

Co-training for Predicting Emotions with Spoken Dialogue Data

Beatriz Maeireizo and Diane Litman and Rebecca Hwa

Department of Computer Science

University of Pittsburgh

Pittsburgh, PA 15260, U.S.A.

beamt@cs.pitt.edu, litman@cs.pitt.edu, hwa@cs.pitt.edu

Abstract

Natural Language Processing applications often require large amounts of annotated training data, which are expensive to obtain. In this paper we investigate the applicability of Co-training to train classifiers that predict emotions in spoken dialogues. In order to do so, we have first applied the wrapper approach with Forward Selection and Naïve Bayes, to reduce the dimensionality of our feature set. Our results show that Co-training can be highly effective when a good set of features are chosen.

1 Introduction

In this paper we investigate the automatic labeling of spoken dialogue data, in order to train a classifier that predicts students' emotional states in a human-human speech-based tutoring corpus. Supervised training of classifiers requires annotated data, which demands costly efforts from human annotators. One approach to minimize this effort is to use Co-training (Blum and Mitchell, 1998), a semi-supervised algorithm in which two learners are iteratively combining their outputs to increase the training set used to re-train each other and generate more labeled data automatically. The main focus of this paper is to explore how Co-training can be applied to annotate spoken dialogues. A major challenge to address is in reducing the dimensionality of the many features available to the learners.

The motivation for our research arises from the need to annotate a human-human speech corpus for the ITSPOKE (Intelligent Tutoring SPOKE dialogue System) project (Litman and Silliman, 2004). Ongoing research in ITSPOKE aims to recognize emotional states of students in order to build a spoken dialogue tutoring system that automatically predicts and adapts to the student's emotions. ITSPOKE uses supervised learning to predict emotions with spoken dialogue data. Although a large set of dialogues have been collected, only 8% of them have been annotated (10 dialogues with a total of 350 utterances), due to

the laborious annotation process. We believe that increasing the size of the training set with more annotated examples will increase the accuracy of the system's predictions. Therefore, we are looking for a less labour-intensive approach to data annotation.

2 Data

Our data consists of the student turns in a set of 10 spoken dialogues randomly selected from a corpus of 128 qualitative physics tutoring dialogues between a human tutor and University of Pittsburgh undergraduates. Prior to our study, the 453 student turns in these 10 dialogues were manually labeled by two annotators as either "Emotional" or "Non-Emotional" (Litman and Forbes-Riley, 2004). Perceived student emotions (e.g. confidence, confusion, boredom, irritation, etc.) were coded based on both what the student said and how he or she said it. For this study, we use only the 350 turns where both annotators agreed on the emotion label. 51.71% of these turns were labeled as Non-Emotional and the rest as Emotional.

Also prior to our study, each annotated turn was represented as a vector of 449 features hypothesized to be relevant for emotion prediction (Forbes-Riley and Litman, 2004). The features represent acoustic-prosodic (pitch, amplitude, temporal), lexical, and other linguistic characteristics of both the turn and its local and global dialogue context.

3 Machine Learning Techniques

In this section, we will briefly describe the machine learning techniques used by our system.

3.1 Co-training

To address the challenge of training classifiers when only a small set of labeled examples is available, Blum and Mitchell (1998) proposed *Co-training* as a way to bootstrap classifiers from a large set of unlabeled data. Under this framework, two (or more) learners are trained iteratively in tandem. In each iteration, the learners classify more unlabeled data to increase the training data

for each other. In theory, the learners must have distinct views of the data (i.e., their features are conditionally independent given the label example), but some studies suggest that Co-training can still be helpful even when the independence assumption does not hold (Goldman, 2000).

To apply Co-training to our task, we develop two *high-precision* learners: Emotional and Non-Emotional. The learners use different features because each is maximizing the precision of its label (possibly with low recall). While we have not proved these two learners are conditionally independent, this division of expertise ensures that the learners are different. The algorithm for our Co-training system is shown in Figure 1. Each learner selects the examples whose predicted labeled corresponds to its expertise class with the highest confidence. The maximum number of iterations and the number of examples added per iteration are parameters of the system.

```

While iteration < MAXITERATION
  Emo_Learner.Train(train)
  NE_Learner.Train(train)

  emo_Predictions = Emo_Learner.Predict(predict)
  ne_Predictions = NE_Learner.Predict(predict)

  emo_sorted_Predictions = Sort_by_confidence(
    emo_Predictions)
  ne_sorted_Predictions = Sort_by_confidence(
    ne_Predictions)

  best_emo = Emo_Learner.select_best(
    emo_sorted_Predictions,
    NUM_SAMPLES_TO_ADD)
  best_ne = NE_Learner.select_best(
    ne_sorted_Predictions,
    NUM_SAMPLES_TO_ADD)

  train = train ∪ best_emo ∪ best_ne
  predict = predict - best_emo - best_ne
end

