arXiv:2006.00937v1 [cs.LG] 1 Jun 2020

# Concept Matching for Low-Resource Classification

**Federico Errica** [1] **Ludovic Denoyer** [2] **Bora Edizel** [2] **Fabio Petroni** [2]
**Vassilis Plachouras** [2] **Fabrizio Silvestri** [2] **Sebastian Riedel** [2 3]

## Abstract

We propose a model to tackle classification tasks in the presence of very little training data. To this aim, we approximate the notion of exact match with a theoretically sound mechanism that computes a probability of matching in the input space. Importantly, the model learns to focus on elements of the input that are relevant for the task at hand; by leveraging highlighted portions of the training data, an error boosting technique guides the learning process. In practice, it increases the error associated with relevant parts of the input by a given factor. Remarkable results on text classification tasks confirm the benefits of the proposed approach in both balanced and unbalanced cases, thus being of practical use when labeling new examples is expensive. In addition, by inspecting its weights, it is often possible to gather insights on what the model has learned.

## 1. Introduction

Gathering and labeling data is a task that can be expensive in terms of time, human effort and resources. When practitioners cannot rely on public large datasets, training a model with acceptable performance on a few data points becomes critical in a variety of applications. It is not uncommon that the data is also imbalanced, and as such the demands of gathering samples of the minority class are high. A natural domain in which these issues arise is, for instance, text classification, with notable tasks being hate-speech (Waseem & Hovy, 2016) and abuse detection (Mishra et al., 2018). For these reasons, the study of techniques that address this problem can have a tangible impact on society.

One effective approach to overcome the lack of training data is to augment the elements of the input with extra annotations, which has proved to be effective when coupled with feature engineering approaches (Zaidan et al., 2007; Zaidan & Eisner, 2008). Such annotations, e.g., highlighted words in a sentence, serve the purposes of guiding the learning process toward good solutions and to prevent overfitting the scarce amount of training samples. The goal of this work is to investigate this idea from a pure representation learning perspective, where there is no human intervention on the raw data but for the extra annotations.

To tackle this challenge, we design an architecture that learns to extract relevant *semantic* concepts from each input sample, such as words in a sentence or nodes in a graph. We assume each input is made by a set of individual representations: in scenarios like natural language processing where words are the main constituents of the input, we can rely on unsupervised pre-trained methods to represent them as vectors (Bojanowski et al., 2017; Devlin et al., 2018; Hu* et al., 2020). As we act solely on the model, the technique is flexible and task-agnostic; this is in contrast with task-dependent feature engineering methods (Zaidan & Eisner, 2008). Here, the task is assumed to be new, and as such labels need be (slowly) gathered by someone with domain-specific expertise.

In particular, we introduce a new mechanism to match concepts in each input sample and an effective error "boosting" technique to exploit the additional annotations. We also provide a theoretical analysis that justifies the choice of our matching mechanism; on the empirical side, we will see how cheap annotation costs can make up for a much larger number of training samples, that is a *desiderata* for low-resource classification.
Additionally, in this scenario, it is of practical importance to have some degree of reassurance about what the model has learned; by direct inspection of the weights, we show how it is possible to gain human-readable insights about its decision process. Results across a consistent number of baselines indicate a significant improvement in performance with respect to neural competitors as well as foundational methods that make use of the given annotations.

To summarize, we make the following contributions: (i) We introduce PARCUS, a new architecture that effectively combines concept matching and error boosting techniques for

[1]Department of Computer Science, University of Pisa, Pisa, Italy [2]Facebook AI [3]Department of Computer Science, University College London, London, United Kingdom. Correspondence to: Federico Errica <federico.errica@phd.unipi.it>.

low-resource classification; (ii) We support our intuition with a theoretical analysis; (iii) We empirically validate the approach against a consistent number of baselines, demonstrating strong performance improvements; (iv) We perform ablation studies to disentangle the contributions of the architecture main constituents; (v) Qualitative analyses show that the model works according to intuition and can be inspected to gain insights into what it has learned.

The rest of the paper is structured as follows: Section 2 reviews the existing literature; Section 3 formally introduces our model; Section 4 details our experiments and discusses our findings; Section 5 analyzes limitations future works; finally, Section 6 summarizes our work.

## 2. Related Works

There are different ways in which extra annotations can be used. Some works *generate* annotations as a way to interpret the model, while others *exploit* them to inform the learning process. Natural language processing is the field in which these techniques have been investigated the most. In particular, the method proposed by Lei et al. (2016) tackles text classification by learning the distribution of annotations given the text and that of the target class given the annotations. Interestingly, an additional regularization term is added to the loss to produce annotations that are short and coherent. The model makes use of high-capacity recurrent neural networks (Schuster & Paliwal, 1997), thus it is tested on large amounts of training data to prevent overfitting. This work was later refined by Bastings et al. (2019), who proposed a probabilistic version of a similar architecture, where a latent model is responsible for the generation of *discrete* annotations. The main advantage of predicting discrete annotations is that it is possible to constrain their maximum number per sample, thus effectively controlling sparsity. However, it usually requires a large number of data points to be effective.

The first to exploit annotations (also called *rationales* in this case) in a low resource scenario were Zaidan et al. (2007) and Zaidan & Eisner (2008), by means of a rationale-constrained SVM (Cortes & Vapnik, 1995) and a probabilistic model. Moreover, the latter is realized as a log-linear classifier that makes heavy use of feature-engineering. On the other hand, when annotations are defined on features rather than on samples, one can use the Generalized Expectation (GE) criteria (Druck et al., 2007; McCallum et al., 2007) to improve the performance of classifiers.

Annotations can also be incorporated in the loss function as done in Barrett et al. (2018), where an attention module (Vaswani et al., 2017) on top of an LSTM (Hochreiter & Schmidhuber, 1997) is forced to attend words in a document. A similar approach has been successfully applied by Bao et al. (2018) to the weak supervision problem. However, the model assumes one *source domain*, with supervised labels, to learn an attention generation module that is then applied to the target domain. In contrast, our method can be built on a given embedding space with minimum supervision.

