arXiv:2009.09112v2 [cs.CL] 31 May 2021

An Interpretable and Uncertainty Aware Multi-Task Framework for Multi-Aspect Sentiment Analysis

TIAN SHI, Virginia Tech, USA
PING WANG, Virginia Tech, USA
CHANDAN K. REDDY, Virginia Tech, USA

In recent years, several online platforms have seen a rapid increase in the number of review systems that request users to provide aspect-level feedback. Document-level Multi-aspect Sentiment Classification (DMSC), where the goal is to predict the ratings/sentiment from a review at an individual aspect level, has become a challenging and an imminent problem. To tackle this challenge, we propose a deliberate self-attention based deep neural network model, named as FEDAR, for the DMSC problem, which can achieve competitive performance while also being able to interpret the predictions made. As opposed to the previous studies, which make use of hand-crafted keywords to determine aspects in rating predictions, our model does not suffer from human bias issues since aspect keywords are automatically detected through a self-attention mechanism. FEDAR is equipped with a highway word embedding layer to transfer knowledge from pre-trained word embeddings, an RNN encoder layer with output features enriched by pooling and factorization techniques, and a deliberate self-attention layer. In addition, we also propose an Attention-driven Keywords Ranking (AKR) method, which can automatically discover aspect keywords and aspect-level opinion keywords from the review corpus based on the attention weights. These keywords are significant for rating predictions by FEDAR. Since crowdsourcing annotation can be an alternate way to recover missing ratings of reviews, we propose a LEcture-AuDience (LEAD) strategy to estimate model uncertainty in the context of multi-task learning, so that valuable human resources can focus on the most uncertain predictions. Our extensive set of experiments on five different open-domain DMSC datasets demonstrate the superiority of the proposed FEDAR and LEAD models. We further introduce two new DMSC datasets in the healthcare domain and benchmark different baseline models and our models on them. Attention weights visualization results and visualization of aspect and opinion keywords demonstrate the interpretability of our model and effectiveness of our AKR method.

CCS Concepts: • **Information systems** → **Sentiment analysis**; **Clustering and classification**; *Information extraction.*

Additional Key Words and Phrases: Multi-task learning, model uncertainty, deep neural network, dropout, classification, online reviews

ACM Reference Format:
Tian Shi, Ping Wang, and Chandan K. Reddy. 2018. An Interpretable and Uncertainty Aware Multi-Task Framework for Multi-Aspect Sentiment Analysis. In *ACM/IMS Transactions on Data Science*. ACM, New York, NY, USA, 20 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

Sentiment analysis plays an important role in many business applications [24]. It is used to identify customers' opinions and emotions toward a particular product/service via identifying polarity (i.e., positive, neutral or negative) of given textual reviews [15, 25]. In the past few years, with the rapid growth of online reviews, the topic of fine-grained aspect-based sentiment analysis (ABSA)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM/IMS TDS, New York, NY, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

Fig. 1. An example of an online review from BeerAdvocate platform. Keywords corresponding to different aspects are highlighted with different colors.

[28] has attracted significant attention since it allows models to predict opinion polarities with respect to aspect-specific terms in a sentence. Different from sentence-level ABSA, document-level multi-aspect sentiment classification (DMSC) aims to predict the sentiment polarity of documents, which are composed of several sentences, with respect to a given aspect [14, 42, 44]. DMSC has become a significant challenge since many websites provide platforms for users to give aspect-level feedback and ratings, such as TripAdvisor[1] and BeerAdvocate[2]. Fig. 1 shows a review example from the BeerAdvocate website. In this example, a beer is rated with four different aspects, i.e., feel, look, smell and taste. The review also describes the beer with four different aspects. There is an overall rating associated with this review. Recent studies have found that users are less motivated to give aspect-level ratings [42, 44], which makes it difficult to analyze their preference, and it takes a lot of time and effort for human experts to manually annotate them.

There are several recent studies that aim to predict the aspect ratings or opinion polarities using deep neural network based models with multi-task learning framework [14, 42, 44, 45]. In this setting, rating predictions for different aspects, which are highly correlated and can share the same review encoder, are treated as different tasks. However, these models rely on hand-crafted aspect keywords to aid in rating/sentiment predictions [14, 42, 45]. Thus, their results, especially case studies of reviews, are biased towards pre-defined aspect keywords. In addition, these models only focus on improving the prediction accuracy, however, knowledge discovery (such as aspect and opinion related keywords) from review corpus still relies on unsupervised [21] and rule-based methods [44], which limits applications of current DMSC models [14, 42, 45]. In the past few years, model uncertainty of deep neural network classifiers has received increasing attention [6, 8], because it can identify low-confidence regions of input space and give more reliable predictions. Uncertainty models have also been applied to deep neural networks for text classification [46]. However, few existing uncertainty methods have been used to improve the overall prediction accuracy of multi-task learning models when crowd-sourcing annotation is involved in the DMSC task. In this paper, we attempt to tackle the above mentioned issues. The primary contributions of this paper are as follows:

- Develop a FEDAR model that achieves competitive results on five benchmark datasets without using hand-crafted aspect keywords. The proposed model is equipped with a highway word embedding layer, a sequential encoder layer whose output features are enriched by pooling and factorization techniques, and a deliberate self-attention layer. The deliberate self-attention layer

[1]https://www.tripadvisor.com
[2]https://www.beeradvocate.com

can boost performance as well as provide interpretability for our FEDAR model. Here, FEDAR represents of some key components of our model, including Feature Enrichment, Deliberate self-Attention, and overall Rating.

- Introduce two new datasets obtained from the RateMDs website https://www.ratemds.com, which is a platform for patients to review the performance of their doctors. We benchmark different models on them.
- Propose an Attention-driven Keywords Ranking (AKR) method to automatically discover aspect and opinion keywords from review corpus based on attention weights, which also provides a new research direction for interpreting self-attention mechanism. The extracted keywords are significant to ratings/polarities predicted by FEDAR.
- Propose a LEcture-Audience (LEAD) method to measure the uncertainty of our FEDAR model for given reviews. This method can also be generally applied to other deep neural networks.

The rest of this paper is organized as follows: In Section 2, we introduce related work of the DMSC task and uncertainty estimation methods. In Section 3, we present details of our proposed FEDAR model, AKR method and LEAD uncertainty estimation approach. In Section 4, we introduce different DMSC datasets, baseline methods and implementation details, as well as analyze experimental results. Our discussion concludes in Section 5.

2 RELATED WORK

Document-level Multi-aspect Sentiment Classification (DMSC) aims to predict ratings/sentiment of reviews with respect to given aspects. It is originated from online review systems which request users to provide aspect-level ratings for a product or service. Most of the early studies in DMSC solved this problem by first extracting features (e.g., *n*-grams) for each aspect and then predicting aspect-level ratings [18, 21] using regression techniques (e.g., Support Vector Regression [33]). More recently, deep learning models formulate DMSC as a multi-task classification problem [14, 42, 45]. In these models, reviews are first encoded to their corresponding vector representation using recurrent neural networks. Then, aspect-specific attention modules and classifiers are built upon the review encoders to predict the sentiment. For example, Yin *et al.* [42] have formulated this task as a machine comprehension problem. Li *et al.* [14] proposed incorporating users' information, overall ratings, and hand-crafted aspect keywords into their model to predict ratings, instead of merely using textual reviews. Zeng *et al.* [44] introduced a variational approach to weakly supervised sentiment analysis. Aspect-based sentiment classification (ABSA) [28] is another research direction that is related to our work. It consists of several fine-grained sentiment classification tasks, including aspect category detection and polarity, and aspect term extraction and polarity. However, these tasks primarily focus on sentence-level sentiment classification [36, 38], and typically need human experts to annotate aspect terms, categories, and entities. In this paper, we focus on the DMSC problem and our model is also based on a multi-task learning framework. In addition, we place more emphasis on the model interpretability, automatic aspect and opinion keywords discovery, and uncertainty estimation.

