

# Accountable Error Characterization

Amita Misra, Zhe Liu and Jalal Mahmud

IBM-Research, Almaden

San Jose, CA, USA

amita.misra1|liuzh|jumahmud@ibm.com

## Abstract

Customers of machine learning systems demand accountability from the companies employing these algorithms for various prediction tasks. Accountability requires understanding of system limit and condition of erroneous predictions, as customers are often interested in understanding the incorrect predictions, and model developers are absorbed in finding methods that can be used to get incremental improvements to an existing system. Therefore, we propose an accountable error characterization method, AEC, to understand when and where errors occur within the existing black-box models. AEC, as constructed with human-understandable linguistic features, allows the model developers to automatically identify the main sources of errors for a given classification system. It can also be used to sample for the set of most informative input points for a next round of training. We perform error detection for a sentiment analysis task using AEC as a case study. Our results on the sample sentiment task show that AEC is able to characterize erroneous predictions into human understandable categories and also achieves promising results on selecting erroneous samples when compared with the uncertainty-based sampling.

## 1 Introduction

As machine learning is becoming the method of choice for many analytics functionalities in industry, it becomes crucial to be able to understand the limits and risks of the existing models. In favour of more accurate AI, the availability of computational resources is coupled with increasing dataset sizes that has resulted in more complex models. Complex models suffer from lack of transparency, which leads to low trust as well as the inability to fix or improve the models output easily. Deep learning algorithms are among the highly accurate and complex models. Most users of deep learning models often treat them as a black box because of

its incomprehensible functions and unclear working mechanism (Liu et al., 2019). However, customers' retention requires accountability for these systems (Galitsky, 2018). Interpreting and understanding what the model has learned, as well as the limits and the risks of the existing model have therefore become a key ingredient of a robust validation (Montavon et al., 2018).

One line of research on model accountability examines the information learned by the model itself to probe the linguistic aspects of language learnt by the models (Shi et al., 2016; Adi et al., 2017; Giulianelli et al., 2018; Belinkov and Glass, 2019; Liu et al., 2019). Other line of research gives machine learning models the ability to explain or to present their behaviours in understandable terms to humans (Doshi-Velez and Kim, 2017) to make the predictions more transparent, and trustworthy. However, very few studies set the focus on error characterization as well as automatic error detection and mitigation. To address the above-mentioned gaps in characterizing model limits and risks, we seek to improve a model's behavior by categorizing incorrect predictions using explainable linguistic features. To accomplish that, we propose a framework called Accountable Error Characterization (AEC) to explain the predictions of a neural network model by constructing an explainable error classifier. The most similar work to ours is by (Nushi et al., 2018). They build interpretable decision-tree classifiers for summarizing failure conditions using human and machine generated features. In contrast, our approach builds upon incorrect predictions on a separate set to provide insights into model failure.

The AEC framework has three key components: A base neural network model, an error characterization model, and a set of interpretable features that serve as the input to the error characterization model. The features used in the error characterization model are based on explainable linguistic and lexical features such as dependency relations, and

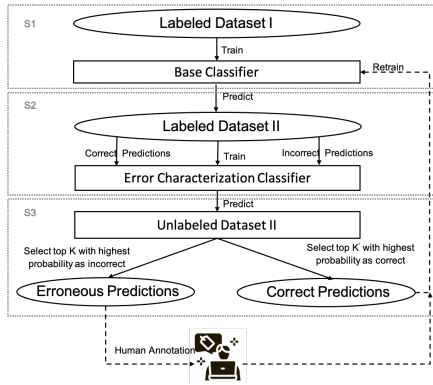


Figure 1: The overall workflow of AEC. Dashed lines represent planned future work

various lexicons that have been inspired by prior art, which allows the users and model developers to identify when a model fails. The error characterization model also offer rankings of informative features to provide insight into where and why the model fails.

By adding the error classification step on top of the base model, AEC can also be adopted to identify the highly confident error cases as the most informative samples for the next round of training. Although uncertainty based sampling can also be adopted to get the most informative samples (Lewis, 1995; Cawley, 2011; Shao et al., 2019), as it selects the examples with the least confidence, Ghai et al. (2020) show that uncertainty sampling led to an increasing challenge for annotators to provide correct labels. AEC avoids such problem by learning from error cases from a validation set. Our results show that AEC outperforms the uncertainty based sampling in terms of selecting erroneous predictions on a sample sentiment dataset (see Table 4).

We first present the overall framework of AEC to construct the error classifier, followed by the experiments and result. Finally, we conclude the paper with future directions and work in progress.