Apart from incorporating prior knowledge in the form of annotations, we also mention other ways in which neural networks can be augmented: first-order logic (Li & Srikumar, 2019; Hu et al., 2016a); a corpora of regular expressions (Luo et al., 2018); or massive linguistic constraints (Hu et al., 2016b). While generally powerful and effective, all these methods require domain-specific expertise to define the additional features and constraints that have to be explicitly incorporated into the network; our method, instead, is designed to be task-agnostic. In a different manner, the SoPA architecture of Schwartz et al. (2018) learns to match surface patterns on text through a differentiable version of finite state machines, which relies on fixed-length and linear-chain patterns to classify a document. Instead, BabbleLabble (BL) (Hancock et al., 2018) is a method for generating weak classifiers from natural language explanations when supervision is scarce. On the one hand, BL works well because it exploits a domain-specific grammar to parse explanations; on the other hand, this grammar must be carefully designed by domain experts.

Perhaps the most similar to this work, the Neural Bag Of Words (NBOW) model (Kalchbrenner et al., 2014) takes an average of the elements belonging to an input sample and applies a logistic regression to classify a document. Its extension, NBOW2 (Sheikh et al., 2016), computes an importance score for each word by comparing it with a single reference vector that is learned. Despite the underlying idea being similar, we propose a more general mechanism to focus on relevant words and use the given annotations.

As a final remark, notice that our setting is substantially different from the more common literature on few-shot learning (Snell et al., 2017; Garcia & Bruna, 2018; Chen et al., 2019), where the goal is to classify classes that were *unseen* at training time. Here, we use annotations to be able to associate concepts with the right class, something which must be known in advance for the method to work properly.

In the following, we describe the architecture. As we shall discuss, the model has a strong inductive bias that reflects our intuition about how a model should work in the absence of large amounts of data. Hereinafter, we refer to our new architecture with the name PARCUS (the Latin word for "*parsimonious*").

## 3. The PARCUS model

Let us consider a classification task in which a very small labelled dataset $\mathcal{D} = \{(x_1, r_1, y_1), \ldots, (x_L, r_L, y_L)\}$ is

given. Here, $x_i$ is an input sample, $r_i$ represents the extra (optional) annotations provided by a human, and $y_i$ is a discrete target label. For the purpose of this paper, an input is a set of tokens $x_i = \{x_i^1, \ldots, x_i^{T_i}\}$ of arbitrary size $T_i$. In addition, $x_i^j \in \mathbb{R}^n$, where $n$ is the size of an embedding space. Finally, each token in the training set may be marked as relevant or not by annotators, i.e., $r_i = \{r_i^1, \ldots, r_i^{T_i}\} \in \{0, 1\}^{T_i}$.

### 3.1. Intuition

When humans are asked to solve a classification problem after seeing a few examples, they tend to look for very simple patterns across the dataset, and text classification is an excellent use case. For instance, assume the word "excellent" is important to classify a movie review as positive; if we were to work in the character space, a straightforward solution would be to match specific (sub-)strings in the input, an instance of the so-called *pattern matching* technique. At the same time, however, humans are able to generalize to *semantically similar* concepts, and our goal is to exploit similar embeddings to reflect this ability.

In this work, we transfer the concept of pattern matching into the embedding space, where semantically similar words are assumed to have similar representations. We achieve this via a mechanism that outputs a probability of matching between an input token and a *prototype vector*, the latter of which is learned to capture discriminative concepts. Differently from bag-of-words methods of Section 2, our model can accommodate multiple prototypes and focus on concepts that are useful for the task.

Moreover, in order to guide the learning process using the extra annotations, it seems sensible to magnify the error for those tokens that have been marked as relevant by annotators. Notwithstanding the simplicity of the idea, the underlying challenge this work addresses is to effectively embed such human knowledge into the prototypes. In other words, each matching probability should be highly correlated with a particular target class. In addition, it would be desirable that a user could understand what the model has learned, something of great interest when working with the uncertainty caused by a scarce number of training samples. In this respect, we will provide a practical example in Section 4.5.

We now show how to compute and combine multiple matching probabilities, and then we introduce a technique to incorporate extra annotations in the training process. It is worth mentioning that both techniques have been designed to coexist, even though the latter is not strictly necessary. To provide a graphical representation of the proposed architecture, Figure 1 depicts a use case for text classification.

### 3.2. Concept Matching

We now present the core mechanism that implements concept matching. Let us define a set $\mathcal{P} = \{p_1, \ldots, p_N \mid p_i \in \mathbb{R}^n\}$ of prototypes to be learned, where $N$ is an hyper-parameter of the model. Each $p_i \in \mathcal{P}$ should ideally adapt to be similar (in the embedding space) to the representation of important tokens.

To learn the $N$ prototypes, we employ the *cosine similarity* metric. Cosine similarity has been often used to measure semantic similarity (Landauer & Dumais, 1997); its co-domain ranges from $-1$, i.e., opposite in meaning, to 1, i.e., same meaning, with 0 indicating uncorrelation. Ideally, we would like our prototypes to have high similarity with the relevant tokens in the input. To this aim, we further define an exponential activation function $g : [-1, 1] \rightarrow [0, 1]$ that takes the distance between a token $x_i^j$ and a prototype $p_k$ and outputs a probability of matching:

$$P(x_i^j \text{ matches } p_k) = g(d(x_i^j, p_k)) = a^{d(x_i^j, p_k)-1} \quad (1)$$

where $a$ is an hyper-parameter and $d(x_i^j, p_k)$ computes the cosine similarity between $x_i^j$ and $p_k$. In practice, the closer to 1 the similarity is, the greater the output of this gated activation, and $g(v) = 1 \Leftrightarrow v = 1$. By choosing a high value of $a$ we strongly penalize tokens that are associated with low similarity scores.