Model uncertainty of deep neural networks (NNs) is another research topic related to this work. Bayesian NNs, which learn a distribution over weights, have been studied extensively and achieved competitive results for measuring uncertainty [1, 17, 23]. However, they are often difficult to implement and computationally expensive compared with standard deep NNs. Gal and Ghahramani [7] proposed using Monte Carlo dropout to estimate uncertainty by applying dropout [34] at testing time, which can be interpreted as a Bayesian approximation of the Gaussian process [30]. This method has gain popularity in practice [10, 20] since it is simple to implement and computationally more efficient. Recently, Zhang *et al.* [46] applied dropout-based uncertainty estimation methods

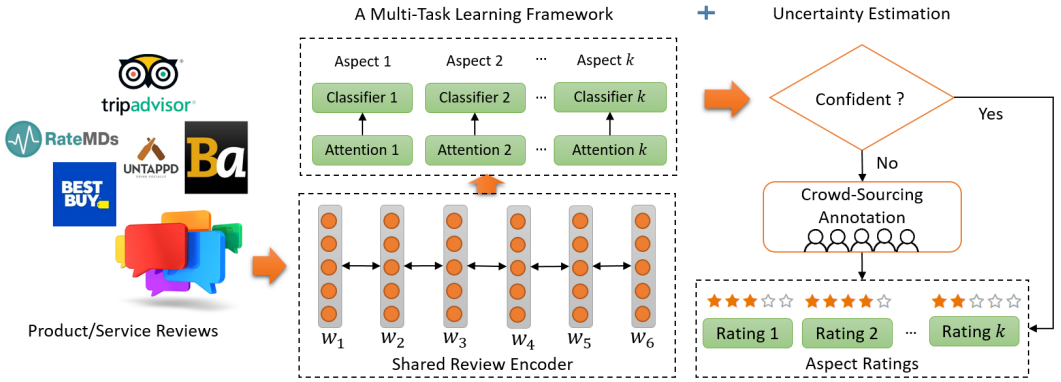

Fig. 2. An overview of our multi-task learning framework with uncertainty estimation for accurate and reliable sentiment classification in DMSC task. Here, sentiment classification for each aspect is treated as a task and different tasks share the same review encoder.

to text classification. Our paper proposes a new method for estimating uncertainty for deep NNs and we use it to measure the uncertainty of our models.

3 PROPOSED METHODS

In this section, we first introduce our FEDAR model (See Fig. 3) for the DMSC task. Then, we describe our AKR method to automatically discover aspect and aspect-level sentiment terms based on the FEDAR model. Finally, we discuss our LEAD method (See Fig. 4) for measuring the uncertainty of the FEDAR model.

3.1 The Proposed FEDAR Model

3.1.1 Problem Formulation. The DMSC problem can be formulated as a multi-task classification problem, where the sentiment classification for each aspect is viewed as a task (See Fig. 2). More formally, the DMSC problem is described as follows: Given a textual review $X = (x_1, x_2, ..., x_T)$, our goal is to predict class labels, i.e., integer ratings/sentiment polarity of the review $y = (y^1, y^2, ..., y^K)$, where T and K are the number of tokens in the review and the number of aspects/tasks, respectively. x_t and y^k are the one-hot vector representations of word t and the class label of aspect k, respectively.

The challenge in this problem is to build a model that can achieve competitive accuracy without losing model interpretability or obtaining biased results. Therefore, we propose improving word embedding, review encoder and self-attention layers to accomplish this goal. We will now introduce our model and provide more details of our architecture in a layer-by-layer manner.

3.1.2 Highway Word Embedding Layer. This layer aims to learn word vectors based on pre-trained word embeddings. We first use word embedding technique [22] to map one-hot representations of tokens $x_1, x_2, ..., x_T$ to a continuous vector space, thus, they are represented as $E_{x_1}, E_{x_2}, ..., E_{x_T}$, where E_{x_t} is the word vector of x_t, pre-trained on a large corpus and fixed during parameter inference. In our experiments, we adopted GloVe word vectors [27], so that they do not need to be trained from random states, which may result in poor embeddings due to the lack of word co-occurrence.

Then, a single layer highway network [35] is used to adapt the knowledge, i.e., semantic information from pre-trained word embeddings, to target DMSC datasets. Formally, the highway network is defined as follows:

$$E'_{x_t} = f(E_{x_t}) \odot g(E_{x_t}) + E_{x_t} \odot (1 - g(E_{x_t})) \quad (1)$$

where $f(\cdot)$ and $g(\cdot)$ are affine transformations with ReLU and Sigmoid activation functions, respectively. $\odot$ represents element-wise product. $g(\cdot)$ is also known as gate, which is used to control the information that is being carried to the next layer. Intuitively, the highway network aims at transferring knowledge from pre-trained word embeddings to the target review corpus. E'_{x_t} can be viewed as a perturbation of E_{x_t}, and $f(\cdot)$ and $g(\cdot)$ have significantly fewer parameters than E_{x_t}. Therefore, training a highway network is more efficient than training a word embedding layer from random parameters.

3.1.3 Review Encoder Layer. This layer describes the review encoder and feature enrichment techniques proposed in our model.

Sequential Encoder Layer: The output of highway word embedding layer $(E'_{x_1}, E'_{x_2}, ..., E'_{x_T})$ is fed into a sequential encoder layer. Here, we adopt a multi-layer bi-directional LSTM encoder [9], which encodes a review into a sequence of hidden states in forward direction $\overrightarrow{H} = (\overrightarrow{h_1}, \overrightarrow{h_2}, ..., \overrightarrow{h_T})$ and backward direction $\overleftarrow{H} = (\overleftarrow{h_1}, \overleftarrow{h_2}, ..., \overleftarrow{h_T})$.

Representative Features: For each hidden state $\overrightarrow{h_t}$ (or $\overleftarrow{h_t}$), we generate three representative features, which will be later used to assist the attention mechanism to learn the overall review representation.

The first and second features, denoted by $\overrightarrow{h_t^{\max}}$ and $\overrightarrow{h_t^{\text{avg}}}$, are the max-pooling and average-pooling of $\overrightarrow{h_t}$, respectively. The third one is obtained using factorization machine [31], where the factorization operation is defined as

$$\mathcal{F}(z) = w_0 + \sum_{i=1}^{N} w_i z_i + \sum_{i=1}^{N} \sum_{j=i+1}^{N} \left\langle V_i, V_j \right\rangle z_i z_j. \tag{2}$$

Here, the model parameters are $w_i \in \mathbb{R}$ and $V \in \mathbb{R}^{N \times F}$. N and F are the dimensions of the input vector z and factorization, respectively. $\langle \cdot, \cdot \rangle$ is the dot product between two vectors. w_0 in Eq. (2) is a global bias, w_i is the strength of the i-th variable, and $\left\langle V_i, V_j \right\rangle$ captures the pairwise interaction between z_i and z_j.

Intuitively, the max-pooling and avg-pooling provide the approximated location (bound and mean) of the hidden state $\overrightarrow{h_t}$ in the N dimensional space, while the factorization captures all single and pairwise interactions. Together they provide the high-level knowledge of that hidden state.

Feature Augmentation: Finally, the aggregated hidden state h_t at time step t is obtained by concatenating hidden states in both directions and all representative features, i.e.,

$$\begin{aligned} \overrightarrow{h_t} &= \overrightarrow{h_t} \oplus \overrightarrow{h_t^{\max}} \oplus \overrightarrow{h_t^{\text{avg}}} \oplus \mathcal{F}(\overrightarrow{h_t}), \\ \overleftarrow{h_t} &= \overleftarrow{h_t} \oplus \overleftarrow{h_t^{\max}} \oplus \overleftarrow{h_t^{\text{avg}}} \oplus \mathcal{F}(\overleftarrow{h_t}), \\ h_t &= \overrightarrow{h_t} \oplus \overleftarrow{h_t}. \end{aligned} \tag{3}$$

Thus, the review is encoded into a sequence of aggregated hidden states $H = (h_1, h_2, \ldots, h_T)$.