## 2 Explainable Framework

Figure 1 summarizes our overall method for constructing a human understandable classifier that can be used to explain the erroneous predictions of a deep neural network classifier and thus to improve the model performance. Our method consists of the following steps:

**S1:** Train a neural network based classifier using labeled dataset I, call it as the base classifier.

**S2:** Apply the base classifier on another labeled dataset II to get correct and incorrect prediction cases, based on which train a second 2-class error identification classifier with a set of human understandable features. Note here labeled dataset I and II can be in the same domain or in different domains.

**S3:** Rank the features according to their individual predictive power. Apply the error identification classifier from step 2, to a set of unlabeled data from the same domain as labeled dataset II and rank the unlabeled instances according to their prediction probability of being erroneous. These represent the most informative samples that can be further used in an active learning setting.

The focus of the current work is to identify and characterize the error cases of a base classifier in an human understandable manner. The following two sections describe the experiments and implementation of the framework using a sentiment prediction task as case study. The integration of these samples into an iterative training set up is a work in progress for future extension.

## 3 Machine Learning Experiments and Results

### 3.1 Data

We adopt a cross-domain sentiment analysis task as case study in this section to demonstrate the AEC method, although the proposed method would also be applicable to datasets from the same domain. We chose the cross-domain sentiment analysis task here as it is a challenging, but necessary task within the NLP domain and there are high chances of observing erroneous predictions. We use data from two different domains, Stanford Sentiment Treebank (SST) (Socher et al., 2013) (Labeled Dataset I) to train the base classifier, and a conversational Kaggle Airlines dataset (Labeled + Unlabeled Dataset II) to build and evaluate the error characterization classifier. The conversation domain represents a new dataset seeking an improvement on the base classifier trained using sentiment reviews.

**SST dataset:** A dataset of movie reviews annotated at 5 levels (very negative, negative, neutral, positive, and very positive). Sentence level annotations are extracted using the python package *pytree-*

DataSet	Negative	Neutral	Positive
SST	3304	1622	3605

Table 1: SST dataset distribution

DataSet	Negative	Neutral	Positive
Airline	7366	2451	1847

Table 2: Airline dataset distribution

bank<sup>1</sup>. We merged the *negative* and *very-negative* class labels into a single negative class, *positive* and *very-positive* into a single positive class, keeping neutral as it is. A preprocessing step to remove near duplicates gives a training set distribution as shown in Table 1. This is the only dataset used to train the base classifier.

**Twitter Airline Dataset:** The dataset is available through the library *Crowdfunder’s Data for Everyone*.<sup>2</sup> Each tweet is classified as either positive, neutral, or negative. The label distribution for the Twitter Airline is shown in Table 2.

### 3.2 Train the Base Classifier

We chose Convolution Neural Network (CNN) as a showcase here, as the base sentiment classifier to be trained using the SST dataset. However, the framework can be easily adapted to more advanced state of the art classifiers such as BERT (Devlin et al., 2019). A multi-channel CNN architecture is employed to train as it has been shown to work well on multiple sentiment datasets including SST (Kim, 2014). The samples are weighted to account for class imbalance.

### 3.3 Train the Error Characterization Classifier

We next applied the trained base classifier on the training set of a cross-domain dataset as described in Table 2 to get the predictions on a sample of 11664 labeled instances of Airlines dataset. Predictions from the base model on this Airlines dataset are further divided into two classes based on the ground truth test labels, correct-prediction and incorrect-prediction. The base classifier has an overall accuracy of 60.09% on the Airline dataset as shown in Table 3. A balanced set is created by undersampling the correct predictions giving a dataset of total 9310 instances. We use a 80/20 split for training and testing giving a training set of 7448 and a test set of 1862 instances. This train

<sup>1</sup><https://pypi.org/project/pytreebank>

<sup>2</sup><https://appen.com/resources/datasets/>

set serves as the input to train the error characterization classifier with erroneous or not as labels and different collections of explainable features as independent variables. A random forest classifier using a 5-fold cross validation was used to train the error characterization classifier. (Pedregosa et al., 2011).

Dataset	Total instances	Correct pred.	InCorrect pred.
Airline dataset	11664	7009	4655

Table 3: Performance of the Base classifier on the Airline dataset

### 3.3.1 Features

Our features have been inspired by previous work on sentiment, disagreement, and conversations. The feature values are normalized by sentence length.

**Generalized Dependency.** Dependency relations are obtained using the python package *spacy*<sup>3</sup>. Relations are generalized by replacing the words in each dependency relation by their corresponding POS tag (Joshi and Penstein-Rosé, 2009; Abbott et al., 2011; Misra et al., 2016).

**Emotion.** Count of words in each of the 8 emotion classes from the NRC emotion lexicon (anger, anticipation, disgust, fear, joy, negative, positive, sadness, surprise, and trust) available from (Mohammad and Turney, 2010).

