite step for building a more generalizable classification model.

## 7. Related Work

**Textual Data Exploration Tools.** While there is need for data exploration and interpretable patterns, not many tools exist. Most of the data exploration tools for text analyze the given data only at the word level (Heimerl et al., 2014; Zainol et al., 2018) or the n-gram level (Davies, 2012; Benoit et al., 2018). Even though there exist tools extracting complicated patterns (like regular expressions) (Bartoli et al., 2016; Locascio et al., 2016), the extracted patterns are not generalizable, i.e., having limited capability to match tokens based on their linguistic attributes. To the best of our knowledge, apart from our work, there is no other publicly available tool to extract and display recurring generalizable patterns from text. Somewhat related to ours, there are pattern-based search engines over data, even ones with web interfaces (Shlain et al., 2020; Resnik and Elkiss, 2005). However, such resources require an expert to come up with the patterns herself in order to query data. Unlike our library, those patterns are not learned by the resources as a helping mechanism for the data exploration. In other cases, patterns are not supplied by an expert. They are however, based on preexisting rules to extract specific information from the data, rather than to understand the data and what patterns occur in it (Valenzuela-Escarcega et al., 2020; Valenzuela-Escarcega et al., 2015).

Despite the fact that there are other algorithms except GrASP to learn patterns, such as (Sen et al., 2020), we could not find their tools or implementations publicly released. It is noteworthy that the outputs of such algorithms could be formatted and plugged into our web-based exploration tool as well.

**Debugging Tools for NLP.** Although GrASP was devised for textual pattern extraction and exploration in general, it can be applied to debug NLP models and datasets as shown in sections 5 and 6, respectively. Since debugging is a crucial and widely studied topic (Lertvittayakumjorn and Toni, 2021), there are several tools available to help humans debug NLP models and



datasets. Some of them are similar to our use cases. For example, Gralinski et al. (2019) propose GEval, which highlights features in inputs, outputs, or expected outputs that often make the model struggle. This helps users notice the model weaknesses. However, GEval does not consider the order of features in text, so it requires users’ effort to form patterns from feature combinations. Wu et al. (2019) propose Errudite, which extracts attributes from input instances and allows expert users to write a rule with the attributes so as to query and group model errors, enabling further model behavior analysis. Concerning dataset artifact discovery, existing works usually focus on word-level artifacts (Gururangan et al., 2018; Gardner et al., 2021). Meanwhile, when it comes to more complex artifacts, humans are required to notice or hypothesize those artifacts with or without supporting tools (McCoy et al., 2019; Han et al., 2020). GrASP, in contrast, learns and provides patterns which might be artifacts to humans, so they only need to check those patterns with the aid of our web-based exploration tool, leading to more efficiency in the artifact discovery process for complicated patterns.

## 8. Conclusion and Outlook

We presented GrASP – a Python library for extracting patterns from textual data together with the novel web-based tool for exploring the resulting patterns and the input examples. Also, we demonstrated the usefulness of both via three use cases in NLP research (classification, machine translation model analysis, and data artifacts identification). Apart from these, GrASP has potential to be applied to several other problems such as explainable text classification (where patterns provide explanations) (Efstathiadis et al., 2022; Ribeiro et al., 2018), human-AI model co-creation (with patterns enabling communication) (Yang et al., 2019), and interpreting deep learning models (e.g., by studying patterns that the models’ neurons capture as reflected by their activation scores) (Lertvittayakumjorn et al., 2020; Albini et al., 2020). Furthermore, the GrASP algorithm and the provided exploration tool are language-agnostic, so, as future work, it would be interesting to apply GrASP to a language besides English. To do so, users need to create custom attributes that augment input tokens with linguistic features of the target language. More generally, GrASP is applicable to other types of sequence data beyond text as long as we can extract suitable attributes for them (Agrawal and Srikant, 1995). Lastly, it is possible and interesting to extend our library with the functionalities of GrASP<sup>lite</sup> (Shnarch et al., 2020) — extracting patterns in an unsupervised way when we have a single list of texts rather than two lists (i.e., positive and negative classes) of texts.