### 3.3. Combining Multiple Prototypes

As we saw, Equation 1 computes the matching probability between a token and a prototype. Because we have $N$ prototypes, we treat the associated $N$ probabilities as a feature vector for the input token, and we denote each feature as $\phi_k(x_i^j) = g(d(x_i^j, p_k)) \; \forall k \in 1, \ldots, N$. Interestingly, working with matching probabilities allows us to combine all of them through AND/OR/XOR logical functions. An approximation of such functions can be straightforwardly implemented through the pseudo-differentiable version of $min$ and $max$ (Paszke et al., 2019), though a fully differentiable version exists:

$$\phi_{AND}(x_i^j) = min(\{\phi_k(x_i^j) \quad \forall k\}) \quad (2)$$

$$\phi_{OR}(x_i^j) = max(\{\phi_k(x_i^j) \quad \forall k\}) \quad (3)$$

$$\phi_{XOR}(x_i^j) = \phi_{OR}(x_i^j) - \phi_{AND}(x_i^j) \quad (4)$$

In our experiments, the use of $min$ and $max$ functions significantly sped up convergence due to the absence of non-linearities. Moreover, we chose to augment $\Delta(x_i^j)$ with the probability of opposite matching: $\phi_{\neg k}(x_i^j) = g(-d(x_i^j, p_k)) \; \forall k \in 1, \ldots, N$. Specifically, when $\phi_{\neg k}(x_i^j) \approx 1$ it means the token $x_i^j$ and $p_k$ have opposite meaning. Finally, notice how our method differs from NBOW2 (Sheikh et al., 2016), as we use prototypes to compute per-token features rather than an importance score.

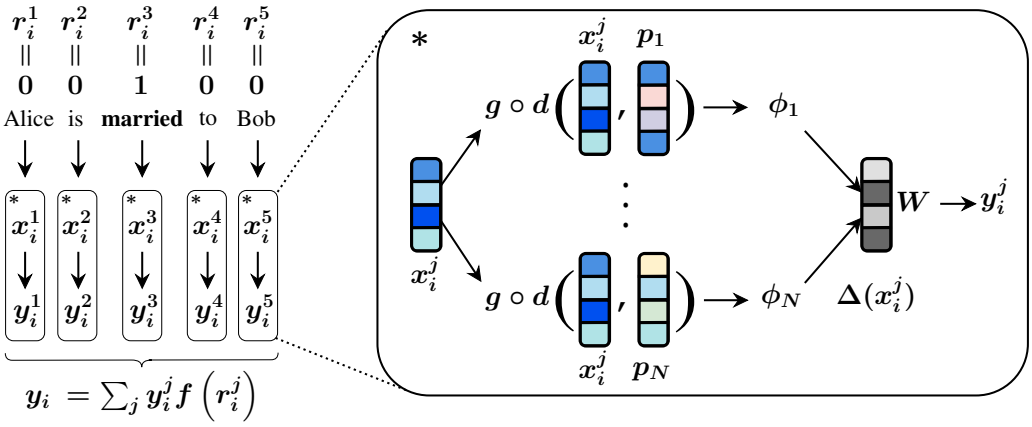


*Figure 1.* The PARCUS architecture is applied to the $i$-th example of a dataset, i.e., "Alice is married to Bob", with "married" being highlighted. We extract features by computing the similarity between the token's embedding and the prototypes of our model. Then, we combine these features with a linear layer that outputs per-token predictions. At training time only, predictions are multiplied by a boosting factor $f(r_i^j)$. Individual tokens' predictions are then summed to yield the sentence prediction $y_i$.

### 3.4. Inference

Once we have a feature vector for each token, we need to combine all $F = N + 3$ features to output a token prediction $y_i^j$. Let us first define an auxiliary term (omitting the argument $x_i^j$ to make notation less cluttered):

$$\Delta(x_i^j) = [\phi_1, \ldots, \phi_N, \phi_{AND}, \phi_{OR}, \phi_{XOR}] \quad (5)$$

where square brackets denote concatenation. Then, we compute token predictions by linearly combining features:

$$y_i^j = \Delta(x_i^j)\mathbf{W} + b \quad (6)$$

where $\mathbf{W} \in \mathbb{R}^{F \times C}$ is a matrix of parameters (multi-class prediction with $C$ classes) and $b$ is the (optional) bias. The linear model is useful when we want the user to analyze the importance given to each matching probability, as discussed in Section 3.1, as well as to restrict the number of parameters of the model (see discussion below). Finally, the input prediction is just a sum of the individual $y_i^j$

$$y_i = \sigma(\sum_j^{T_i} y_i^j), \quad (7)$$

where $\sigma$ is the *softmax* activation.

**Discussion** Regularization of the matrix $\mathbf{W}$ plays an important role to answer our research questions. We use both L1 and L2 regularization terms on $\mathbf{W}$, as done in (Zou & Hastie, 2005), for two main purposes. First, the L1 term enforces sparsity and discourages the mixing of too many concepts. Secondly, L2 limits the magnitude of the weights, hence avoiding over-compensation of low cosine similarity scores. Consequently, in order to increase one of the matching probability features, the model is encouraged to make changes to the prototypes rather than to the linear weights; In other words, relevant information for the task will be stored inside prototypes in the form of semantic embeddings.