3.1.4 Deliberate Self-Attention Layer. Once the aggregated hidden states for each review are obtained, we apply a self-attention layer for each task to learn an overall review representation for that task. Compared with pooling and convolution operations, self-attention mechanism is more interpretable, since it can capture relatively important words for a given task. However, a standard self-attention layer merely relies on a single global alignment vector across different reviews, which results in sub-optimal representations. Therefore, we propose a deliberate self-attention alignment

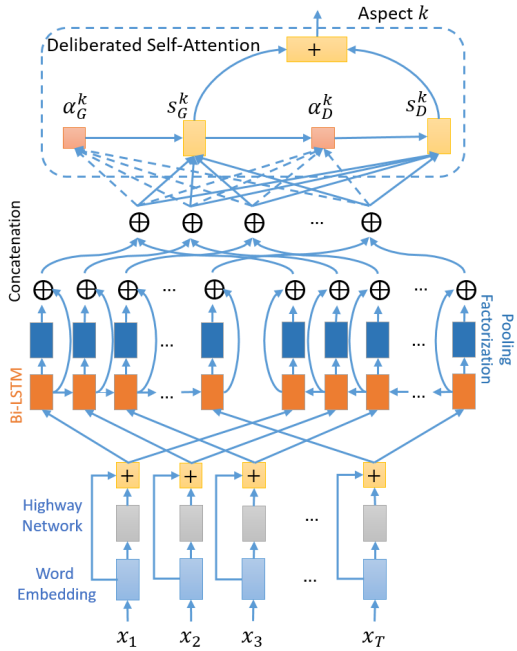

Fig. 3. Review encoder and deliberate self-attention for aspect k. Each hidden state is enriched by three features, i.e., max-pooling, average-pooling and factorization.

method to refine the review representations while maintaining the network interpretability. In this section, we will first introduce the self-attention mechanism, and then provide the details of the deliberation counterpart.

Global Self-Attention: For each aspect k, the self-attention mechanism [41] is used to learn the relative importance of tokens in a review to the sentiment classification task. Formally, given the aggregated hidden states H for a review, the alignment score $u_{t,G}^k$ and attention weight $\alpha_{t,G}^k$ are calculated as follows:

$$u_{t,G}^k = (v_G^k)^\top \tanh(W_G^k h_t + b_G^k), \ \alpha_{t,G}^k = \frac{\exp(u_{t,G}^k)}{\sum_{\tau=1}^T \exp(u_{\tau,G}^k)}, \tag{4}$$

where W_G^k, v_G^k and b_G^k are model parameters. G represents global, as the above attention mechanism is also known as global attention [19]. v_G^k is viewed as a global aspect-specific base-vector in this paper, since it has been used in calculating the alignment with different hidden states across different reviews. It can also be viewed as a global aspect-specific filter that is designed to capture important information for a certain aspect from different reviews. Therefore, we also refer a regular self-attention layer as a global self-attention layer. With attention weights, the global review representation is calculated by taking the weighted sum of all aggregated hidden states, i.e., $s_G^k = \sum_{t=1}^T \alpha_{t,G}^k h_t$. Traditionally, s_G^k is used for the sentiment classification task.

Deliberate Attention: As we can see from Eq. (4), the importance of a token t is measured by the similarity between $\tanh(W_G^k h_t + b_G^k)$ and the base-vector v_G^k. However, a single base-vector v_G^k is difficult to capture the variability in the reviews, and hence, such alignment results in sub-optimal

representations of reviews. In this paper, we attempt to alleviate this problem by reusing the output of the global self-attention, i.e., s_G^k, as a document-level aspect-specific base-vector to produce better review representations. Notably, s_G^k already incorporates the knowledge of the review content and aspect k. We refer this step as deliberation.

Given the hidden states H and review representation s_G^k, we first calculate the alignment scores and attention weights as follows:

$$u_{t,D}^k = (s_G^k)^\top \tanh(W_D^k h_t + b_D^k),\ \alpha_{t,D}^k = \frac{\exp(u_{t,D}^k)}{\sum_{\tau=1}^T \exp(u_{\tau,D}^k)}, \tag{5}$$

where W_D^k and b_D^k are parameters. D represents deliberation. Similarly, we can calculate the aspect-specific review representation by deliberation as $s_D^k = \sum_{t=1}^T \alpha_{t,D}^k h_t$.

Review Representation: Finally, the review representation for aspect k can be obtained as follows[3]:

$$s^k = s_G^k + s_D^k = \sum_{t=1}^T (\alpha_{t,G}^k + \alpha_{t,D}^k) h_t. \tag{6}$$

From the above equation, we not only get refined review representations but also maintain the interpretability of our model. Here, we did not use the concatenation of two vectors since we would like to maintain the interpretability as well. Notably, we can use the accumulated attention weights, i.e., $\frac{1}{2}(\alpha_{t,G}^k + \alpha_{t,D}^k)$, to interpret our experimental results.

3.1.5 Sentiment Classification Layer. Finally, we pass the representation of each review for aspect k into an aspect-specific classifier to get the probability distribution over different class labels. Here, the classifier is defined as a two layer feed-forward network with a ReLU activation followed by a softmax layer, i.e.,

$$\begin{aligned} y_{\text{out}}^k &= \text{ReLU}(W_{\text{out}}^k s^k + b_{\text{out}}^k), \\ y_{\text{pred}}^k &= \text{softmax}(W_{\text{pred}}^k y_{\text{out}}^k + b_{\text{pred}}^k), \end{aligned} \tag{7}$$

where W_{out}^k, W_{pred}^k, b_{out}^k, and b_{pred}^k are learnable parameters.

Given the ground-truth labels $\hat{y}^k$, which is a one-hot vector, our goal is to minimize the averaged cross-entropy error between y_{pred}^k and $\hat{y}^k$ across all aspects, i.e.,

$$\mathcal{L}_\theta = -\sum_{k=1}^K \sum_{i=1}^N \hat{y}_i^k \log(y_{\text{pred},i}^k), \tag{8}$$

where K and N represents the number of aspects and class labels, respectively. The model is trained in an end-to-end manner using back-propagation.

3.2 Aspect and Sentiment Keywords

Traditionally, aspect and sentiment keywords are obtained using unsupervised clustering methods, such as topic models [21, 32]. However, these methods cannot automatically build correlations between keywords and aspects or sentiment due to the lack of supervision. Aspect and opinion term extractions in fine-grained aspect-based sentiment analysis tasks [5, 28, 29, 37] focus on extracting terms and phrases from sentences. However, they require a number of labeled reviews to train deep learning models. In this paper, we propose a fully automatic Attention-driven Keywords Ranking (AKR) method to discover aspect and opinion keywords, which are important to predicted ratings,

[3]In this paper, we also consider models that repeat the deliberation for multiple times. However, we did not observe significant performance improvement.

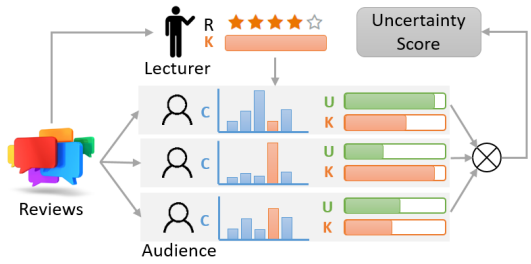

Fig. 4. The LEcture-AuDience (LEAD) model for uncertainty estimation. 'R', 'K', 'C', 'U' represent rating, knowledge, probability distribution of different class labels (see Eq. (7)), and uncertainty score, respectively.

from a review corpus based on self-attention (or deliberate self-attention) mechanism in the context of DMSC.