## 9. Acknowledgements

The first author wishes to thank the support from Anandamahidol Foundation, Thailand. The last author was

partially supported by funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 101020934, ADIX, Argumentation-based Deep Interactive eXplanations).

## 10. Bibliographical References

- Agrawal, R. and Srikant, R. (1995). Mining sequential patterns. In *Proceedings of the Eleventh International Conference on Data Engineering*, pages 3–14.
- Albini, E., Lertvittayakumjorn, P., Rago, A., and Toni, F. (2020). Dax: Deep argumentative explanation for neural networks. *arXiv preprint arXiv:2012.05766*.
- Almeida, T. A., Hidalgo, J. M. G., and Yamakami, A. (2011). Contributions to the study of sms spam filtering: new collection and results. In *Proceedings of the 11th ACM symposium on Document engineering*, pages 259–262.
- Avramidis, E., Macketanz, V., Strohriegel, U., and Uszkoreit, H. (2019). Linguistic evaluation of German-English machine translation using a test suite. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 445–454, Florence, Italy, August. Association for Computational Linguistics.
- Barrault, L., Biesialska, M., Bojar, O., Costa-jussà, M. R., Federmann, C., Graham, Y., Grundkiewicz, R., Haddow, B., Huck, M., Joanis, E., Kocmi, T., Koehn, P., Lo, C.-k., Ljubešić, N., Monz, C., Morishita, M., Nagata, M., Nakazawa, T., Pal, S., Post, M., and Zampieri, M. (2020). Findings of the 2020 conference on machine translation (WMT20). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1–55, Online, November. Association for Computational Linguistics.
- Bartoli, A., Lorenzo, A. D., Medvet, E., and Tarlao, F. (2016). Inference of regular expressions for text extraction from examples. *IEEE Transactions on Knowledge and Data Engineering*, 28(5):1217–1230, May.
- Benoit, K., Watanabe, K., Wang, H., Nulty, P., Obeng, A., Müller, S., and Matsuo, A. (2018). quanteda: An r package for the quantitative analysis of textual data. *Journal of Open Source Software*, 3(30):774.
- Bowman, S. R., Angeli, G., Potts, C., and Manning, C. D. (2015). A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal, September. Association for Computational Linguistics.
- Choshen, L. and Abend, O. (2019). Automatically extracting challenge sets for non-local phenomena in neural machine translation. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 291–303, Hong Kong, China, November. Association for Computational Linguistics.