### 3.5. Annotation-driven Error Boosting

So far, we have not made use of annotations, which are of fundamental importance to guide the learning process in low-resource scenarios. To learn prototypes that match relevant concepts, the proposed technique should weight the importance of tokens rather than whole samples. It follows a boosting approach (Freund et al., 1999) is not feasible in this scenario; instead, our method exploits prior information in an efficient way. The idea is to modify the error associated with specific tokens to encourage prototypes to be similar to them. To be more precise, *at training time* we modify Equation 7 to take into account the given annotations:

$$y_i = \sigma(\sum_j^{T_i} y_i^j \cdot f(r_i^j)), \quad (8)$$

where $f : [0, 1] \rightarrow \mathbb{R}$ is an arbitrary exponential function of our choice that boosts the error, e.g., $f(r_i^j) = e^{r_i^j}$. In terms of learning, $f(r_i^j)$ boosts the gradient of highlighted tokens while leaving unchanged the rest (i.e., if $r_i^j$ is 0, our $f(r_i^j)$ outputs a multiplicative factor of 1). From a mathematical standpoint, we cannot achieve the same result as Equation 8 by means of an additional loss term, as done in Lei et al. (2016), because gradients would be summed and not multiplied as done here.

### 3.6. Model complexity and inductive bias

We conclude with remarks on the model complexity. The total number of parameters is $\Theta(Nn + FC)$, which could be much larger than that of a linear model ($\Theta(nC)$) when $N$ is high and $C$ is small. Usually, a restricted number of parameters serves to counteract overfitting by limiting the hypotheses space of the model (Vapnik, 1998). However, this work tackles the problem from a novel perspective, as we prevent the prototype weights from *freely* changing. Specifically, prototype weights vary in a way that depends on the given embedding space, because the learning process makes them similar to some token $x_i^j$. If we allowed the weights to freely change, we would get something similar to an MLP; our experiments show how this way of constraining the weights fits particularly well the use case we are considering. Finally, notice that PARCUS ignores the structural dependencies between input tokens; this is intended, as it is not feasible to learn complex interactions with only a few data samples. Nonetheless, if semantic representations $x_i^j$ are obtained with a pre-trained model, they will usually carry some structural information as well.

### 3.7. Theoretical analysis

The choice behind the concept matching mechanism of Section 3.2 is backed up by a theoretical explanation. Indeed, in the limit of the gating parameter $a$, Equation 1 converges to the discontinuous Kronecker delta function $\delta(x_i^j, p_k)$ that is 1 when its arguments are equal and 0 otherwise; hence, Eq. 1 is a sound approximation of a "hard match" function.

**Proposition 3.1.** *Let $x, y \in \mathbb{R}^n$ and $d : \mathbb{R}^n \times \mathbb{R}^n \to (-\infty, 1]$ be a function such that $d(x, y) = 1 \iff x = y$. Then, the sequence of functions $\{f_a\}_{a>1}$ with $f_a = a^{d(x,y)-1}$ is pointwise convergent to $\delta(x, y)$.*

*Proof.* To prove pointwise convergence, it is sufficient to show that

$$\lim_{a \to +\infty} a^{d(x,y)-1} = \delta(x, y) \;\; \forall x, y.$$

Because the $d$ cannot take values greater than 1, it follows that $d(x, y)-1 \leq 0$, and the equivalence holds if and only if $x = y$. Therefore, $f_a(x, y) \to 0$ for $x \neq y$ and $f_a(x, y) \to 1$ when $x = y$, that is $\delta(x, y)$. □

From this proposition we can make another important consideration. Given that it is not possible to have *uniform* convergence to any discontinuous function, some parts of $\delta(x, y)$ will be approximated more easily than others. Specifically for cosine similarity, it can be shown that the area comprised between the two functions, i.e., the error of our approximation, is $\frac{a^2-1}{a^2 \ln a}$; nonetheless, reasonable values of $a$ guarantee good performances and stable learning curves in our experiments. In summary, this result reveals that the best we can do is to look for approximations that satisfy desirable properties, for example being more accurate near the discontinuity and more "permissive" elsewhere.

## 4. Experiments

This section reports the experimental setting as well as our experimental findings. We compare PARCUS against a large number of baselines. Additionally, we perform an in-depth analysis of our model through ablation studies and qualitative analyses of the effect of some hyper-parameters. Then, we consider a practical scenario in which a user wants to gather insights on how PARCUS predicts a class for each input sample. We use natural language processing benchmarks to validate our model, and all code to reproduce and extend our experiments is made available[1].

### 4.1. Experimental Setting

**Datasets** We empirically validate our method on two different datasets. First, the MOVIEREVIEW dataset (Zaidan et al., 2007) contains balanced positive and negative movie reviews with annotations. Secondly, we use the highly imbalanced (8% of positive samples) SPOUSE dataset from Hancock et al. (2018), where the task is to tell whether two entities in a given piece of news are married or not. This is a harder task than standard classification, as the same document can appear multiple times with different given entities and the background context greatly varies. Datasets statistics are reported in Table 1. We provide annotations for 60 randomly chosen positive samples of the SPOUSE dataset; this process is fast and aims at replicating real world scenarios where labels are scarce and hard to collect.

*Table 1.* Datasets' statistics.

| | Train | Valid. | Test | Annotations |
|---|---|---|---|---|
| SPOUSE | 22195 | 2796 | 2697 | 60 |
| MOVIEREVIEW | 1800 | - | 200 | 1800 |

**Setup** We measure performances on the given test set while varying the number of training data points. We use balanced train splits for all models; on MOVIEREVIEW, the validation set is taken as big as the training one to simulate a real scenario. As for SPOUSE, we use the given validation set for model selection to fairly compare with the results of Hancock et al. (2018). We chose the pre-trained (unsupervised) base version of BERT (Devlin et al., 2018) to provide an embedding space to our method and to other neural baselines.

[1] https://github.com/facebookresearch/parcus.