3.2.1 Aspect Keywords Ranking. The significance of a word w to an aspect k can be described by a conditional probability $p_C(w|k)$ on a review corpus C. Intuitively, given an aspect k, if a word w_1 is more frequent than w_2 across the corpus, then, w_1 is more significant to aspect k. We can further expand this probability as follows:

$$p_C(w|k) = \sum_{\xi \in C} p_C(w, \xi|k), \tag{9}$$

where ξ is a review in corpus C. For each $\xi \in C$, probability $p_C(w, \xi|k)$ indicates the importance of word w to the aspect k, which can be defined using attention weights, i.e.,

$$p_C(w, \xi|k) = \frac{\sum_{t=1}^{T} \alpha_t^{\xi} \cdot \delta(w_t, w)}{\sum_{\xi' \in C} f_{\xi'}(w) + \gamma}, \tag{10}$$

where $f_{\xi'}(w)$ is frequency of w in document ξ' and γ is a smooth factor. $\delta(w_t, w) = \begin{cases} 1 & \text{if } w_t = w \\ 0 & \text{otherwise} \end{cases}$ is a delta function. Attention weight α_t^{ξ} is defined as $\alpha_t^{\xi} = \frac{1}{2}(\alpha_{t,G}^{k} + \alpha_{t,D}^{k})$ for the deliberation self-attention mechanism. In Eq. (10), the denominator is applied to reduce the noise from stop-words and punctuation. After obtaining the score $p_C(w|k)$ for every member in the vocabulary, we collect top-ranked words (with part-of-speech tags: NOUN and PROPN) as aspect keywords.

3.2.2 Aspect-level Opinion Keywords. Similarly, we can estimate the significance of a word w to an aspect-level opinion label/rating $\hat{y}^k$ by a conditional probability $p_C(w|\hat{y}^k)$. Let us use $C_{\hat{y}^k}$ to denote reviews with rating $\hat{y}^k$ for aspect k, then, the following equivalence holds, i.e.,

$$p_C(w|\hat{y}^k) = p_{C_{\hat{y}^k}}(w|k), \tag{11}$$

which can be further calculated by Eqs. (9) and (10). Intuitively, we first construct a subset $C_{\hat{y}^k} \subset C$ of the review corpus, then, we use attention weights of aspect k to calculate the significance of word w to that aspect. Finally, we collect top-ranked words (with part-of-speech tags: ADJ, ADV and VERB) as aspect-level opinion keywords.

3.3 The Proposed Uncertainty Model

Although our FEDAR model has achieved competitive prediction accuracy and our AKR method allows us to explore aspect and sentiment keywords, it is still difficult to deploy such a model in

real-world applications. In DMSC datasets, we find that there are many typos and abbreviations in reviews and many reviews describe the product or service from only one aspect. However, deep learning models cannot capture these problems in the datasets, therefore, the predictions are not reliable. One way to tackle this challenge is by estimating the uncertainty of model predictions. If a model returns ratings with high uncertainty, we can pass the review to human experts for annotation. In this section, we propose a LEcture-AuDience (LEAD) method (See Fig.4) to measure the uncertainty of our FEDAR model in the context of multi-task learning.

3.3.1 Lecturer and Audiences. We use a lecturer (denoted by $\mathcal{M}^L$) to represent any well-trained deep learning model, e.g., FEDAR model. Audiences are models (denoted by $\mathcal{M}^A$) with partial knowledge of the lecturer, where *knowledge can be interpreted as relationships between an input review and output ratings* which are inferred by $\mathcal{M}^L$. Here, $\mathcal{M}^A = \{\mathcal{M}^{A_1}, \mathcal{M}^{A_2}, ..., \mathcal{M}^{A_{|A|}}\}$, where $|A|$ is the number of audiences. *Partial knowledge determines the eligibility of audiences to provide uncertainty scores.* For example, eligible audiences can be: (1) Models obtained by pruning some edges (e.g., dropout with small dropout rate) of the lecturer model. (2) Models obtained by continuing training the lecturer model with very small learning rate for a few batches. Ineligible audiences include: (1) Random models trained on the same or a different review corpus. (2) Models with the same or similar structure as lecturer but initialized with different parameters and trained on a different corpus.

3.3.2 Uncertainty Scores. Given a review, suppose the lecturer $\mathcal{M}^L$ predicts the class label as $\tilde{y}^{L,k}$ for aspect k, where $\tilde{y}^{L,k}$ is an one-hot vector. An audience $\mathcal{M}^{A_\mu}$ obtains the probability distribution over different class labels as $y_{\text{pred}}^{A_\mu,k}$ (See Eq. (7)). Then, the uncertainty score is defined as the cross entropy between $\tilde{y}^{L,k}$ and $y_{\text{pred}}^{A_\mu,k}$, which is calculated by

$$\psi^{A_\mu,k} = -\sum_{i=1}^{N} \tilde{y}_i^{L,k} \log(y_{\text{pred},i}^{A_\mu,k}). \tag{12}$$

Intuitively, the audience is more uncertain about the lecturer's prediction if it gets lower probability for that prediction. For example, in Fig. 4, the lecturer model predicts rating/label as 4. Three audiences obtain probability 0.1, 0.8, 0.5 for that label, respectively. Then, their uncertainty scores are $\psi^{A_1,k} = 2.30$, $\psi^{A_2,k} = 0.22$, and $\psi^{A_3,k} = 0.69$.

With the uncertainty score from a single audience and for a single aspect, we can calculate the final uncertainty score as

$$\psi = \exp \sum_{\mu=1}^{|A|} \zeta \log \left(\exp \sum_{k=1}^{k} \log \left(\psi^{A_\mu,k} + \lambda \right) + \eta \right), \tag{13}$$

where λ and $\eta \geq 1$ are smoothing factor and set to 1 in our experiments. ζ is an empirical factor for knowledge. If audience networks are obtained by applying dropout to the lecturer network, the higher the dropout rate, the lower the factor ζ. In this case, the audiences have less knowledge to the lecturer.

After obtaining uncertainty scores for all reviews in the testing set, we can select either certain percent of reviews with higher scores or reviews with scores over a threshold for crowdsourcing annotation. Human experts are expected to analyze the reviews and decide the aspect ratings for them.

Table 1. Statistics of different DMSC datasets. † indicates the datasets collected and prepared by us.

Dataset	# docs	# aspects	Scale
TripAdvisor-R	29,391	7	1-5
TripAdvisor-RU	58,632	7	1-5
TripAdvisor-B	28,543	7	1-2
BeerAdvocate-R	50,000	4	1-10
BeerAdvocate-B	27,583	4	1-2
RateMDs-R†	155,995	4	1-5
RateMDs-B†	120,303	4	1-2

4 EXPERIMENTAL RESULTS

In this section, we present the results from an extensive set of experiments and demonstrate the effectiveness of our proposed FEDAR model, AKR and LEAD methods.

4.1 Research Questions

Our empirical analysis aims at the following Research Questions (RQs):

- **RQ1**: What is the overall performance of FEDAR? Does it outperform state-of-the-art baselines?
- **RQ2**: What is the overall performance of LEAD method compared with uncertainty estimation baselines?
- **RQ3**: How does each component in FEDAR contribute to the overall performance?
- **RQ4**: Is the deliberate self-attention module interpretable? Does it learn meaningful aspect and opinion terms from a review corpus?

4.2 Datasets

We first conduct our experiments on five benchmark datasets, which are obtained from TripAdvisor and BeerAdvocate review platforms. TripAdvisor based datasets have seven aspects (*value, room, location, cleanliness, check in/front desk, service, and business service*), while BeerAdvocate based datasets have four aspects (*feel, look, smell, and taste*). TripAdvisor-R [42], TripAdvisor-U [14] and BeerAdvocate-R [13, 42] use the original rating scores as sentiment class labels. In TripAdvisor-B and BeerAdvocate-B [44], the original scale is converted to a binary scale, where 1 and 2 correspond to negative and positive sentiment, respectively. Neutral has been ignored in both datasets. All datasets have been tokenized and split into train/development/test sets with a proportion of 8:1:1. In our experiments, we use the same datasets that are provided by the previous studies in the literature [14, 42, 44]. Statistics of the datasets are summarized in Table 1.