- Choshen, L., Nikolaev, D., Berzak, Y., and Abend, O. (2020). Classifying syntactic errors in learner language. In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 97–107, Online, November. Association for Computational Linguistics.
- Davies, M. (2012). Word and phrase. [www.wordandphrase.info](http://www.wordandphrase.info). Accessed: 2022-01-06.
- Efstathiadis, I. S., Paulino-Passos, G., and Toni, F. (2022). Explainable patterns for distinction and prediction of moral judgement on reddit. *arXiv preprint arXiv:2201.11155*.
- Ein-Dor, L., Shnarch, E., Dankin, L., Halfon, A., Sznajder, B., Gera, A., Alzate, C., Gleize, M., Choshen, L., Hou, Y., Bilu, Y., Aharonov, R., and Slonim, N. (2020). Corpus wide argument mining—a working solution. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7683–7691, Apr.
- Fonseca, E., Yankovskaya, L., Martins, A. F. T., Fishel, M., and Federmann, C. (2019). Findings of the WMT 2019 shared tasks on quality estimation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 1–12, Florence, Italy, August. Association for Computational Linguistics.
- Gardner, M., Merrill, W., Dodge, J., Peters, M., Ross, A., Singh, S., and Smith, N. A. (2021). Competency problems: On finding and removing artifacts in language data. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1801–1813, Online and Punta Cana, Dominican Republic, November. Association for Computational Linguistics.
- Gralinski, F., Wróblewska, A., Stanislawek, T., Grabowski, K., and Górecki, T. (2019). Geval: Tool for debugging nlp datasets and models. In *ACL 2019*.
- Gururangan, S., Swayamdipta, S., Levy, O., Schwartz, R., Bowman, S., and Smith, N. A. (2018). Annotation artifacts in natural language inference data. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Han, X., Wallace, B. C., and Tsvetkov, Y. (2020). Explaining black box predictions and unveiling data artifacts through influence functions. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5553–5563, Online, July. Association for Computational Linguistics.
- Hearst, M. A. (1992). Automatic acquisition of hyponyms from large text corpora. In *COLING 1992 Volume 2: The 15th International Conference on Computational Linguistics*.
- Heimerl, F., Lohmann, S., Lange, S., and Ertl, T. (2014). Word cloud explorer: Text analytics based on word clouds. In *2014 47th Hawaii International Conference on System Sciences*, pages 1833–1842. IEEE.
- Houvardas, J. and Stamatatos, E. (2006). N-gram feature selection for authorship identification. In *International conference on artificial intelligence: Methodology, systems, and applications*, pages 77–86. Springer.
- Hu, M. and Liu, B. (2004). Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177.
- Hutchins, J. (2001). Machine translation and human translation: in competition or in complementation. *International Journal of Translation*, 13(1-2):5–20.
- Joulin, A., Grave, E., Bojanowski, P., and Mikolov, T. (2017). Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain, April. Association for Computational Linguistics.
- Kim, Y. (2014). Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar, October. Association for Computational Linguistics.
- Lang, K. (1995). Newsweeder: Learning to filter news. In *Proceedings of the Twelfth International Conference on Machine Learning*, pages 331–339.
- Lawrence, J. and Reed, C. (2020). Argument mining: A survey. *Computational Linguistics*, 45(4):765–818.
- Lertvittayakumjorn, P. and Toni, F. (2021). Explanation-Based Human Debugging of NLP Models: A Survey. *Transactions of the Association for Computational Linguistics*, 9:1508–1528, 12.
- Lertvittayakumjorn, P., Specia, L., and Toni, F. (2020). FIND: Human-in-the-Loop Debugging Deep Text Classifiers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 332–348, Online, November. Association for Computational Linguistics.
- Lesk, M. (1986). Automatic sense disambiguation using machine readable dictionaries: how to tell a pine cone from an ice cream cone. In *Proceedings of the 5th annual international conference on Systems documentation*, pages 24–26.
- Levy, R., Gretz, S., Sznajder, B., Hummel, S., Aharonov, R., and Slonim, N. (2017). Unsupervised corpus-wide claim detection. In *Proceedings of the 4th Workshop on Argument Mining*, pages 79–84, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Locascio, N., Narasimhan, K., DeLeon, E., Kushman, N., and Barzilay, R. (2016). Neural generation of regular expressions from natural language with minimal domain knowledge. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1918–1923, Austin, Texas,

