

Weakly Supervised Attention Networks for Fine-Grained Opinion Mining and Public Health

Giannis Karamanolakis, Daniel Hsu, Luis Gravano

Columbia University, New York, NY 10027, USA

{gkaraman, djhsu, gravano}@cs.columbia.edu

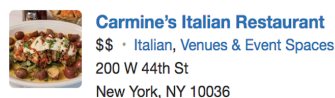
Abstract

In many review classification applications, a fine-grained analysis of the reviews is desirable, because different segments (e.g., sentences) of a review may focus on different aspects of the entity in question. However, training supervised models for segment-level classification requires segment labels, which may be more difficult or expensive to obtain than review labels. In this paper, we employ Multiple Instance Learning (MIL) and use only weak supervision in the form of a single label per review. First, we show that when inappropriate MIL aggregation functions are used, then MIL-based networks are outperformed by simpler baselines. Second, we propose a new aggregation function based on the sigmoid attention mechanism and show that our proposed model outperforms the state-of-the-art models for segment-level sentiment classification (by up to 9.8% in F1). Finally, we highlight the importance of fine-grained predictions in an important public-health application: finding actionable reports of foodborne illness. We show that our model achieves 48.6% higher recall compared to previous models, thus increasing the chance of identifying previously unknown foodborne outbreaks.

1 Introduction

Many applications of text review classification, such as sentiment analysis, can benefit from a fine-grained understanding of the reviews. Consider the Yelp restaurant review in Figure 1. Some segments (e.g., sentences or clauses) of the review express positive sentiment towards some of the items consumed, service, and ambience, but other segments express a negative sentiment towards the price and food. To capture the nuances expressed in such reviews, analyzing the reviews at the segment level is desirable.

In this paper, we focus on segment classification when only review labels—but not segment



4/11/2017

Waited at the bar to be seated. Drink was very nice. Very strong delicious drink. People were all friendly. Our server Papa was amazing. Unfortunately I have been up half the night and suffering all day due to food poisoning. I'm assuming it was the shrimp. Its been a waterfall out of both ends and for the price I would expect better quality. Thus even making me late for school drop off and pick up today. My "medium rare" steak was too tough, more like medium well and the shrimp also was slightly over cooked. Both to the point I had to spit them out. Manager did take 50% off the steak. Great atmosphere. Just wish my bf and I weren't suffering.

Figure 1: A Yelp review discussing both positive and negative aspects of a restaurant, as well as food poisoning.

labels—are available. The lack of segment labels prevents the use of standard supervised learning approaches. While review labels, such as user-provided ratings, are often available, they are not directly relevant for segment classification, thus presenting a challenge for supervised learning.

Existing weakly supervised learning frameworks have been proposed for training models such as support vector machines (Andrews et al., 2003; Yessenalina et al., 2010; Gärtner et al., 2002), logistic regression (Kotzias et al., 2015), and hidden conditional random fields (Täckström and McDonald, 2011). The most recent state-of-the-art approaches employ the Multiple Instance Learning (MIL) framework (Section 2.2) in hierarchical neural networks (Pappas and Popescu-Belis, 2014; Kotzias et al., 2015; Angelidis and Lapata, 2018; Pappas and Popescu-Belis, 2017; Ilse et al., 2018). MIL-based hierarchical networks combine the (unknown) segment labels through an aggregation function to form a single review label. This enables the use of ground-truth review labels as a weak form of supervision for training segment-level classifiers. However, it remains unanswered whether performance gains in current models stem from the hierarchical struc-

ture of the models or from the representational power of their deep learning components. Also, as we will see, the current modeling choices for the MIL aggregation function might be problematic for some applications and, in turn, might hurt the performance of the resulting classifiers.

As a first contribution of this paper, we show that non-hierarchical, deep learning approaches for segment-level sentiment classification—with only review-level labels—are strong, and they equal or exceed in performance hierarchical networks with various MIL aggregation functions.

As a second contribution of this paper, we substantially improve previous hierarchical approaches for segment-level sentiment classification and propose the use of a new MIL aggregation function based on the sigmoid attention mechanism to jointly model the relative importance of each segment as a product of Bernoulli distributions. This modeling choice allows multiple segments to contribute with different weights to the review label, which is desirable in many applications, including segment-level sentiment classification. We demonstrate that our MIL approach outperforms all of the alternative techniques.

As a third contribution, we experiment beyond sentiment classification and apply our approach to a critical public health application: the discovery of foodborne illness incidents in online restaurant reviews. Restaurant patrons increasingly turn to social media—rather than official public health channels—to discuss food poisoning incidents (see Figure 1). As a result, public health authorities need to identify such rare incidents among the vast volumes of content on social media platforms. We experimentally show that our MIL-based network effectively detects segments discussing food poisoning and has a higher chance than all previous models to identify unknown foodborne outbreaks.