In addition to the aforementioned five datasets, we also propose two new datasets, i.e., RateMDs-R and RateMDs-B, and benchmarked our models on them. RateMDs dataset was collected from https://www.ratemds.com website which has textual reviews along with numeric ratings for medical experts primarily in the North America region. Each review comes with ratings of four different aspects, i.e., *staff, punctuality, helpfulness, and knowledge*. The overall rating is the average of these aspect ratings. To obtain a more refined dataset for our experiments, we removed reviews with missing aspect ratings and selected the rest of the reviews whose lengths are between 72 and 250 tokens (outliers [4]), since short reviews may not have information on all the four aspects. The original data has a rating-imbalance problem, i.e., 60% and 17% of reviews are rated as 5 and 1, respectively, and more than 50% of reviews have identical aspect ratings. Therefore, similar to [13],

[4]The average number of tokens for all reviews is 72 tokens and there are very few reviews with more than 250 tokens.

we chose reviews with different aspect ratings, i.e., at least three of aspect ratings are different. The statistics of our dataset is shown in Table 1. For RateMDs-R, we tokenized reviews with Stanford corenlp[5] and randomly split the dataset into training, development and testing by a proportion of 135,995/10,000/10,000. For RateMDs-B, we followed the process in [44] by converting original scales to binary and sampling data according to the overall polarities to avoid the imbalance issue. The statistics of the RateMDs-B dataset is also shown in Table 1. Similarly, we split the dataset into training, development and testing by a proportion of 100,303/10,000/10,000.

4.3 Comparison Methods

To demonstrate the effectiveness of our methods, we compare the proposed models with following baseline methods:

- **MAJOR** simply uses the majority sentiment labels or polarities in training data as predictions.
- **GLVL** first calculates the document representation by averaging the word vectors of all keywords in a review, where pre-trained word vectors are obtained from GloVe [27]. Then, a LIBLINEAR package [4] is used for the classification task.
- **BOWL** feeds normalized Bag-of-Words (BOW) representation of reviews into the LIBLINEAR package for the sentiment classification. In our experiments, stop-words and punctuation are removed in order to enable the model to capture the keywords more efficiently.
- **MCNN** is an extension of the CNN model in the multi-task learning framework. For each task, CNN [11] extracts key features from a review by applying convolution and max-over-time pooling [2] operations over the shared word embeddings layer.
- **MLSTM** extends a multi-layer Bi-LSTM model [9], which captures both forward and backward semantic information, with the multi-task learning framework, where different tasks have their own classifiers and share the same Bi-LSTM encoder.
- **MBERT** is a multi-task version of the BERT classification model [3]. Different tasks share the same BERT encoder [39].
- **MATTN** is a multi-task version of self-attention based models. Similar to MLSTM, different tasks share the same Bi-LSTM encoder. For each task, we first apply a self-attention layer, and then pass the document representations to a sentiment classifier.
- **DMSCMC** [42] introduces a hierarchical iterative attention model to build aspect-specific document representations by frequent and repeated interactions between documents and aspect questions.
- **HRAN** [14] incorporates hand-crafted aspect keywords and the overall rating into a hierarchical network to build sentence and document representations.
- **AMN** [45] first uses attention-based memory networks to incorporate hand-crafted aspect keywords information into the aspect and sentence memories. Then, recurrent attention operation and multi-hop attention memory networks are employed to build document representations.
- **FEDAR** We name our model as FEDAR, where FE, DA and R represent Feature Enrichment, Deliberate self-Attention, and overall Rating, respectively.

We compare our LEAD method with the following uncertainty estimation approaches:

- **Max-Margin** is the maximal activation of the sentiment classification layer (after softmax normalization).

[5]https://stanfordnlp.github.io/CoreNLP/

Table 2. Averaged Accuracy (ACC) and MSE of different models on TripAdvisor-R (Trip-R), TripAdvisor-U (Trip-U), TripAdvisor-B (Trip-B), BeerAdvocate-R (Beer-R), and BeerAdvocate-B (Beer-B) testing sets. For MSE, smaller is better. † indicates that results are obtained from previous published papers and NA indicates that results are not available in those papers. we use bold font to highlight the best performance values and underline to highlight the second best values

Method	**Trip-R**		**Trip-U**		**Trip-B**	**Beer-R**		**Beer-B**
	ACC	**MSE**	**ACC**	**MSE**	**ACC**	**ACC**	**MSE**	**ACC**
MAJOR	29.12	2.115	39.73	1.222	62.42	26.29	4.252	67.26
GLVL	38.94	1.795	48.04	0.879	78.15	30.59	2.774	79.73
BOWL	40.14	1.708	48.68	0.888	78.38	31.02	2.715	79.14
MCNN	41.75	1.458	51.21	0.714	81.31	34.11	2.016	82.37
MLSTM	42.74	1.401	48.64	0.791	80.56	34.48	2.167	82.07
MATTN	42.13	1.427	50.53	0.679	80.82	35.78	1.962	84.86
MBERT	44.41	1.250	54.50	0.617	82.84	35.94	1.963	84.73
DMSCMC†	46.56	1.083	55.49	0.583	83.34	38.06	1.755	86.35
HRAN†	47.43	1.169	58.15	0.528	NA	39.11	1.700	NA
AMN†	48.66	1.109	NA	NA	NA	40.19	1.686	NA
FEDAR (Ours)	**48.92**	**1.072**	**58.50**	**0.522**	**85.50**	**40.62**	**1.530**	**87.40**

- **PL-Variance** (Penultimate Layer Variance) [43] uses the variance of the output of the sentiment classification layer (before softmax normalization) as the uncertainty score.
- **Dropout** [7] apply dropout to deep neural networks during training and testing. The dropout can be used as an approximation of Bayesian inference in deep Gaussian processes, which aims to identify low-confidence regions of input space.

All methods are based on our FEDAR model.

4.4 Implementation Details

We implemented all deep learning models using Pytorch [26] and the best set of parameters are selected based on the development set. Word embeddings are pre-loaded with 300-dimensional GloVe embeddings [27] and fixed during training. For MCNN, filter sizes are chosen to be 3, 4, 5 and the number of filters are 400 for each size. For all LSTM based models, the dimension of hidden states is set to 600 and the number of layers is 4. All parameters are trained using ADAM optimizer [12] with an initial learning rate of 0.0005. The learning rate decays by 0.8 every 2 epochs. Dropout with a dropout-rate 0.2 is applied to the classifiers. Gradient clipping with a threshold of 2 is also applied to prevent gradient explosion. For MBERT, we leveraged the pre-trained BERT encoder from HuggingFace's Transformers package [39] and fixed its weights during training. We also adopted the learning rate warmup heuristic [16] and set the warmup step to 2000. For dropout-based uncertainty estimation methods, we set the dropout-rate to 0.5. The number of samples for Dropout are 50. The number of audiences is 20 for our LEAD model. ζ is set to 1.0. Our codes and datasets are available at https://github.com/tshi04/DMSC_FEDA.

4.5 Prediction Performance

For research question **RQ1**, we use accuracy (ACC) and mean squared error (MSE) as our evaluation metrics to measure the prediction performance of different models. All results are shown in Tables 2 and 3, where we use bold font to highlight the best performance values and underline to highlight the second best values.

Table 3. Averaged accuracy (ACC) and MSE of different models on RateMDs-R (RMD-R) and RateMDs-B (RMD-B) testing sets. For MSE, smaller is better.

Method	RMD-R		RMD-B
	ACC	MSE	ACC
MAJOR	31.42	3.393	57.18
GLVL	43.11	1.882	76.93
BOWL	44.78	1.704	78.68
MCNN	46.19	1.333	81.60
MLSTM	48.37	1.148	82.40
MATTN	49.08	1.157	82.66
MBERT	48.65	1.160	83.39
FEDAR (Ours)	**55.57**	**0.794**	**88.63**

For the DMSC problem, it has been demonstrated that deep neural network (DNN) based models perform much better than conventional machine learning methods that rely on n-gram or embedding features [14, 42]. In our experiments, we have also demonstrated this by comparing different DNN models with MAJOR, GLVL and BOWL. Compared to simple DNN classification models, multi-task learning DNN models (MDNN) can achieve better results with fewer parameters and training time [42]. Therefore, we focused on comparing the performance of our model with different MDNN models. From Table 2, DMSCMC achieves better results on all five datasets compared with baselines MCNN, MLSTM, MBERT, and MATTN. HRAN and AMN leverage the power of overall rating and get significantly better results than other compared methods. From both tables, we observed our FEDAR model achieves the best performance on all seven datasets. These results demonstrate the effectiveness of our methods.