*Table 2.* Hyper-parameters tried during model selection.

| | LINEAR | MLP/NBOW(2)/DAN | BERT+FINETUNE | OURS |
|---|---|---|---|---|
| LEARNING RATE | $\{1e\text{-}2, 1e\text{-}3, 1e\text{-}4\}$ | $\{1e\text{-}3, 1e\text{-}4\}$ | $\{2e\text{-}5, 3e\text{-}5, 5e\text{-}5\}$ | $\{1e\text{-}2\}$ |
| L1 | - | - | - | $\{1e\text{-}2, 1e\text{-}3\}$ |
| L2 | $\{1e\text{-}1, 1e\text{-}2, 1e\text{-}4\}$ | $\{1e\text{-}2, 1e\text{-}4\}$ | - | $\{1e\text{-}3, 1e\text{-}4\}$ |
| EPOCHS | $\{50, 100, 150\}$ | $\{100, 500\}$ | $\{2, 4, 10\}$ | $\{500\}$ |
| HIDDEN UNITS | - | $\{8, 16, 32\}$ | - | - |
| BATCH SIZE | 32 | 32 | 8 | 32 |
| $N$ | - | - | - | $\{5, 10\}$ |
| $f(r)$ | - | - | - | $\{e^r, 5^r, 10^r\}$ |
| $a$ | - | - | - | $\{10, 100\}$ |

We repeat each experiment 10 times with different random splits; importantly, we train and validate different models on the same data splits. The hyper-parameters for (hold-out) model selection are reported in Table 2. Moreover, to avoid bad initializations of the final re-training with the selected configuration, we average test performances over 3 training runs. The optimized measure is Accuracy for MOVIEREVIEW and F1-score for SPOUSE. PARCUS is trained by gradient descent in an end-to-end fashion, from the prototypes to the linear weights. We optimize the Cross-Entropy loss using Adam (Kingma & Ba, 2015).

**Methods** To have a good comparison with embedding-based models other than those reported in the literature, we trained a linear model (Linear) and a single-layer MLP, as well as NBOW (Kalchbrenner et al., 2014), NBOW2 (Sheikh et al., 2016) and the Deep Averaging Network (DAN) of (Iyyer et al., 2015). We also fine-tune BERT using the suggested hyper-parameters (Devlin et al., 2018), adding 10 to the possible training epochs.

On SPOUSE, we devised a regular expression that associates specific sub-strings (*"wife", "husb", "marr"* and *"knot"*) to the positive class; ideally, models should be able to focus on such words but also generalize. Moreover, Traditional Supervision (TS) and Babble Labble (BL-DM) were taken from the work of Hancock et al. (2018): the former method is a logistic regression using n-gram features, whereas the latter is a complex pipeline tested on 30 natural language explanations provided by humans. Notably, BL-DM exploits the relational information of the SPOUSE dataset via task-specific grammar and parser, while PARCUS simply ignores sentences where the entities of interest are not present.

On MOVIEREVIEW, we also report results of an SVM (Zaidan et al., 2007) and a log-linear model on language features (Zaidan & Eisner, 2008), both of which are specifically designed to exploit additional annotations.

Finally, we performed a number of ablation studies to isolate the effect of different techniques: (i) an MLP with the error boosting technique (MLP-W. H.) to validate the use of prototypes; (ii) our method without highlights (PARCUS-WO H.) to assess the impact of rationales; (iii) our method with no logical features (PARCUS-NO-LOGIC); (iv) our method with $\phi_k$ features only (PARCUS-$\phi_k$); (v) our method with bilinear rather than cosine similarity (PARCUS-BILINEAR) to show the importance of constrained weights; (vi) Parcus where the input is the average of all input tokens (PARCUS-AVG); (vii) Parcus where centroids are pre-computed using the unsupervised k-means algorithm (PARCUS-KMEANS).

### 4.2. Results & Discussion

Table 3 presents all our empirical results, including the ablation studies. Results highlight that PARCUS has strong performances in a low data regime, validating intuition and theoretical results. On SPOUSE, our model strongly outperforms other neural baselines and reaches the *manually tuned* regular expression with just 60 training points. Moreover, TS needs $\approx$*50x* more data to achieve similar performance. We also found that TS performs much worse than our linear baseline (hence the need for a fair comparison in the embedding space). Surprisingly, only 10 data points are sufficient to perform better than almost all baselines with a training size of 300, a $>$*30x* improvement which does not depend on the chosen embedding space. With 300 datapoints and no annotations, our model has an average F1 score very close to that of BL-DM. Notice that the reported result (BL-DM, 46.5) is not averaged over multiple runs, and one of our random splits achieves a test score of 46.3; this indicates the need for robust evaluation when it comes to experimenting with few data points/natural language explanations. Overall, we found that the proposed approach can be helpful when data is greatly imbalanced and diverse in nature, and outperforms powerful models like BERT that are quite performing when fine-tuned on relatively small datasets (Devlin et al., 2018; Howard & Ruder, 2018).

Similar arguments apply to MOVIEREVIEW, where our model improves over the baselines. Interestingly, PARCUS is able to improve the state of the art by a large margin when very few data points are used. Here, NBOW and NBOW2 models proved to be the strongest competitors, as they rely on the mean representation of a document.

Overall, the gap is more evident as training size is very scarce, even when compared to other baselines that use extra annotations. This suggests the model could be a good fit for all those practical scenarios where the data gathering

*Table 3.* Results for all datasets. Standard deviation is shown in brackets. We report the F1-score as the evaluation metric for SPOUSE and the accuracy for MOVIEREVIEW.