2 Background and Problem Definition

We now summarize relevant work on fully supervised (Section 2.1) and weakly supervised models (Section 2.2) for segment classification. We also describe a public health application for our model evaluation (Section 2.3). Finally, we define our problem of focus (Section 2.4).

2.1 Fully Supervised Models

State-of-the-art supervised learning methods for segment classification use segment embedding techniques followed by a classification model. During segment encoding, a segment $s_i = (x_{i1}, x_{i2}, \dots, x_{iN_i})$ composed of N_i words is encoded as a fixed-size real vector $h_i \in \mathbb{R}^\ell$ using transformations such as the average of word embeddings (Wieting et al., 2015; Arora et al., 2017), Recurrent Neural Networks (RNNs) (Wieting and Gimpel, 2017; Yang et al., 2016), Convolutional Neural Networks (CNNs) (Kim, 2014), or self-attention blocks (Devlin et al., 2019; Radford et al., 2018). We refer to the whole segment encoding procedure as $h_i = \text{ENC}(s_i)$. During segment classification, the segment s_i is assigned to one of C predefined classes $[C] := \{1, 2, \dots, C\}$. To provide a probability distribution $p_i = \langle p_i^1, \dots, p_i^C \rangle$ over the C classes, the segment encoding h_i is fed to a classification model: $p_i = \text{CLF}(h_i)$. Supervised approaches require ground-truth *segment labels* for training.

2.2 Weakly Supervised Models

State-of-the-art weakly supervised approaches for segment and review classification employ the Multiple Instance Learning (MIL) framework (Zhou et al., 2009; Pappas and Popescu-Belis, 2014; Kotzias et al., 2015; Pappas and Popescu-Belis, 2017; Angelidis and Lapata, 2018). In contrast to traditional supervised learning, where *segment labels* are required to train segment classifiers, MIL-based models can be trained using *review labels* as a weak source of supervision, as we describe next.

MIL is employed for problems where data are arranged in groups (bags) of instances. In our setting, each review is a group of segments: $r = (s_1, s_2, \dots, s_M)$. The key assumption followed by MIL is that the observed review label is an aggregation function of the unobserved segment labels: $p = \text{AGG}(p_1, \dots, p_M)$. Hierarchical MIL-based models (Figure 2) work in three main steps: (1) encode the review segments into fixed-size vectors $h_i = \text{ENC}(s_i)$, (2) provide segment predictions $p_i = \text{CLF}(h_i)$, and (3) aggregate the predictions to get a review-level probability estimate $p = \text{AGG}(p_1, \dots, p_M)$. Supervision during training is provided in the form of review labels.

Different modeling choices have been taken for each part of the MIL hierarchical architecture. Kotzias et al. (2015) encoded sentences as the in-

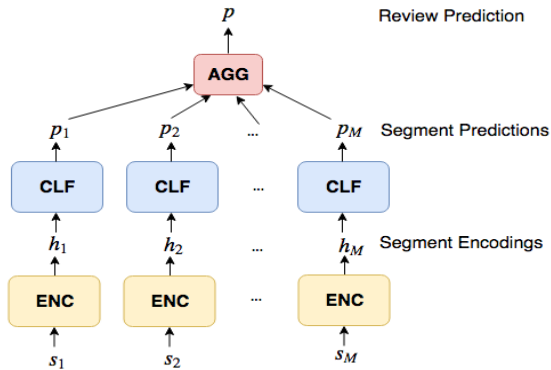


Figure 2: MIL-based hierarchical models.

ternal representations of a hierarchical CNN that was pre-trained for document-level sentiment classification (Denil et al., 2014). For sentence-level classification, they used logistic regression, while the aggregation function was the uniform average. Pappas and Popescu-Belis (2014, 2017) employed Multiple Instance Regression, evaluated various models for segment encoding, including feed forward neural networks and Gated Recurrent Units (GRUs) (Bahdanau et al., 2015), and used the weighted average for the aggregation function, where the weights were computed by linear regression or a one-layer neural network. Angelidis and Lapata (2018) proposed an end-to-end Multiple Instance Learning Network (MILNET), which outperformed previous models for sentiment classification using CNNs for segment encoding, a softmax layer for segment classification, and GRUs with attention (Bahdanau et al., 2015) to aggregate segment predictions as a weighted average. Our proposed model (Section 4) also follows the MIL hierarchical structure of Figure 2 for both sentiment classification and an important public health application, which we discuss next.