4.6 Uncertainty Performance

Uncertainty estimation can help users identify reviews for which the models are not confident of their predictions. More intuitively, prediction models are prone to mistakes on the reviews that they are uncertain about. In Table 4, we first selected the most uncertain predictions (denoted by **top-n%**) based uncertainty scores from the testing sets of TripAdvisor-R, BeerAdvocate-R and RateMDs-R datasets. Then, we evaluated the uncertainty performance by comparing the mis-classification rate (i.e., error rate) of our FEDAR model for the selected reviews. The more incorrect predictions that can be captured, the better the uncertainty method will be. From these results, we can observe that Dropout method achieves significantly better results than Max-Margin and PL-Variance. Our LEAD method outperforms all these baseline methods on three datasets, which shows our method is superior in identifying less confident predictions and answers research question **RQ2**.

4.7 Ablation Study of FEDAR

For research question **RQ3**, we attribute the performance improvement of our FEDAR model to: 1) Better review encoder, including a highway word embedding layer and a feature enriched encoder. 2) Deliberate self-attention mechanism. 3) Overall rating. Therefore, we systematically conducted ablation studies to demonstrate the effectiveness of these components, and provided the results in Table 5, Table 6 and Fig. 5.

We first observe that FEDAR significantly outperforms model-OR (FEDAR w/o OR), which indicates that overall rating can help the model make better predictions. Secondly, we compare model-OR with model-ORFE (FEDAR w/o OR, FE), which is equipped with a regular word embedding

Table 4. Performance of various uncertainty methods on different datasets.

TripAdvisor-R					
Method	**top-5%**	**top-10%**	**top-15%**	**top-20%**	**top-25%**
Max-Margin	35.40	36.00	37.47	39.15	40.68
PL-Variance	40.20	42.40	43.00	44.25	44.84
Dropout	53.80	53.50	53.33	53.35	53.60
LEAD	**65.40**	**62.60**	**60.93**	**60.85**	**60.20**
BeerAdvocate-R					
Method	**top-5%**	**top-10%**	**top-15%**	**top-20%**	**top-25%**
Max-Margin	38.80	43.80	46.53	48.30	49.44
PL-Variance	44.00	47.00	48.33	49.65	50.68
Dropout	57.00	57.90	58.60	58.70	59.28
LEAD	**71.60**	**69.50**	**67.93**	**67.55**	**67.20**
RateMDs-R					
Method	**top-5%**	**top-10%**	**top-15%**	**top-20%**	**top-25%**
Max-Margin	20.20	23.80	26.80	28.15	29.40
PL-Variance	28.60	29.70	30.53	30.85	31.60
Dropout	51.00	50.70	50.60	49.60	48.88
LEAD	**66.00**	**62.70**	**60.40**	**59.05**	**58.32**

Table 5. Ablation study results. Different models are evaluated by Averaged Accuracy (ACC) and MSE metrics on five public DMSC testing sets. For MSE, smaller is better. **FE**, **DA** and **OR** represent Feature Enrichment, Deliberated self-Attention, Overall Rating, respectively.

Method	**Trip-R**		**Trip-U**		**Trip-B**	**Beer-R**		**Beer-B**
	ACC	**MSE**	**ACC**	**MSE**	**ACC**	**ACC**	**MSE**	**ACC**
FEDAR	**48.92**	**1.072**	**58.50**	**0.522**	**85.50**	**40.62**	**1.530**	**87.40**
w/o OR	46.72	1.178	55.82	0.574	84.23	39.66	1.617	86.52
w/o OR, DA	45.70	1.224	55.39	0.584	83.43	38.85	1.633	85.99
w/o OR, FE	44.50	1.300	53.41	0.632	82.39	38.92	1.714	84.99
w/o OR, DA, FE	42.13	1.427	50.53	0.679	80.82	35.78	1.962	84.86

Table 6. Ablation study results. Different models are evaluated by Averaged Accuracy (ACC) and MSE metrics on RateMDs-R (RMD-R) and RateMDs-B (RMD-B) testing sets.

Method	**RMD-R**		**RMD-B**
	ACC	**MSE**	**ACC**
FEDAR	**55.82**	**0.786**	**88.63**
w/o OR	49.80	1.106	83.89
w/o OR, DA	49.68	1.108	83.62
w/o OR, FE	49.28	1.123	83.47
w/o OR, DA, FE	49.08	1.157	82.66

layer and a multi-layer Bi-LSTM encoder. Obviously, model-OR obtained better results than model-ORFE. Similarly, we also compare model-ORDA (FEDAR w/o OR, DA) with model-BASE (FEDAR w/o OR, DA, FE), since model-ORDA adopts the same self-attention mechanism as model-BASE.

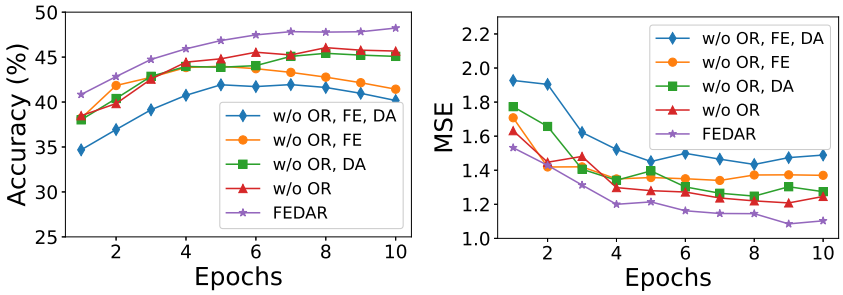

(a) Accuracy vs. Epochs (b) MSE vs. Epochs

Fig. 5. This figure shows (a) Averaged Accuracy and (b) MSE for FEDAR and its variants on the TripAdvisor-R dataset during the training process.

It can be observed that model-ORDA performs significantly better than model-BASE on all the datasets. This experiment shows that we can improve the performance by using highway word embedding layer and feature enrichment technique. Furthermore, we compared model-OR with model-ORDA, which does not have a deliberate self-attention layer. It can been seen that model-OR outperforms model-ORDA in all the experiments. In addition, we have also compared the results of model-ORFE and model-BASE, which are equipped with a deliberate self-attention layer and a regular self-attention layer. We observed that model-ORFE has a better performance compared to model-BASE. This experiment indicates the effectiveness of deliberate self-attention mechanism. In Fig. 5, we show the accuracy and MSE of different models during training in order to demonstrate that FEDAR can get consistently higher accuracy and lower MSE after training for several epochs than its basic variants.

4.8 Attention Visualization

The attention mechanism enables a model to selectively focus on important parts of the reviews, and hence, visualization of the attention weights can help in interpreting our model and analyze the experimental results [40, 42]. To answer research question **RQ4**, we need to investigate whether our model attends to relevant keywords when it is making aspect-specific rating predictions for the DMSC problem.

In Fig. 6 (a), we show a review example from the BeerAdvocate-R testing set, for which our model has successfully predicted all aspect-specific ratings. In this Figure, we highlighted the review with deliberate attention weights. The review contains keywords of all four aspects, thus, we only need to verify whether our model can successfully detect those aspect-specific keywords. We observed that deliberate self-attention attend to "*creamy and luscious mouthfeel*" for **feel**. For the **look** aspect, it captures "*dark murky brown with a ..., leave some lacing on the glass*", which is quite relevant to the appearance of the beer. Our model also successfully detects "*very rich and spicy*" for **smell**. For **taste**, it attends to "*taste is a bit disappointing, ... too prominent*", which yields a slightly lower rating. Similarly, we show an example from the RateMDs-R testing set in Fig .6 (b). Our model detects "*unfortunately, the office staff is very lousy! I do think ...*" for **staff**, which expresses negative opinion to the office staff. For **punctuality**, it captures "*true that you have to wait a long time for her*", which is also negative. Finally, it attends to "*is by far the best doctor, she does get a lot of patient and may get overwhelmed. but when it comes to knowledge, communicating, the best*" for the **knowledge**, and "*she is patient and caring, patience and caring attitude*" for the **helpfulness** of the doctor. Both