SPOUSE

| MODEL/TRAIN SIZE | 10 | 30 | 60 | 150 | 300 | 3K | 10K | |
|---|---|---|---|---|---|---|---|---|
| TUNED REGEXP | - | - | - | - | - | - | - | 40.5 |
| TS | - | 15.5 | 15.9 | 16.4 | 17.2 | 41.8 | 55.0 | |
| BL-DM (30 EXPL.) | - | - | - | - | - | - | - | **46.5** |
| LINEAR | 18.2 (1.3) | 20.6 (1.4) | 22.5 (1.4) | 26.1 (1.1) | 26.1 (1.2) | - | - | - |
| MLP | 17.9 (2.4) | 20.2 (3.1) | 18.3 (0.6) | 23.3 (1.2) | 24.1 (1.3) | - | - | - |
| NBOW | 21.0 (2.3) | 21.8 (1.7) | 24.0 (1.0) | 27.4 (2.0) | 28.2 (1.8) | - | - | - |
| NBOW2 | 19.5 (2.6) | 22.3 (1.9) | 25.9 (1.4) | 29.6 (1.5) | 31.7 (2.1) | - | - | - |
| DAN | 21.8 (3.2) | 24.1 (2.5) | 26.6 (1.7) | 28.2 (1.6) | 29.2 (1.5) | - | - | - |
| BERT+FINETUNING | 16.9 (2.6) | 20.2 (2.1) | 23.4 (1.2) | 32.1 (2.0) | 35.5 (3.2) | - | - | - |
| (ABL.) MLP W. H. | 16.7 (1.4) | 20.8 (2.7) | 20.9 (1.6) | 22.7 (1.8) | 23.1 (2.0) | - | - | - |
| (ABL.) PARCUS-WO H. | 27.0 (2.2) | 31.6 (2.5) | 34.2 (2.3) | 41.8 (2.1) | **44.0** (1.2) | - | - | - |
| (ABL.) PARCUS-$\phi_k$ | 32.4 (4.5) | 34.4 (4.2) | 37.8 (2.7) | 42.7 (1.0) | 41.4 (2.4) | - | - | - |
| (ABL.) PARCUS-NO-LOGIC | 32.7 (3.4) | 34.5 (3.9) | 36.8 (2.6) | 42.7 (1.6) | 42.0 (1.9) | - | - | - |
| (ABL.) PARCUS-AVG | 22.9 (3.7) | 26.5 (2.9) | 28.8 (2.2) | 30.5 (1.1) | 32.7 (0.9) | - | - | - |
| (ABL.) PARCUS-KMEANS | 30.3 (2.0) | 33.5 (0.5) | 32.93 (1.0) | 32.8 (0.9) | 34.2 (1.3) | - | - | - |
| (ABL.) PARCUS-BILINEAR | 29.1 (4.5) | 31.4 (5.9) | 36.0 (5.4) | 36.1 (5.1) | 33.1 (3.0) | - | - | - |
| PARCUS | **34.0** (4.5) | **36.6** (4.3) | **40.3** (2.5) | **43.7** (1.7) | 42.9 (1.6) | - | - | - |

MOVIEREVIEW

| MODEL/TRAIN SIZE | 10 | 20 | 50 | 100 | 200 |
|---|---|---|---|---|---|
| SVM + RATIONALES | - | 65.4 | - | 75 | 83.2 |
| LOG-LINEAR + RATIONALES | - | 65.8 | - | 76 | 83.8 |
| LINEAR | 60.4 (3.4) | 64.0 (3.5) | 70.2 (2.0) | 77.2 (2.6) | 80.3 (3.1) |
| MLP | 59.1 (4.1) | 62.6 (4.2) | 69.7 (2.4) | 73.3 (3.8) | 80.0 (3.0) |
| NBOW | 62.6 (4.6) | 65.7 (4.8) | 73.9 (1.6) | 78.0 (2.0) | 81.2 (3.6) |
| NBOW2 | 61.5 (4.5) | 64.3 (4.9) | 72.9 (1.4) | 78.9 (4.4) | 83.6 (1.8) |
| DAN | 61.5 (6.2) | 62.3 (4.8) | 72.9 (3.3) | 78.7 (3.2) | 82.35 (2.7) |
| BERT+FINETUNING | 53.5 (2.0) | 54.8 (4.9) | 59.7 (4.5) | 67.7 (4.3) | 79.2 (2.5) |
| (ABL.) MLP W. H. | 61.5 (4.6) | 63.1 (5.8) | 68.9 (7.0) | 72.4 (8.5) | 74.6 (5.8) |
| (ABL.) PARCUS-WO H. | 61.2 (4.3) | 64.9 (5.0) | 74.3 (2.4) | 78.6 (2.3) | **84.6** (2.8) |
| (ABL.) PARCUS-$\phi_k$ | 66.1 (5.7) | 68.4 (3.5) | **77.8** (2.0) | 80.7 (3.0) | 83.4 (2.4) |
| (ABL.) PARCUS-NO-LOGIC | 66.9 (5.9) | 67.9 (3.5) | 75.5 (4.0) | **81.0** (2.4) | 83.7 (2.7) |
| (ABL.) PARCUS-AVG | 62.1 (4.9) | 62.5 (4.4) | 71.0 (3.5) | 73.3 (3.1) | 79.0 (3.4) |
| (ABL.) PARCUS-KMEANS | 54.4(5.0) | 53.2 (3.4) | 54.2 (2.6) | 53.6 (2.7) | 58.0 (2.4) |
| (ABL.) PARCUS-BILINEAR | 57.5 (5.1) | 61.9 (6.7) | 70.4 (3.7) | 75.3 (2.9) | 78.3 (3.6) |
| PARCUS | **67.2** (5.5) | **70.1** (5.6) | 76.6 (2.4) | 80.0 (2.6) | 83.8 (2.8) |

process is just started and one wants to boost performances by means of extra annotations.