2.3 Foodborne Illness Discovery in Online Restaurant Reviews

Health departments nationwide have started to analyze social media content (e.g., Yelp reviews, Twitter messages) to identify foodborne illness outbreaks originating in restaurants. In Chicago (Harris et al., 2014), New York City (Efland et al., 2018), Nevada (Sadilek et al., 2016), and St. Louis (Harris et al., 2018), text classification systems have been successfully deployed for the detection of social media documents mentioning foodborne illness. (Figure 1 shows a Yelp re-

view discussing a food poisoning incident.) After such social media documents are flagged by the classifiers, they are typically examined manually by epidemiologists, who decide if further investigation (e.g., interviewing the restaurant patrons who became ill, inspecting the restaurant) is warranted. This manual examination is time-consuming, and hence it is critically important to (1) produce accurate review-level classifiers, to identify foodborne illness cases while not showing epidemiologists large numbers of false-positive cases; and (2) annotate the flagged reviews to help the epidemiologists in their decision-making.

We propose to apply our segment classification approach to this important public health application. By identifying which review segments discuss food poisoning, epidemiologists could focus on the relevant portions of the review and safely ignore the rest. As we will see, our evaluation will focus on Yelp restaurant reviews. Discovering foodborne illness is fundamentally different from sentiment classification, because the mentions of food poisoning incidents in Yelp are rare. Furthermore, even reviews mentioning foodborne illness often include multiple sentences unrelated to foodborne illness (see Figure 1).

2.4 Problem Definition

Consider a text review for an entity, with M contiguous segments $r = (s_1, \dots, s_M)$. Segments may have a variable number of words and different reviews may have a different number of segments. A discrete label $y_r \in [C]$ is provided for each review but the individual segment labels are not provided. Our goal is to train a segment-level classifier that, given an unseen test review $r^t = (s_1^t, s_2^t, \dots, s_{M_t}^t)$, predicts a label $p_i \in [C]$ for each segment and then aggregates the segment labels to infer the review label $y_r^t \in [C]$ for r^t .

3 Non-Hierarchical Baselines

We can address the problem described in Section 2.4 without using hierarchical approaches such as MIL. In fact, the hierarchical structure of Figure 2 for the MIL-based deep networks adds a level of complexity that has not been empirically justified, giving rise to the following question: do performance gains in current MIL-based models stem from their hierarchical structure or just from the representational power of their deep learning components?

We explore this question by evaluating a class of simpler non-hierarchical baselines: deep neural networks trained at the *review level* (without encoding and classifying individual segments) and applied at the *segment level* by treating each test segment as if it were a short “review.” While the distribution of input length is different during training and testing, we will show that this class of non-hierarchical models is quite competitive and sometime outperforms MIL-based networks with inappropriate modeling choices.

4 Hierarchical Sigmoid Attention Networks

We now describe the details of our MIL-based hierarchical approach, which we call Hierarchical Sigmoid Attention Network (HSAN). HSAN works in three steps to process a review, following the general architecture in Figure 2: (1) each segment s_i in the review is encoded as a fixed-size vector using word embeddings and CNNs (Kim, 2014): $h_i = \text{CNN}(s_i) \in \mathbb{R}^\ell$; (2) each segment encoding h_i is classified using a softmax classifier with parameters $W \in \mathbb{R}^\ell$ and $b \in \mathbb{R}$: $p_i = \text{softmax}(Wh_i + b)$; and (3) a review prediction p is computed as an aggregation function of the segment predictions p_1, \dots, p_M from the previous step. A key contribution of our work is the motivation, definition, and evaluation of a suitable aggregation function for HSAN, a critical design issue for MIL-based models.

The choice of aggregation function has a substantial impact on the performance of MIL-based models and should depend on the specific assumptions about the relationship between bags and instances (Carbonneau et al., 2018). Importantly, the performance of MIL algorithms depends on the witness rate (WR), which is defined as the proportion of positive instances in positive bags. For example, when WR is very low (which is the case in our public health application), using the uniform average as an aggregation function in MIL is not an appropriate modeling choice, because the contribution of the few positive instances to the bag label is outweighed by that of the negative instances.

The choice of the uniform average of segment predictions (Kotzias et al., 2015) is also problematic because particular segments of reviews might be more informative than other segments for the task at hand and thus should contribute with higher weights to the computation of the review label.

For this reason, we opt for the weighted average (Pappas and Popescu-Belis, 2014; Angelidis and Lapata, 2018):

$$p = \frac{\sum_{i=1}^M \alpha_i \cdot p_i}{\sum_{i=1}^M \alpha_i}. \quad (1)$$

The weights $\alpha_1, \dots, \alpha_M \in [0, 1]$ define the relative contribution of the corresponding segments s_1, \dots, s_M to the review label. To estimate the segment weights, we adopt the attention mechanism (Bahdanau et al., 2015). In contrast to MIL-NET (Angelidis and Lapata, 2018), which uses the traditional softmax attention, we propose to use the sigmoid attention. Sigmoid attention is both functionally and semantically different from softmax attention and is more suitable for our problem, as we show next.