- November. Association for Computational Linguistics.
- Macketanz, V., Avramidis, E., Burchardt, A., and Uszkoreit, H. (2018). Fine-grained evaluation of German-English machine translation based on a test suite. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 578–587, Belgium, Brussels, October. Association for Computational Linguistics.
- Madnani, N., Heilman, M., Tetreault, J., and Chodorow, M. (2012). Identifying high-level organizational elements in argumentative discourse. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 20–28, Montréal, Canada, June. Association for Computational Linguistics.
- McCoy, T., Pavlick, E., and Linzen, T. (2019). Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy, July. Association for Computational Linguistics.
- Nilsson, J. and Nivre, J. (2008). Malteval: an evaluation and visualization tool for dependency parsing. In *LREC*. Citeseer.
- Peng, Y., Torii, M., Wu, C. H., and Vijay-Shanker, K. (2014). A generalizable nlp framework for fast development of pattern-based biomedical relation extraction systems. *BMC bioinformatics*, 15(1):1–18.
- Pennington, J., Socher, R., and Manning, C. (2014). GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October. Association for Computational Linguistics.
- Renduchintala, A. and Williams, A. (2021). Investigating failures of automatic translation in the case of unambiguous gender. *ArXiv*, abs/2104.07838.
- Resnik, P. and Elkiss, A. (2005). The linguist’s search engine: An overview. In *ACL*.
- Ribeiro, M. T., Singh, S., and Guestrin, C. (2016). “why should i trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’16, page 1135–1144, New York, NY, USA. Association for Computing Machinery.
- Ribeiro, M. T., Singh, S., and Guestrin, C. (2018). Anchors: High-precision model-agnostic explanations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Rinott, R., Dankin, L., Alzate Perez, C., Khapra, M. M., Aharoni, E., and Slonim, N. (2015). Show me your evidence - an automatic method for context dependent evidence detection. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 440–450, Lisbon, Portugal, September. Association for Computational Linguistics.
- Sen, P., Danilevsky, M., Li, Y., Brahma, S., Boehm, M., Chiticariu, L., and Krishnamurthy, R. (2020). Learning explainable linguistic expressions with neural inductive logic programming for sentence classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4211–4221, Online, November. Association for Computational Linguistics.
- Shlain, M., Taub-Tabib, H., Sadde, S., and Goldberg, Y. (2020). Syntactic search by example. In *ACL*.
- Shnarch, E., Levy, R., Raykar, V., and Slonim, N. (2017). GRASP: Rich patterns for argumentation mining. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1345–1350, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Shnarch, E., Alzate, C., Dankin, L., Gleize, M., Hou, Y., Choshen, L., Aharonov, R., and Slonim, N. (2018). Will it blend? blending weak and strong labeled data in a neural network for argumentation mining. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 599–605, Melbourne, Australia, July. Association for Computational Linguistics.
- Shnarch, E., Choshen, L., Moshkovich, G., Aharonov, R., and Slonim, N. (2020). Unsupervised expressive rules provide explainability and assist human experts grasping new domains. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2678–2697, Online, November. Association for Computational Linguistics.
- Specia, L., Blain, F., Fomicheva, M., Fonseca, E., Chaudhary, V., Guzmán, F., and Martins, A. F. T. (2020). Findings of the wmt 2020 shared task on quality estimation. In *Proceedings of the Fifth Conference on Machine Translation*, pages 743–764, Online, November. Association for Computational Linguistics.
- Stanovsky, G., Smith, N. A., and Zettlemoyer, L. (2019). Evaluating gender bias in machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1679–1684, Florence, Italy, July. Association for Computational Linguistics.
- Taravella, A. and Villeneuve, A. O. (2013). Acknowledging the needs of computer-assisted translation tools users: the human perspective in human-machine translation. *The Journal of Specialised Translation*, 19(January):62–74.
- Valenzuela-Escarcega, M. A., Hahn-Powell, G., Surdeanu, M., and Hicks, T. (2015). A domain-independent rule-based framework for event extraction. In *ACL*.
- Valenzuela-Escarcega, M. A., Hahn-Powell, G., and Bell, D. (2020). Odinson: A fast rule-based information extraction framework. In *LREC*.

- Wu, T., Ribeiro, M. T., Heer, J., and Weld, D. (2019). Errudite: Scalable, reproducible, and testable error analysis. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 747–763, Florence, Italy, July. Association for Computational Linguistics.
- Yang, Y., Kandogan, E., Li, Y., Sen, P., and Lasecki, W. S. (2019). A study on interaction in human-in-the-loop machine learning for text analytics. In *IUI Workshops*.
- Zainol, Z., Jaymes, M. T., and Nohuddin, P. N. (2018). Visualurtext: a text analytics tool for unstructured textual data. In *Journal of Physics: Conference Series*, volume 1018, page 012011. IOP Publishing.