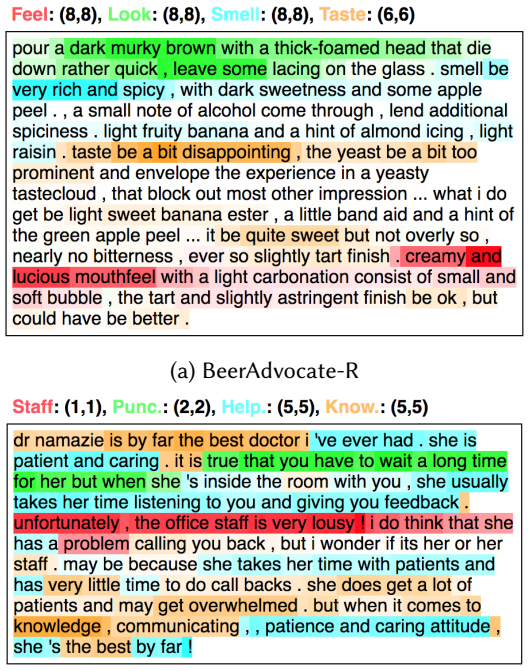

(a) BeerAdvocate-R

(b) RateMDs-R

Fig. 6. Visualization of attention weights. In parentheses, the first and second numbers represent ground-truth and predicted ratings, respectively. Different aspects are labeled with different colors. The figure is best viewed in color.

aspects have positive sentiment. Therefore, these two examples show good interpretability of our model.

4.9 Aspect and Opinion Keywords

In Fig. 7, we first show aspect keywords detected by our AKR method for TripAdivsor-B, BeerAdvocate-B, and RateMDs-B corpus. From Fig. 7 (Top row), we observe that **value** related keywords include "*price, money, rate, overprice*". Keywords related to a **room** are "*air conditioning, comfy, leak, mattress, bathroom, modern, ceiling*" and others. For **cleanliness**, people are interested in "*housekeeping, spotless, cleaning, hair, stain, smell*" and so on. **Service** is related with "*staff, service, employee, receptionist, personnel*". From Fig. 7 (Middle row), we observe that **feel** is usually related with keywords, like "*mouthfeel, mouth, smooth, watery*", which describe feel of beers in mouth. **Look** is the appearance of beers, thus, the model captures "*appearance, retention, white, head, foam, color*" and others. **Smell** related aspect keywords include "*smell, aroma, scent, fruity*" and more. Finally, representative keywords for **taste** are "*taste, balance, complex, flavor*" and so on. From Fig. 7 (Bottom row), we observe that **staff** related keywords are "*staff, assistant, secretary, receptionist*" and so on. For **punctuality**, people usually concern "*waits, hour, hours, retard*". The **helpfulness** of a doctor is related to "*compassion, manner, empathy, attitude, condescending*" and so on. Finally, **knowledge** related keywords are "*knowledge, expertise, surgeon, skill*" and others.

We also obtain aspect-specific opinion keywords from Trip-B, Beer-B, and RMD-B datasets, and show them in Fig. 8. From this figure (Top row), we observe that reviewers with positive

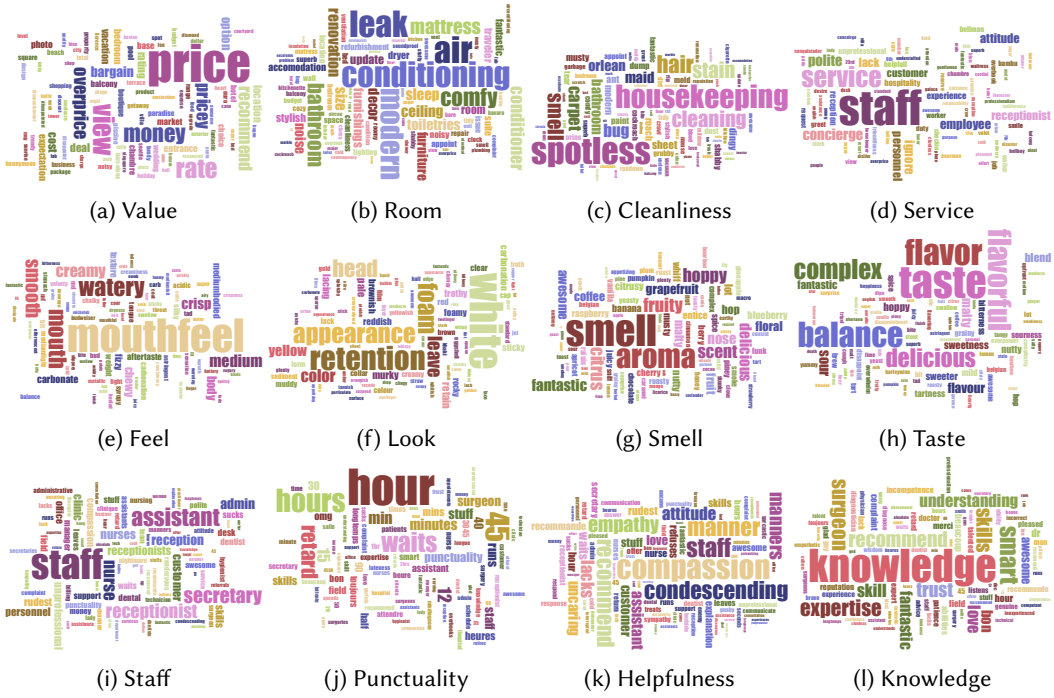

(a) Value (b) Room (c) Cleanliness (d) Service

(e) Feel (f) Look (g) Smell (h) Taste

(i) Staff (j) Punctuality (k) Helpfulness (l) Knowledge

Fig. 7. Word-cloud visualization of aspect keywords for TripAdvisor-B (Top row), BeerAdvocate-B (Middle row) and RateMDs-B datasets (Bottom row).

experience usually live in "*comfortable, beautiful, spacious, lovely and gorgeous*" rooms, and the staff are "*helpful, friendly, courteous and attentive*". While reviewers with negative experience may live in "*uncomfortable, small, cramped and tiny*" rooms. Something may "*leak*" and there are also problems with "*air conditioning*". The staff are "*rude, unhelpful and unfriendly*" and the service is "*poor*". From Fig. 8 (Middle row), we learn that good beers should have "*great, amazing, wonderful, pleasant, aromatic, fresh, rich, and incredible*" smell, and the taste may be "*tasty, great, balanced, enjoyable, and flavorful*". The smell of low-rated beers is "*faint, weak, pungent, odd, funky, and rotten*", and the taste may be "*bland, unbalanced, disappointed, and sour*". From Fig. 8 (Bottom row), we find that good doctors usually have "*sincerely, friendly, helpful, and wonderful*" staff and are "*knowledgeable, competent, intelligent, and excellent*". In a low-rated clinic, staff may be "*incompetent, rude, horrible, terrible, and unfriendly*", and doctors may "*misdiagnose*" conditions of patients and can be not "*competent, knowledgeable, or trusted*".

From these figures, we can conclude that our deliberate self-attention mechanism is interpretable, and by leveraging our AKR method, it is a powerful knowledge discovery tool for online multi-aspect reviews, which answers research question **RQ4**.