```

Figure 1. Algorithm for Co-training System

3.2 Wrapper Approach with Forward Selection

As described in Section 2, 449 features have been currently extracted from each utterance of the ITSPOKE corpus (where an utterance is a student’s turn in a dialogue). Unfortunately, high dimensionality, i.e. large amount of input features, may lead to a large variance of estimates, noise, overfitting, and in general, higher complexity and inefficiencies in the learners. Different approaches have been proposed to address this problem. In this work, we have used the Wrapper Approach with Forward Selection.

The Wrapper Approach, introduced by John et al. (1994) and refined later by Kohavi and John (1997), is a method that searches for a good subset of relevant features using an induction algorithm as part of the evaluation function. We can apply

different search algorithms to find this set of features.

Forward Selection is a greedy search algorithm that begins with an empty set of features, and greedily adds features to the set. Figure 2 shows our algorithm implemented for the forward wrapper approach.

```

bestFeatures = []
while dim(bestFeatures) < MINFEATURES
  for iterations = 1: MAXITERATIONS
    split train into training/development
    parameters = computeParameters(training)
    for feature = 1:MAXFEATURES
      evaluate(parameters,development,
        [bestFeatures + feature])

      keep validation performance
    end
  end
  average_performance and keep average_performance
end
B = best average_performance
bestFeatures ← B ∪ bestFeatures
end

```

Figure 2. Implemented algorithm for forward wrapper approach. The variables underlined are the ones whose parameters we have changed in order to test and improve the performance.

We can use different criteria to select the feature to add, depending on the object of optimization.

Earlier, we have explained the basis of the Co-training system. When developing an expert learner in one class, we want it to be correct most of the time when it guesses that class. That is, we want the classifier to have high precision (possibly at the cost of lower overall accuracy). Therefore, we are interested in finding the best set of features for precision in each class. In this case, we are focusing on Emotional and Non-Emotional classifiers.

Figure 3 shows the formulas used for the optimization criterion on each class. For the Emotional Class, our optimization criterion was to maximize the PPV (Positive Predictive Value), and for the Non-Emotional Class our optimization criterion was to maximize the NPV (Negative Predictive Value).

		MODEL	
		1	0
TARGET	1	TP	FN
	0	FP	TN

$$PPV = \frac{TP}{TP + FP}$$

$$NPV = \frac{TN}{TN + FN}$$

Figure 3. Confusion Matrix, Positive Predictive Value (Precision for Emotional) and Negative Predictive Value (Precision for Non-Emotional)

4 Experiments

For the following experiments, we fixed the size of our training set to 175 examples (50%), and the size of our test set to 140 examples (40%). The remaining 10% has been saved for later experiments.

4.1 Selecting the features

The first task was to reduce the dimensionality and find the best set of features for maximizing the PPV for Emotional class and NPV for Non-Emotional class. We applied the Wrapper Approach with Forward Selection as described in section 3.2, using Naïve Bayes to evaluate each subset of features.

We have used 175 examples for the training set (used to select the best features) and 140 for the test set (used to measure the performance). The training set is randomly divided into two sets in each iteration of the algorithm: One for training and the other for development (65% and 35% respectively). We train the learners with the training set and we evaluate the performance to pick the best feature with the development set.

Number of Features	Naïve Bayes	AdaBoost-j48 Decision Trees
All Features	74.5 %	83.1 %
3 best for PPV	92.9 %	92.9 %

Table 1. Precision of Emotional with all features and 3 best features for PPV using Naïve Bayes (used for Feature Selection) and AdaBoost-j48 Decision Trees (used for Co-training)

The selected features that gave the best PPV for Emotional Class are 2 lexical features and one acoustic-prosodic feature. By using them we increased the precision of Naïve Bayes from 74.5% (using all 449 features) to 92.9%, and of AdaBoost-j48 Decision Trees from 83.1% to 92.9% (see Table 1).

Number of Features	Naïve Bayes	AdaBoost-j48 Decision Trees
All Features	74.2 %	90.7 %
1 best for NPV	100.0 %	100.0 %

Table 2. Precision of Non-Emotional with all features and best feature for NPV using Naïve Bayes (used for Feature Selection) and AdaBoost-j48 Decision Trees (used for Co-training)

For the Non-Emotional Class, we increased the NPV of Naïve Bayes from 74.2% (with all features) to 100% just by using one lexical feature,

and the NPV of AdaBoost-j48 Decision Trees from 90.7% to 100%. This precision remained the same with the set of 3 best features, one lexical and two non-acoustic prosodic features (see Table 2).