The probabilistic interpretation of softmax attention is that of a categorical latent variable $z \in \{1, \dots, M\}$ that represents the index of the segment to be selected from the M segments (Kim et al., 2017). The attention probability distribution is:

$$p(z = i | e_1, \dots, e_M) = \frac{\exp(e_i)}{\sum_{i=1}^M \exp(e_i)}, \quad (2)$$

where:

$$e_i = u_a^T \tanh(W_a h'_i + b_a), \quad (3)$$

where h'_i are context-dependent segment vectors computed using bi-directional GRUs (Bi-GRUs), $W_a \in \mathbb{R}^{m \times n}$ and $b_a \in \mathbb{R}^n$ are the attention model’s weight and bias parameter, respectively, and $u_a \in \mathbb{R}^m$ is the “attention query” vector parameter. The probabilistic interpretation of Equation 2 suggests that, when using the softmax attention, exactly one segment should be considered important under the constraint that the weights of all segments sum to one. This property of the softmax attention to prioritize one instance explains the successful application of the mechanism for problems such as machine translation (Bahdanau et al., 2015), where the role of attention is to align each target word to (usually) one of the M words from the source language. However, softmax attention is not well suited for estimating the aggregation function weights for our problem, where multiple segments usually affect the review-level prediction.

We hence propose using the sigmoid attention mechanism to compute the weights $\alpha_1, \dots, \alpha_M$.

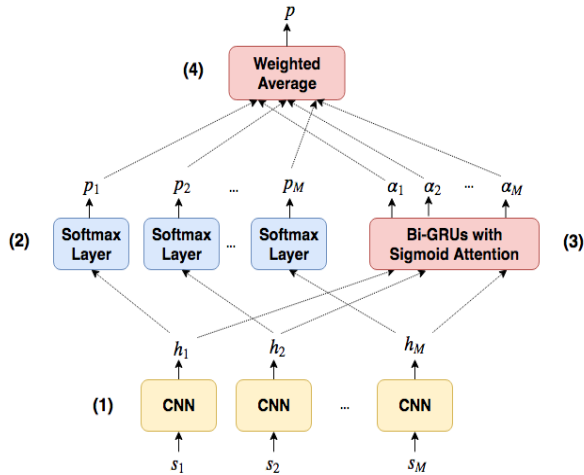


Figure 3: Our Hierarchical Sigmoid Attention Network.

In particular, we replace softmax in Equation (2) with the sigmoid (logistic) function:

$$\alpha_i = \sigma(e_i) = \frac{1}{1 + \exp(-e_i)}. \quad (4)$$

With sigmoid attention, the computation of the attention weight α_i does not depend on scores e_j for $j \neq i$. Indeed, the probabilistic interpretation of sigmoid attention is a vector z of discrete latent variables $z = [z_1, \dots, z_M]$, where $z_i \in \{0, 1\}$ (Kim et al., 2017). In other words, the relative importance of each segment is modeled as a Bernoulli distribution. The sigmoid attention probability distribution is:

$$p(z_i = 1 \mid e_1, \dots, e_M) = \sigma(e_i). \quad (5)$$

This probabilistic model indicates that z_1, \dots, z_M are conditionally independent given e_1, \dots, e_M . Therefore, sigmoid attention allows multiple segments, or even no segments, to be selected. This property of sigmoid attention explains why it is more appropriate for our problem. Also, as we will see in the next sections, using the sigmoid attention is the key modeling change needed in MIL-based hierarchical networks to outperform non-hierarchical baselines for segment-level classification. Attention mechanisms using sigmoid activation have also been recently applied for tasks different than segment-level classification of reviews (Shen and Lee, 2016; Kim et al., 2017; Rei and Søgaard, 2018). Our work differs from these approaches in that we use the sigmoid attention

mechanism for the MIL aggregation function of Equation 1, i.e., we aggregate segment labels p_i (instead of segment vectors h_i) into a single review label p (instead of review vectors h).

We summarize our HSN architecture in Figure 3. HSN follows the MIL framework and thus it does not require segment labels for training. Instead, we only use ground-truth review labels and jointly learn the model parameters by minimizing the negative log-likelihood of the model parameters. Even though a single label is available for each review, our model allows different segments of the review to receive different labels. Thus, we can appropriately handle reviews such as that in Figure 1 and assign a mix of positive and negative segment labels, even when the review as a whole has a negative (2-star) rating.

5 Experiments

We now turn to another key contribution of our paper, namely, the evaluation of critical aspects of hierarchical approaches and also our HSN approach. For this, we focus on two important and fundamentally different, real-world applications: segment-level sentiment classification and the discovery of foodborne illness in restaurant reviews.

5.1 Experimental Settings

For segment-level sentiment classification, we use the Yelp’13 corpus with 5-star ratings (Tang et al., 2015) and the IMDB corpora with 10-star ratings (Diao et al., 2014). We do not use segment labels for training any models except the fully supervised Seg-* baselines (see below). For evaluating the segment-level classification performance on Yelp’13 and IMDB, we use the SPOT-Yelp and SPOT-IMDB datasets, respectively (Angelidis and Lapata, 2018), annotated at two levels of granularity, namely, sentences (SENT) and Elementary Discourse Units (EDUs)¹ (see Table 1). For dataset statistics and implementation details, see the supplementary material.