### 4.3. Ablation Studies

We performed ablation studies on both datasets to understand whether the improvements are only due to prototypes, error boosting technique or both. Overall, we observe that the use of prototypes provides a consistent improvement with respect to the other baseline, and this is especially evident on the SPOUSE dataset. Interestingly, MLP W. H. does not benefit from error boosting, which is in accord with the fact that unconstrained weights make it more difficult to select and isolate the contribution of relevant tokens. In addition, it seems that the logical and opposite matching features can help to boost the average performance, as PARCUS-$\phi_k$ and PARCUS-NO-LOGIC always perform worse than PARCUS on SPOUSE. Because annotations guide the learning process, these are most important in the extremely low resource scenario, but their effect slowly fades as the training size increases; contrarily to our expectations, PARCUS performs even better on larger amounts of training points without annotations. This indicates that, at a certain point, annotations may regularize the model too much, and it suggests future works on *adaptive* error boosting functions. Finally, note that neither averaging tokens nor pre-computing centroids seem beneficial; indeed, models like DAN better exploit the average using an MLP on top of the averaged representation, while we force the model to align to some relevant input token. Also, the use of pre-computed centroids will make the model focus on the most common semantics in the dataset, which are not necessarily the most adequate to solve the task.

#### 4.3.1. MORE GENERAL DISTANCE FUNCTIONS

In Section 3.6 and in the above discussion, we argued that the inductive bias of our architecture is favorable for the specific problem we are tackling. Here, we empirically validate our statement by showing that the use of a more general distance function $d$ tends to overfit the data and achieves significantly worse performances. In particular, we substitute the cosine similarity with its bilinear counterpart $d_{W_b}(x, p) = \tanh(x^T W_b p)$, where $W_b \in \mathbb{R}^d \times \mathbb{R}^d$,

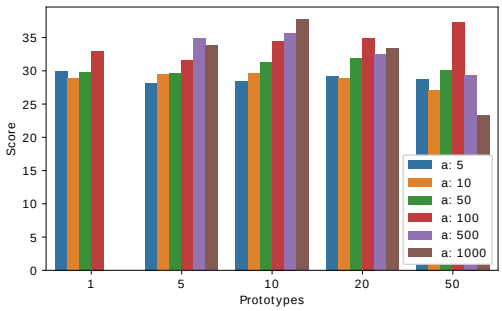


*Figure 2.* Here, 60 data points are used to train PARCUS on the SPOUSE dataset. Our analysis reveals how larger values of $a$ should be associated with a reasonable number of prototypes (much less than the size of the training set) to achieve good performances.

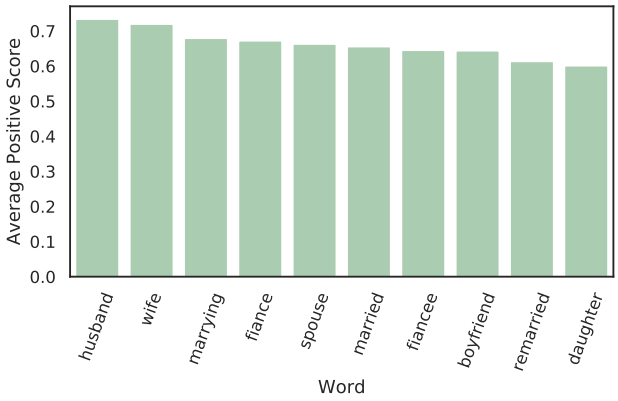


*Figure 3.* Top-10 most relevant tokens for positive prediction, averaged on *unseen* data.

and we ran the experiments on SPOUSE and MOVIEREVIEW (shown in Table 3 as PARCUS-BILINEAR). Bilinear similarity can be seen as a generalization of cosine similarity when individual features are given different importance (specified by the matrix $W_b$). However, the number of parameters is quadratic in the dimension of the given embeddings, and this matrix is unconstrained, unlike prototypes. Overall, we observe that the use of bilinear similarity still yields good performances on the SPOUSE task, but it is not capable of generalizing well on MOVIEREVIEW where the average number of tokens in each sentence is much higher. The reason may be that since SPOUSE contains pieces of news related to different topics, focusing solely on those concepts related to marriage may help.
These empirical results reinforce the belief that constraining the weights to match specific concepts in a low-resource scenario helps to generalize to new instances.

### 4.4. Qualitative analysis on the effect of $a$

The parameter $a$ plays an important role in controlling how strict the model is in considering a matching to be highly probable. Larger values of $a$ should produce prototypes that are more specific to a single concept, while smaller values (but still greater than 1, see Proposition 3.1) allow a prototype to match less similar tokens. To further confirm our intuition, we run an experiment on the SPOUSE dataset where we analyzed the trade-off between the value of $a$ and the number of prototypes. Figure 2 shows our results for 60 data points. We immediately see that using just 1 prototype with a large value of $a$ may be too restrictive to solve the task, which is in accord with common sense. However, the general trend we observe is that enforcing separation of concepts is usually beneficial, provided the number of prototypes is sufficiently high.

### 4.5. Gaining insights from the learned weights

The learned weights of the proposed model can be inspected to gain insights on what concepts it focuses on and how they are related. To show this, we train a model using $N$=3 prototypes on 60 examples taken from the SPOUSE dataset. Then, we rank the tokens' outputs of sentences belonging to unseen data, so that the outputs with the highest rank correspond to semantic concepts that have been considered relevant for the task by the model. Specifically, as shown in Figure 3, the model learns to focus on words related to marriage, as well as syntactic variations associated with similar semantics. Importantly, some of the words were not given as annotations in the training set, meaning that the model is also able to recognize similar concepts.
We additionally show that annotators' knowledge has been effectively incorporated into the prototypes, and how the features of Equation 5 have been combined together. We start by inspecting the magnitude of the linear weights $\mathbf{W} \in R^{F\times 2}$; specifically, if the $i$-th feature is discriminative for a class $c$, then the $i$-th row of $\mathbf{W}$ will have the $c$-th element larger than the others. In our example, we find that $\phi_1$ was important for positive predictions, whereas the other features did not contribute much to a particular class. We then perform top-10 cosine similarity ranking between tokens and the prototype $p_1$. From the most similar to the least one, we obtain: *husband*; *marriage*; *marrying*; *wife*; *married*; *marry*; *fiance*; *wedding*; *fiancee*; and *girlfriend*. This result gives insights on how PARCUS has learned to match concepts similar to those provided in natural language form by BL-DM (see Appendix of Hancock et al. (2018)).