These two set of features for each learner are disjoint.

4.2 Co-training experiments

The two learners are initialized with only 6 labeled examples in the training set. The Co-training system added examples from the 140 “pseudo-labeled” examples¹ in the Prediction Set. The size of the training set increased in each iteration by adding the 2 best examples (those with the highest confidence scores) labeled by the two learners. The Emotional learner and the Non-Emotional learner were set to work with the set of features selected by the wrapper approach to optimize the precision (PPV and NPV) as described in section 4.1.

We have applied Weka’s (Witten and Frank, 2000) AdaBoost’s version of j48 decision trees (as used in Forbes-Riley and Litman, 2004) to the 140 unseen examples of the test set for generating the learning curve shown in figure 4.

Figure 4 illustrates the learning curve of the accuracy on the test set, taking the union of the set of features selected to label the examples. We used the 3 best features for PPV for the Emotional Learner and the best feature for NPV for the Non-Emotional Learner (see Section 4.1). The x-axis shows the number of training examples added; the y-axis shows the accuracy of the classifier on test instances. We compare the learning curve from Co-training with a baseline of majority class and an upper-bound, in which the classifiers are trained on human-annotated data. Post-hoc analyses reveal that four incorrectly labeled examples were added to the training set: example numbers 21, 22, 45, and 51 (see the x-axis). Shortly after the inclusion of example 21, the Co-training learning curve diverges from the upper-bound. All of them correspond to Non-Emotional examples that were labeled as Emotional by the Emotional learner with the highest confidence.

The Co-training system stopped after adding 58 examples to the initial 6 in the training set because the remaining data cannot be labeled by the learners with high precision. However, as we can see, the training set generated by the Co-training technique can perform almost as well as the upper-bound, even if incorrectly labeled examples are included in the training set.

¹ This means that although the example has been labeled, the label remains unseen to the learners.




Figure 4. Learning Curve of Accuracy using best features for Precision of Emotional/Non-Emotional

5 Conclusion

We have shown Co-training to be a promising approach for predicting emotions with spoken dialogue data. We have given an algorithm that increased the size of the training set producing even better accuracy than the manually labeled training set, until it fell behind due to its inability to add more than 58 examples.

We have shown the positive effect of selecting a good set of features optimizing precision for each learner and we have shown that the features can be identified with the Wrapper Approach.

In the future, we will verify the generalization of our results to other partitions of our data. We will also try to address the limitation of noise in our Co-training System, and generalize our solution to a corresponding corpus of human-computer data (Litman and Forbes-Riley, 2004). We will also conduct experiments comparing Co-training with other semi-supervised approaches such as self-training and Active learning.

6 Acknowledgements

Thanks to R. Pelikan, T. Singliar and M. Hauskrecht for their contribution with Feature Selection, and to the NLP group at University of Pittsburgh for their helpful comments. This research is partially supported by NSF Grant No. 0328431.

References

- A. Blum and T. Mitchell. 1998. *Combining Labeled and Unlabeled Data with Co-training*. Proceedings of the 11th Annual Conference on Computational Learning Theory: 92-100.
- K. Forbes-Riley and D. Litman. 2004. *Predicting Emotion in Spoken Dialogue from Multiple Knowledge Sources*. Proceedings of Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT/NAACL).
- S. Goldman and Y. Zhou. 2000. *Enhancing Supervised Learning with Unlabeled Data*. International Joint Conference on Machine Learning, 2000.
- G. H. John, R. Kohavi and K. Pleger. 1994. *Irrelevant Features and the Subset Selection Problem*. Machine Learning: Proceedings of 11th International Conference:121-129, Morgan Kaufmann Publishers, San Francisco, CA.
- R. Kohavi and G. H. John. 1997. *Wrappers for Feature Subset Selection*. Artificial Intelligence, Volume 97, Issue 1-2.
- D. J. Litman and K. Forbes-Riley, 2004. *Annotating Student Emotional States in Spoken Tutoring Dialogues*. Proc. 5th Special Interest Group on Discourse and Dialogue Workshop on Discourse and Dialogue (SIGdial).
- D. J. Litman and S. Silliman, 2004. *ITSPOKE: An Intelligent Tutoring Spoken Dialogue System*. Companion Proceedings of Human Language Technology conf. of the North American Chapter of the Association for Computational Linguistics (HLT/NAACL).
- I. H. Witten and E. Frank. 2000. *Data Mining: Practical Machine Learning Tools and Techniques with Java implementations*. Morgan Kaufmann, San Francisco.