For the discovery of foodborne illness, we use a dataset of Yelp restaurant reviews, manually labeled by epidemiologists in the New York City Department of Health and Mental Hygiene. Each review is assigned a binary label (“Sick” vs. “Not Sick”). To test the models at the sentence level, epidemiologists have manually annotated each

¹The use of EDUs for sentiment classification is motivated in (Angelidis and Lapata, 2018).

Statistic	SPOT-Yelp		SPOT-IMDB	
	SENT	EDU	SENT	EDU
# Segments	1,065	2,110	1,029	2,398
Positive segments (%)	39.9	32.9	37.9	25.6
Neutral segments (%)	21.7	34.3	29.2	47.7
Negative segments (%)	38.4	32.8	32.9	26.7
Witness positive (# segs)	7.9	12.1	6.0	8.5
Witness negative (# segs)	7.3	11.6	6.6	11.2
Witness salient (# segs)	8.5	14.0	7.6	12.6
WR positive	0.74	0.58	0.55	0.36
WR negative	0.68	0.53	0.63	0.43
WR salient	0.80	0.65	0.76	0.55

Table 1: Label statistics for the SPOT datasets. “WR (x)” is the witness rate, meaning the proportion of segments with label x in a review with label x . “Witness (x)” is the average number of segments with label x in a review with label x . “Salient” is the union of the “positive” and “negative” classes.

sentence for a subset of the test reviews (see the supplementary material). In this sentence-level dataset, the WR of the “Sick” class is 0.25, which is significantly lower than the WR on sentiment classification datasets (Table 1). In other words, the proportion of “Sick” segments in “Sick” reviews is relatively low; in contrast, in sentiment classification the proportion of positive (or negative) segments is relatively high in positive (or negative) reviews.

For a robust evaluation of our approach (HSAN), we compare against state-of-the-art models and baselines:

- **Rev-***: non-hierarchical models, trained at the review level and applied at the segment level (see Section 3); this family includes a logistic regression classifier trained on review embeddings, computed as the element-wise average of word embeddings (“Rev-LR-EMB”), a CNN (“Rev-CNN”) (Kim, 2014), and a Bi-GRU with attention (“Rev-RNN”) (Bahdanau et al., 2015). For foodborne classification we also report a logistic regression classifier trained on bag-of-words review vectors (“Rev-LR-BoW”), because it is the best performing model in previous work (Effland et al., 2018).
- **MIL-***: MIL-based hierarchical deep learning models with different aggregation functions. “MIL-avg” computes the review label as the average of the segment-level predictions (Kotzias et al., 2015). “MIL-softmax” uses the softmax attention mechanism –this is

the best performing MILNET model reported in (Angelidis and Lapata, 2018) (“MIL-NETgt”). “MIL-sigmoid” uses the sigmoid attention mechanism as we propose in Section 4 (HSAN model). All MIL-* models have the hierarchical structure of Figure 2 and for comparison reasons we use the same functions for segment encoding (ENC) and segment classification (CLF), namely, a CNN and a softmax classifier, respectively.

For the evaluation of hierarchical non-MIL networks such as the hierarchical classifier of Yang et al. (2016), see Angelidis and Lapata (2018). Here, we ignore this class of models as they have been outperformed by MILNET.

The above models require only review-level labels for training, which is the scenario of focus of this paper. For comparison purposes, we also evaluate a family of fully supervised baselines trained at the *segment* level:

- **Seg-***: fully supervised baselines using SPOT segment labels for training. “Seg-LR” is a logistic regression classifier trained on segment embeddings, which are computed as the element-wise average of the corresponding word embeddings. We also report the CNN baseline (“Seg-CNN”), which was evaluated in Angelidis and Lapata (2018). Seg-* baselines are evaluated using 10-fold cross-validation on the SPOT dataset.

For sentiment classification, we evaluate the models using the macro-averaged F1 score. For foodborne classification, we report both macro-averaged F1 and recall scores (for more metrics, see the supplementary material).

5.2 Experimental Results

Sentiment Classification: Table 2 reports the evaluation results on SPOT datasets for both sentence- and EDU-level classification.

The Seg-* baselines are not directly comparable with other models, as they are trained at the segment level on the (relatively small) SPOT datasets with segment labels. The more complex Seg-CNN model does not significantly improve over the simpler Seg-LR, perhaps due to the small training set available at the segment level.