## 5. Limitations and future works

Though PARCUS performs very well and its learning dynamics follow our intuition, there are some inherent limitations to the method. The first is that it is not possible to uniformly approximate the Kronecher delta function of Proposition 3.1, and as such we can only study further approximations that work better around the discontinuity. The second is the need to map the input into embedding space before training, which can be restrictive for less common application domains. This has an impact on how easily we can inspect the weights as done in Section 4.5; however, all domains for which a pre-trained method exists should benefit from our technique. Also, notice that cosine similarity is just one of the functions that can be used: if we are interested in the magnitude of the vectors when computing similarities, a normalized Euclidean norm can be a valid choice. Interesting future works will be the investigation of PARCUS performance on larger training sets and its extension to an adaptive version of the error boosting function $f$.

## 6. Conclusions

In this work, we presented PARCUS, a new representation learning methodology to perform classification in the low data regime. We coupled matching probabilities with error boosting to focus on concepts that are important for the task at hand. After comparing it with a large number of baselines, the model performed very well and outperformed most of them. We provided theoretical insights on the design of our matching technique, and we make an in-depth analysis of some characteristics of the model as well as many ablation studies. Moreover, we showed with a practical example that the weights can be inspected to see what concepts the model focuses on. In summary, our model can be very useful in tasks where gathering data is challenging, and it can be used to assist users in training a classifier for a very specific task.

## References


Bao, Y., Chang, S., Yu, M., and Barzilay, R. Deriving machine attention from human rationales. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 1903–1913, 2018.

Barrett, M., Bingel, J., Hollenstein, N., Rei, M., and Søgaard, A. Sequence classification with human attention. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pp. 302–312, 2018.

Bastings, J., Aziz, W., and Titov, I. Interpretable neural predictions with differentiable binary variables. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL*, 2019.

Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146, 2017.

Chen, W.-Y., Liu, Y.-C., Kira, Z., Wang, Y.-C. F., and Huang, J.-B. A closer look at few-shot classification. In *International Conference on Learning Representations*, 2019.

Cortes, C. and Vapnik, V. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

Druck, G., Mann, G., and McCallum, A. Reducing annotation effort using generalized expectation criteria. Technical report, Massacusetts University, Department of Computer Science, Amherst, 2007.

Freund, Y., Schapire, R., and Abe, N. A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence*, 14(771-780):1612, 1999.

Garcia, V. and Bruna, J. Few-shot learning with graph neural networks. In *International Conference on Learning Representations*, 2018.

Hancock, B., Varma, P., Wang, S., Bringmann, M., Liang, P., and Ré, C. Training classifiers with natural language explanations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pp. 1884–1895, 2018.

Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

Howard, J. and Ruder, S. Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 328–339, 2018.

Hu*, W., Liu*, B., Gomes, J., Zitnik, M., Liang, P., Pande, V., and Leskovec, J. Strategies for pre-training graph neural networks. In *International Conference on Learning Representations*, 2020.

Hu, Z., Ma, X., Liu, Z., Hovy, E., and Xing, E. Harnessing deep neural networks with logic rules. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pp. 2410–2420, 2016a.

Hu, Z., Yang, Z., Salakhutdinov, R., and Xing, E. Deep neural networks with massive learned knowledge. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1670–1679, 2016b.

Iyyer, M., Manjunatha, V., Boyd-Graber, J., and Daumé III, H. Deep unordered composition rivals syntactic methods for text classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pp. 1681–1691, 2015.

Kalchbrenner, N., Grefenstette, E., and Blunsom, P. A convolutional neural network for modelling sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 655–665, 2014.

Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

Landauer, T. and Dumais, S. A solution to plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological review*, 104(2):211, 1997.

Lei, T., Barzilay, R., and Jaakkola, T. Rationalizing neural predictions. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 107–117, 2016.

Li, T. and Srikumar, V. Augmenting neural networks with first-order logic. *arXiv preprint arXiv:1906.06298*, 2019.

Luo, B., Feng, Y., Wang, Z., Huang, S., Yan, R., and Zhao, D. Marrying up regular expressions with neural networks: A case study for spoken language understanding. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2083–2093, 2018.

McCallum, A., Mann, G., and Druck, G. Generalized expectation criteria. *Computer science technical note, University of Massachusetts, Amherst, MA*, 94(95):159, 2007.

Mishra, P., Yannakoudakis, H., and Shutova, E. Neural character-based composition models for abuse detection. In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*, pp. 1–10, 2018.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pp. 8024–8035, 2019.

Schuster, M. and Paliwal, K. K. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681, 1997.

Schwartz, R., Thomson, S., and Smith, N. A. Bridging cnns, rnns, and weighted finite-state machines. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL*, 2018.

Sheikh, I., Illina, I., Fohr, D., and Linares, G. Learning word importance with the neural bag-of-words model. In *ACL, Representation Learning for NLP (Repl4NLP) workshop*, 2016.

Snell, J., Swersky, K., and Zemel, R. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pp. 4077–4087, 2017.

Vapnik, V. Statistical learning theory wiley. *New York*, pp. 156–160, 1998.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.

Waseem, Z. and Hovy, D. Hateful symbols or hateful people? predictive features for hate speech detection on twitter. In *Proceedings of the NAACL student research workshop*, pp. 88–93, 2016.

Zaidan, O., Eisner, J., and Piatko, C. Using annotator rationales to improve machine learning for text categorization. In *Human language technologies 2007: The conference of the North American chapter of the association for computational linguistics; proceedings of the main conference*, pp. 260–267, 2007.

Zaidan, O. F. and Eisner, J. Modeling annotators: A generative approach to learning from annotator rationales. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 31–40. Association for Computational Linguistics, 2008.

Zou, H. and Hastie, T. Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)*, 67(2):301–320, 2005.