**Named Entities.** The count of named entities of each entity type obtained from the python package *spacy*.

**Conversation.** Lexical indicators indicating greetings, thank, apology, second person reference, questions starting with do, did, can, could, with who, what, where as described by (Oraby et al., 2017).

### 3.4 Predict erroneous predictions from unlabeled data

Once the error characterization classifier was trained with the set of correctly and incorrectly predicted instances, we then apply it to the 20% test set of the Twitter Airline data, which consists of a total of 1862 instances as described in section 3.3. We selected the top K instances with the highest probability of being incorrectly predicted as the erroneous cases. We hide the actual labels on this

<sup>3</sup><https://spacy.io>

test set when selecting the instances. The actual labels will be later used to evaluate the performance of the error characterization classifier.

## 4 Evaluation and Results

In terms of identifying erroneous predictions, in our evaluation, we compare the performance of AEC with uncertainty-based sampling, in which the learner computes a probabilistic output for each sample, and select the samples that the base classifier is the most uncertain about based on probability scores.

### 4.1 Most informative samples for labeling.

As we are interested in generating a ranking of incorrect predictions for the base classifier from error characterization classifier, we use *precision at top k* as the evaluation metrics in here, which is a commonly used metric in information retrieval, and defined as  $P@K=N/K$ , where N is the actual number of errors samples among top K predicted. We compare the performance of the error characterization classifier and the uncertainty based sampling on the test set of 1832 instances as shown in Table 4. It shows the precision at top K where K varies from 10 to 50. For the first initial 10 samples, the uncertainty based sampling performs marginally better but as we select more samples (rows 2-5) the proposed approach starts outperforming the baseline.

TOP K	uncertainty-based P@K	AEC P@K
10	0.8	0.7
20	0.75	0.8
30	0.77	0.83
40	0.75	0.83
50	0.74	0.76

Table 4: Comparison of uncertainty-based sampling (Baseline) with proposed AEC on the test set.

### 4.2 Feature Characterization

When using uncertainty based sampling, it is not always evident why certain samples were selected, or how these samples map to actual errors of the base classifier. In contrast, AEC framework incorporates explainability into sample selection by mapping highly ranked feature sets from the error characterization model with the selected error samples.

S.No	Text	Base Pred.	Actual Label	Error. Prob
1	@username <b>if you could</b> change your name to @southwestair and <b>do what</b> they do...that'd be <b>awesome</b> . Also this plane <b>smells</b> like onion rings.	Neutral	Negative	0.84
2	@username now on hold for <b>90 minutes</b>	Neutral	Negative	0.82
3	@username user is a <b>compassionate</b> professional! Despite the flight challenges she made passengers feel like priorities!!	Neutral	Positive	0.79

Table 5: A subset of most informative samples for the Base classifier based on error characterization classifier probability score for the error class.

Table 5 shows a few examples of actual errors from the base classifier that are also predicted to be errors on the test set from the error characterization classifier. Words in bold show a few of these feature mappings. For example, feature set of Row-1 has higher values for the feature *question-starters*, text of Row-3 contains *Named Entity type: time*, a feature present in highly ranked feature-set of the error characterization classifier as shown in Table 6.

Feature Type	Highly ranked features
Lexical	second_person, question_yesno, question_wh !, ?,thanks, no
NRC	positive, negative, trust, fear, anger,
Entities	Org, Time , Date, Cardinal
Dependency	amod-NN-JJ, nummod-NNS,CD, compound-NN-NN, ROOT-NNP-NNP, advmod-VB-RB compound-NN-NNP, neg-VB-RB, amod-NNS, JJ, ROOT-VBN-VBN

Table 6: A subset of top 100 Features from Random Forest.

## 5 Conclusion and Future Work

We present an error characterization framework, called AEC, which allows the model users and developers to understand when and where a model fails. AEC is trained on human understandable linguistic features with erroneous predictions from the base classifier as training input. We used a cross-domain sentiment analysis task as case study to showcase the effectiveness of AEC in terms of error detection and characterization. Our experiments showed that AEC outperformed uncertainty based sampling in terms of selecting the erroneous samples for continuous model improvements (a strong active learning baseline for selecting the most uncertain samples for continuous model improvements) for the task of predicting errors which can act as most informative samples of the base

classifier. In addition, errors automatically detected by AEC seemed to be more understandable to the model developers. Having these explanations lets the end users make a more informed decision, as well as guide the labeling decisions for next round of training. As our initial results on sentiment dataset look promising for both performance and explainability, we are in the process of extending the framework to run the algorithm iteratively on multiple datasets. While applying the error characterization classifier on the unlabeled datasets, not only we will select the top  $K'$  instances with the highest prediction probability of being correctly predicted and add them back to the original training dataset for retraining purpose, but we will also select top  $K$  instances with the highest prediction probability of being incorrectly predicted. We will assign those instances to human annotators for labels and add them back to the original labeled data as well for the next iteration of training process. We will continuously feed these samples to train the base network, and evaluate the actual performance gains for the base classifier.