Rev-CNN outperforms Seg-CNN in three out of the four datasets. Although Rev-CNN is trained at the review level (but is applied at the segment

Method	SPOT-Yelp		SPOT-IMDB	
	SENT	EDU	SENT	EDU
Seg-LR	55.6	59.2	60.5	62.8
Seg-CNN	56.2	60.0	58.3	63.0
Rev-LR-EMB	51.2	49.3	52.7	48.6
Rev-CNN	60.6	61.5	60.8	60.1
Rev-RNN	58.5	53.9	55.3	50.8
MIL-avg	51.8	46.8	45.7	38.4
MIL-softmax	63.4	59.9	64.0	59.9
MIL-sigmoid	64.6	63.3	66.2	65.7

Table 2: F1 score for segment-level sentiment classification.

level), it is trained with 10 times as many examples as Seg-CNN. This suggests that, for the non-hierarchical CNN models, review-level training may be advantageous with more training examples. In addition, Rev-CNN outperforms Rev-LR-EMB, indicating that the fine-tuned features extracted by the CNN are an improvement over the pre-trained embeddings used by Rev-LR-EMB.

Rev-CNN outperforms MIL-avg and has comparable performance to MILNET: non-hierarchical deep learning models trained at the review level and applied at the segment level are strong baselines, because of their representational power. Thus, the Rev-* model class should be evaluated and compared with MIL-based hierarchical models for applications where segment labels are not available.

Interestingly, MIL-sigmoid (HSAN) consistently outperforms all models, including MIL-avg, MIL-softmax (MILNET), and the Rev-* baselines. This shows that:

1. the choice of aggregation function of MIL-based classifiers heavily impacts classification performance; and
2. MIL-based hierarchical networks can indeed outperform non-hierarchical networks when the appropriate aggregation function is used.

We emphasize that we use the same ENC and CLF functions across all MIL-based models to show that performance gains stem solely from the choice of aggregation function. Given that HSAN consistently outperforms MILNET in all datasets for segment-level sentiment classification, we conclude that the choice of sigmoid attention for aggregation is a better fit than softmax for this task.

The difference in performance between HSAN and MILNET is especially pronounced on the *-

EDU datasets. We explain this behavior with the statistics of Table 1: “Witness (Salient)” is higher in *-EDU datasets compared to *-SENT datasets. In other words, *-EDU datasets contain more segments that should be considered important than *-SENT datasets. This implies that the attention model needs to “attend” to more segments in the case of *-EDU datasets: as we argued in Section 4, this is best modeled by sigmoid attention.

Foodborne Illness Discovery: Table 3 reports the evaluation results for both review- and sentence-level foodborne classification.² For more detailed results, see the supplementary material. Rev-LR-EMB has significantly lower F1 score than Rev-CNN and Rev-RNN: representing a review as the uniform average of the word embeddings is not an appropriate modeling choice for this task, where only a few segments in each review are relevant to the positive class.

MIL-sigmoid (HSAN) achieves the highest F1 score among all models for review-level classification. MIL-avg has lower F1 score compared to other models: as discussed in Section 2.2, in applications where the value of WR is very low (here WR=0.25), the uniform average is not an appropriate aggregation function for MIL.

Applying the best classifier reported in Effland et al. (2018) (Rev-LR-BoW) for sentence-level classification leads to high precision but very low recall. On the other hand, the MIL-* models outperform the Rev-* models in F1 score (with the exception of MIL-avg, which has lower F1 score than Rev-RNN): the MIL framework is appropriate for this task, especially when the weighted average is used for the aggregation function. The significant difference in recall and F1 score between different MIL-based models highlights once again the importance of choosing the appropriate aggregation function. MIL-sigmoid consistently outperforms MIL-softmax in all metrics, showing that the sigmoid attention properly encodes the hierarchical structure of reviews. MIL-sigmoid also outperforms all other models in all metrics. Also, MIL-sigmoid’s recall is 48.6% higher than that of Rev-LR-BoW. In other words, MIL-sigmoid detects more sentences relevant to foodborne illness than Rev-LR-BoW, which is especially desirable

²We report review-level classification results because epidemiologists rely on the *review-level* predictions to decide whether to investigate restaurants; in turn, *segment-level* predictions help epidemiologists focus on the relevant portions of positively labeled reviews.

Method	REV	SENT			
	F1	Prec	Rec	F1	AUPR
Rev-LR-BoW	86.7	82.1	58.8	68.6	80.9
Rev-LR-EMB	63.3	50.0	84.3	62.8	48.9
Rev-CNN	84.8	79.3	59.4	67.9	24.7
Rev-RNN	86.7	81.0	74.5	77.6	11.3
MIL-avg	59.8	75.0	78.0	76.5	73.6
MIL-softmax	87.6	75.5	83.3	79.2	81.6
MIL-sigmoid	89.6	76.4	87.4	81.5	84.0

Table 3: Review-level (left) and sentence-level (right) evaluation results for discovering foodborne illness.

for this application, as discussed next.

Important Segment Highlighting Fine-grained predictions could potentially help epidemiologists to quickly focus on the relevant portions of the reviews and safely ignore the rest. Figure 4 shows how the segment predictions and attention scores predicted by HSAN—with the highest recall and F1 score among all models that we evaluated—could be used to highlight important sentences of a review. We highlight sentences in red if the corresponding attention scores exceed a pre-defined threshold. In this example, high attention scores are assigned by HSAN to sentences that mention food poisoning or symptoms related to food poisoning. (For more examples, see the supplementary material.) This is particularly important because reviews on Yelp and other platforms can be long, with many irrelevant sentences surrounding the truly important ones for the task at hand. The fine-grained predictions produced by our model could inform a graphical user interface in health departments for the inspection of candidate reviews. Such an interface would allow epidemiologists to examine reviews more efficiently and, ultimately, more effectively.

6 Conclusions and Future Work

We presented a Multiple Instance Learning-based model for fine-grained text classification that requires only review-level labels for training but produces both review- and segment-level labels. Our first contribution is the observation that non-hierarchical deep networks trained at the review level and applied at the segment level (by treating each test segment as if it were a short “review”) are surprisingly strong and perform comparably or better than MIL-based hierarchical networks with a variety of aggregation functions. Our second contribution is a new MIL aggregation

Pred	Att	Text
✓	0.00	I wish I could give it zero stars. 🤢 : Sick ✓ : Not Sick
✓	0.00	I actually created a yelp account to write this review!
✓	0.00	At first I thought it was great that we got a table for 5 morning of on a Saturday.
✓	0.00	The food was okay- the poached eggs on the Benedict were a little over cooked, but nothing to complain about.
✓	0.00	The service was good, it was overall fine.
✓	0.00	That is- until I got home and me and boy friend spent the rest of the day/night and into the morning hunched over or sitting on the toilet!
🤢	0.18	I have never experienced such violent food poisoning in my life!
✓	0.00	That was the only place we ate or drank anything at that day, so I know it was from this restaurant.
🤢	0.82	By far the most miserable I've been- chills and crippling abdominal pain along with uncontrollable vomiting and something worse out the other end for my boyfriend!
🤢	0.00	Whatever you do, do not eat here, it is not worth the risk of ending up so unwell.
✓	0.00	To clarify what I believe caused this- we both had carrot juice randomly.
🤢	0.00	I know more than one person who has gotten food poisoning recently from carrot juice- especially if its raw or cold pressed.

Figure 4: HSAN’s fine-grained predictions for a Yelp review: for each sentence, HSAN provides one binary label (Pred) and one attention score (Att). A sentence is highlighted if its attention score is greater than 0.1.

function based on the sigmoid attention mechanism, which explicitly allows multiple segments to contribute to the review-level classification decision with different weights. We experimentally showed that the sigmoid attention is the key modeling change needed for MIL-based hierarchical networks to outperform the non-hierarchical baselines for segment-level sentiment classification. Our third contribution is the application of our weakly supervised approach to the important public health application of foodborne illness discovery in online restaurant reviews. We showed that our MIL-based approach has a higher chance than all previous models to identify unknown foodborne outbreaks, and demonstrated how its fine-grained segment annotations can be used to highlight the segments that were considered important for the computation of the review-level label.

In future work, we plan to consider alternative techniques for segment encoding (ENC), such as pre-trained transformer-based language models (Devlin et al., 2019; Radford et al., 2018), which we expect to further boost our method’s performance. We also plan to quantitatively evaluate the extent to which the fine-grained predictions of our model help epidemiologists to efficiently examine candidate reviews and to interpret classification decisions. Indeed, choosing segments of the review text that explain the review-level decisions can help interpretability (Lei et al., 2016; Yessenalina et al., 2010; Biran and Cotton, 2017). Another important direction for future work is to study if minimal supervision at the fine-grain

level, either in the form of expert labels or rationales (Bao et al., 2018), could effectively guide the weakly supervised models. This kind of supervision is especially desirable to satisfy prior beliefs about the intended role of fine-grained predictions in downstream applications. We believe that building this kind of fine-grained models is particularly desirable when model predictions are used by humans to take concrete actions in the real world.

Acknowledgments

We thank the anonymous reviewers for their constructive feedback. This material is based upon work supported by the National Science Foundation under Grant No. IIS-15-63785.