## References

- Rob Abbott, Marilyn Walker, Jean E. Fox Tree, Pranav Anand, Robeson Bowmani, and Joseph King. 2011. How can you say such things!?: Recognizing Disagreement in Informal Political Argument. In *Proc. of the ACL Workshop on Language and Social Media*.
- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. Fine-grained analysis of sentence embeddings using auxiliary prediction tasks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Yonatan Belinkov and James Glass. 2019. Analysis methods in neural language processing: A survey. *Transactions of the Association for Computational Linguistics*, 7:49–72.
- Gavin C. Cawley. 2011. Baseline methods for active learning. In *Active Learning and Experimental Design workshop, In conjunction with AISTATS 2010, Sardinia, Italy, May 16, 2010*, volume 16 of *JMLR Proceedings*, pages 47–57. JMLR.org.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Finale Doshi-Velez and Been Kim. 2017. Towards a rigorous science of interpretable machine learning. *arXiv: Machine Learning*.
- Boris Galitsky. 2018. Customers’ retention requires an explainability feature in machine learning systems they use. In *2018 AAAI Spring Symposia, Stanford University, Palo Alto, California, USA, March 26-28, 2018*. AAAI Press.
- Bhavya Ghai, Q. Vera Liao, Yunfeng Zhang, Rachel K. E. Bellamy, and Klaus Mueller. 2020. Explainable active learning (XAL): toward AI explanations as interfaces for machine teachers. *Proc. ACM Hum. Comput. Interact.*, 4(CSCW3):1–28.
- Mario Giulianelli, Jack Harding, Florian Mohnert, Dieuwke Hupkes, and Willem H. Zuidema. 2018. Under the hood: Using diagnostic classifiers to investigate and improve how language models track agreement information. In *Proceedings of the Workshop: Analyzing and Interpreting Neural Networks for NLP, BlackboxNLP EMNLP 2018, Brussels, Belgium, November 1, 2018*, pages 240–248. Association for Computational Linguistics.
- M. Joshi and C. Penstein-Rosé. 2009. Generalizing dependency features for opinion mining. In *Proc. of the ACL-IJCNLP 2009 Conference Short Papers*, pages 313–316.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1746–1751.
- David D. Lewis. 1995. A sequential algorithm for training text classifiers: Corrigendum and additional data. *SIGIR Forum*, 29(2):13–19.
- Hui Liu, Qingyu Yin, and William Yang Wang. 2019. Towards explainable NLP: A generative explanation framework for text classification. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5570–5581. Association for Computational Linguistics.
- Amita Misra, Brian Ecker, and Marilyn A. Walker. 2016. Measuring the similarity of sentential arguments in dialogue. In *Proceedings of the SIGDIAL 2016 Conference, The 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 13-15 September 2016, Los Angeles, CA, USA*, pages 276–287. The Association for Computer Linguistics.

- Saif M Mohammad and Peter D Turney. 2010. Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon. In *Proceedings of the NAACL HLT 2010 workshop on computational approaches to analysis and generation of emotion in text*, pages 26–34. Association for Computational Linguistics.
- Grégoire Montavon, Wojciech Samek, and Klaus-Robert Müller. 2018. Methods for interpreting and understanding deep neural networks. *Digit. Signal Process.*, 73:1–15.
- Besmira Nushi, Ece Kamar, and Eric Horvitz. 2018. Towards accountable AI: hybrid human-machine analyses for characterizing system failure. In *Proceedings of the Sixth AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2018, Zürich, Switzerland, July 5-8, 2018*, pages 126–135. AAAI Press.
- Shereen Oraby, Pritam Gundecha, Jalal Mahmud, Mansurul Bhuiyan, and Rama Akkiraju. 2017. "how may I help you?": Modeling twitter customer service conversations using fine-grained dialogue acts. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces, IUI 2017, Limassol, Cyprus, March 13-16, 2017*, pages 343–355. ACM.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Jingyu Shao, Qing Wang, and Fangbing Liu. 2019. Learning to sample: An active learning framework. In *2019 IEEE International Conference on Data Mining, ICDM 2019, Beijing, China, November 8-11, 2019*, pages 538–547. IEEE.
- Xing Shi, Inkit Padhi, and Kevin Knight. 2016. Does string-based neural MT learn source syntax? In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proc. of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642.