References

- Stuart Andrews, Ioannis Tsochantaridis, and Thomas Hofmann. 2003. Support vector machines for multiple-instance learning. In *Advances in Neural Information Processing Systems*, pages 577–584.
- Stefanos Angelidis and Mirella Lapata. 2018. Multiple instance learning networks for fine-grained sentiment analysis. *Transactions of the Association for Computational Linguistics*, 6:17–31.
- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. A simple but tough-to-beat baseline for sentence embeddings. In *Proceedings of the 5th International Conference on Learning Representations*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the 3rd International Conference on Learning Representations*.
- Yujia Bao, Shiyu Chang, Mo Yu, and Regina Barzilay. 2018. Deriving machine attention from human rationales. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Or Biran and Courtenay Cotton. 2017. Explanation and justification in machine learning: A survey. In *IJCAI-17 Workshop on Explainable AI (XAI)*.
- Marc-André Carbonneau, Veronika Cheplygina, Eric Granger, and Ghyslain Gagnon. 2018. Multiple instance learning: A survey of problem characteristics and applications. *Pattern Recognition*, 77:329–353.
- Misha Denil, Alban Demiraj, and Nando de Freitas. 2014. Extraction of salient sentences from labelled documents. *Technical report, University of Oxford*.
- 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*.
- Qiming Diao, Minghui Qiu, Chao-Yuan Wu, Alexander J Smola, Jing Jiang, and Chong Wang. 2014. Jointly modeling aspects, ratings and sentiments for movie recommendation (JMARS). In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 193–202.
- Thomas Effland, Anna Lawson, Sharon Balter, Katelynn Devinney, Vasudha Reddy, HaeNa Waechter, Luis Gravano, and Daniel Hsu. 2018. Discovering foodborne illness in online restaurant reviews. *Journal of the American Medical Informatics Association*.
- Thomas Gärtner, Peter A Flach, Adam Kowalczyk, and Alexander J Smola. 2002. Multi-instance kernels. In *Proceedings of the International Conference on Machine Learning*, volume 2, pages 179–186.
- Jenine K Harris, Leslie Hinyard, Kate Beatty, Jared B Hawkins, Elaine O Nsoesie, Raed Mansour, and John S Brownstein. 2018. Evaluating the implementation of a Twitter-based foodborne illness reporting tool in the City of St. Louis Department of Health. *International Journal of Environmental Research and Public Health*, 15(5).
- Jenine K Harris, Raed Mansour, Bechara Choucair, Joe Olson, Cory Nissen, and Jay Bhatt. 2014. Health department use of social media to identify foodborne illness—Chicago, Illinois, 2013–2014. *Morbidity and Mortality Weekly Report*, 63(32):681–685.
- Maximilian Ilse, Jakub M Tomczak, and Max Welling. 2018. Attention-based deep multiple instance learning. In *Proceedings of the 36th International Conference on Machine Learning*.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1746–1751.
- Yoon Kim, Carl Denton, Luong Hoang, and Alexander M Rush. 2017. Structured attention networks. In *Proceedings of the 5th International Conference on Learning Representations*.
- Dimitrios Kotzias, Misha Denil, Nando De Freitas, and Pádraic Smyth. 2015. From group to individual labels using deep features. In *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 597–606.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. Rationalizing neural predictions. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 107–117.

- Nikolaos Pappas and Andrei Popescu-Belis. 2014. Explaining the stars: Weighted multiple-instance learning for aspect-based sentiment analysis. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 455–466.
- Nikolaos Pappas and Andrei Popescu-Belis. 2017. Explicit document modeling through weighted multiple-instance learning. *Journal of Artificial Intelligence Research*, 58:591–626.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. <https://blog.openai.com/language-unsupervised>.
- Marek Rei and Anders Søgaard. 2018. Zero-shot sequence labeling: transferring knowledge from sentences to tokens. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 1, pages 293–302.
- Adam Sadilek, Henry A Kautz, Lauren DiPrete, Brian Labus, Eric Portman, Jack Teitel, and Vincent Silenzio. 2016. Deploying nEmesis: Preventing foodborne illness by data mining social media. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3982–3990.
- Sheng-syun Shen and Hung-yi Lee. 2016. Neural attention models for sequence classification: Analysis and application to key term extraction and dialogue act detection. In *Proceedings of INTERSPEECH*, pages 2716–2720.
- Oscar Täckström and Ryan McDonald. 2011. Discovering fine-grained sentiment with latent variable structured prediction models. In *European Conference on Information Retrieval*, pages 368–374.
- Duyu Tang, Bing Qin, and Ting Liu. 2015. Document modeling with gated recurrent neural network for sentiment classification. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1422–1432.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2015. Towards universal paraphrastic sentence embeddings. In *Proceedings of the 3rd International Conference on Learning Representations*.
- John Wieting and Kevin Gimpel. 2017. Revisiting recurrent networks for paraphrastic sentence embeddings. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 2078–2088.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489.
- Ainur Yessenalina, Yisong Yue, and Claire Cardie. 2010. Multi-level structured models for document-level sentiment classification. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1046–1056.
- Zhi-Hua Zhou, Yu-Yin Sun, and Yu-Feng Li. 2009. Multi-instance learning by treating instances as non-IID samples. In *Proceedings of the 26th International Conference on Machine Learning*, pages 1249–1256.