5 CONCLUSION

In this paper, we proposed a multi-task deep learning model, namely FEDAR, for the problem of document-level multi-aspect sentiment classification. Different from previous studies, our model does not require hand-crafted aspect-specific keywords to guide the attention and boost model performance for the task of sentiment classification. Instead, our model relies on (a) a highway word embedding layer to transfer knowledge from pre-trained word vectors on a large corpus, (b)

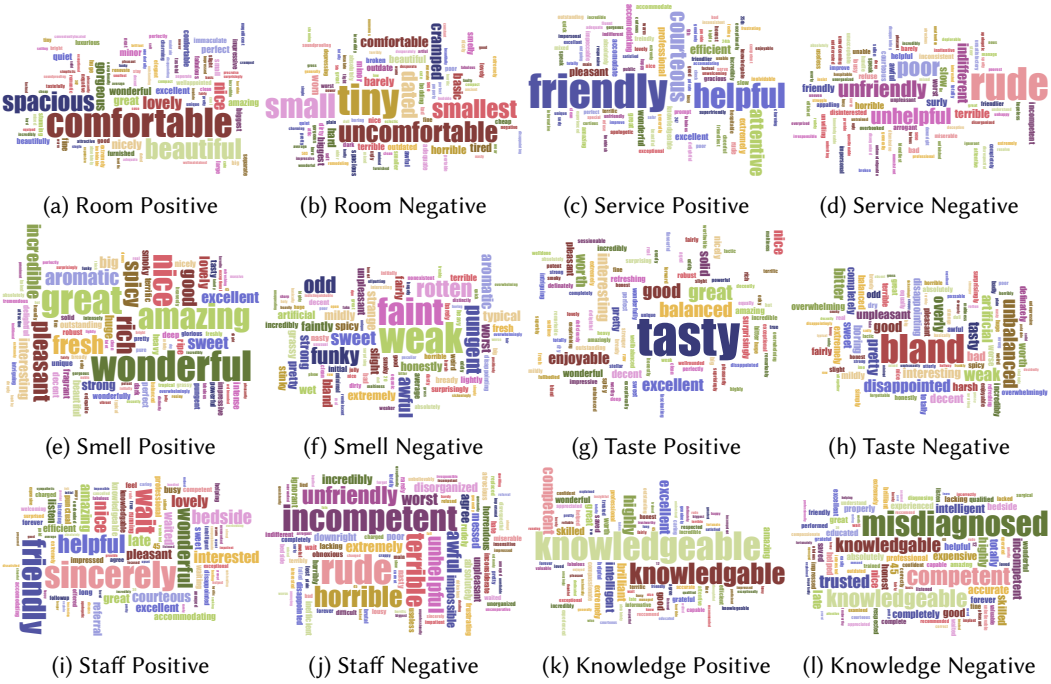

(a) Room Positive (b) Room Negative (c) Service Positive (d) Service Negative

(e) Smell Positive (f) Smell Negative (g) Taste Positive (h) Taste Negative

(i) Staff Positive (j) Staff Negative (k) Knowledge Positive (l) Knowledge Negative

Fig. 8. Word-cloud visualization of aspect-level opinion keywords for TripAdvisor-B (Top row), BeerAdvocate-B (Middle row) and RateMDs-B datasets (Bottom row).

a sequential encoder layer whose output features are enriched by pooling and feature factorization techniques, and (c) a deliberate self-attention layer which maintains the interpretability of our model. Experiments on various DMSC datasets have demonstrated the superior performance of our model. In addition, we also developed an Attention-driven Keywords Ranking (AKR) method, which can automatically discover aspect and opinion keywords from the review corpus based on attention weights. Attention weights visualization and aspect/opinion keywords word-cloud visualization results have demonstrated the interpretability of our model and effectiveness of our AKR method. Finally, we also proposed a LEcture-AuDience (LEAD) method to measure the uncertainty of deep neural networks, including our FEDAR model, in the context of multi-task learning. Our experimental results on multiple real-world datasets demonstrate the effectiveness of the proposed work.

ACKNOWLEDGMENTS

This work was supported in part by the US National Science Foundation grants IIS-1707498, IIS-1838730, and Amazon AWS credits.

REFERENCES

[1] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. 2015. Weight Uncertainty in Neural Network. In *International Conference on Machine Learning*. 1613–1622.

[2] Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research* 12, Aug (2011), 2493–2537.

[3] 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, Volume 1 (Long and Short Papers). 4171–4186.

[4] Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. 2008. LIBLINEAR: A library for large linear classification. *Journal of machine learning research* 9, Aug (2008), 1871–1874.

[5] Zhifang Fan, Zhen Wu, Xinyu Dai, Shujian Huang, and Jiajun Chen. 2019. Target-oriented opinion words extraction with target-fused neural sequence labeling. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. 2509–2518.

[6] Yarin Gal. 2016. Uncertainty in deep learning. *University of Cambridge* 1 (2016), 3.

[7] Yarin Gal and Zoubin Ghahramani. 2015. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. *arXiv preprint arXiv:1506.02142* (2015).

[8] Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*. 1050–1059.

[9] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.

[10] Alex Kendall, Yarin Gal, and Roberto Cipolla. 2018. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7482–7491.

[11] Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1746–1751.

[12] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[13] Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. Rationalizing Neural Predictions. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. 107–117.

[14] Junjie Li, Haitong Yang, and Chengqing Zong. 2018. Document-level Multi-aspect Sentiment Classification by Jointly Modeling Users, Aspects, and Overall Ratings. In *Proceedings of the 27th International Conference on Computational Linguistics*. 925–936.

[15] Bing Liu. 2012. Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies* 5, 1 (2012), 1–167.

[16] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. 2019. On the Variance of the Adaptive Learning Rate and Beyond. In *International Conference on Learning Representations*.

[17] Christos Louizos and Max Welling. 2016. Structured and efficient variational deep learning with matrix gaussian posteriors. In *International Conference on Machine Learning*. 1708–1716.

[18] Bin Lu, Myle Ott, Claire Cardie, and Benjamin K Tsou. 2011. Multi-aspect sentiment analysis with topic models. In *2011 11th IEEE International Conference on Data Mining Workshops*. IEEE, 81–88.

[19] Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. 1412–1421.

[20] Rowan McAllister, Yarin Gal, Alex Kendall, Mark Van Der Wilk, Amar Shah, Roberto Cipolla, and Adrian Weller. 2017. Concrete problems for autonomous vehicle safety: Advantages of bayesian deep learning. International Joint Conferences on Artificial Intelligence, Inc.

[21] Julian McAuley, Jure Leskovec, and Dan Jurafsky. 2012. Learning attitudes and attributes from multi-aspect reviews. In *2012 IEEE 12th International Conference on Data Mining*. IEEE, 1020–1025.

[22] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*. 3111–3119.

[23] Radford M Neal. 2012. *Bayesian learning for neural networks*. Vol. 118. Springer Science & Business Media.

[24] Bo Pang, Lillian Lee, et al. 2008. Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval* 2, 1–2 (2008), 1–135.

[25] Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up?: sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10*. Association for Computational Linguistics, 79–86.

[26] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. In *NIPS-W*.

[27] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.

[28] Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, AL-Smadi Mohammad, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, et al. 2016. SemEval-2016 task 5: Aspect based sentiment analysis. In *SemEval-2016*. 19–30.

[29] Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Harris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. SemEval-2014 Task 4: Aspect Based Sentiment Analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*. 27–35.

[30] Carl Edward Rasmussen. 2003. Gaussian processes in machine learning. In *Summer School on Machine Learning*. Springer, 63–71.
[31] Steffen Rendle. 2010. Factorization machines. In *2010 IEEE International Conference on Data Mining*. IEEE, 995–1000.
[32] Tian Shi, Vineeth Rakesh, Suhang Wang, and Chandan K Reddy. 2019. Document-Level Multi-Aspect Sentiment Classification for Online Reviews of Medical Experts. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 2723–2731.
[33] Alex J Smola and Bernhard Schölkopf. 2004. A tutorial on support vector regression. *Statistics and computing* 14, 3 (2004), 199–222.
[34] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
[35] Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. 2015. Highway networks. *arXiv preprint arXiv:1505.00387* (2015).
[36] Duyu Tang, Bing Qin, and Ting Liu. 2016. Aspect Level Sentiment Classification with Deep Memory Network. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. 214–224.
[37] Wenya Wang, Sinno Jialin Pan, Daniel Dahlmeier, and Xiaokui Xiao. 2017. Coupled multi-layer attentions for co-extraction of aspect and opinion terms. In *Thirty-First AAAI Conference on Artificial Intelligence*.
[38] Yequan Wang, Minlie Huang, Li Zhao, et al. 2016. Attention-based lstm for aspect-level sentiment classification. In *Proceedings of the 2016 conference on empirical methods in natural language processing*. 606–615.
[39] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R'emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace's Transformers: State-of-the-art Natural Language Processing. *ArXiv* abs/1910.03771 (2019).
[40] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*. 2048–2057.
[41] 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*. 1480–1489.
[42] Yichun Yin, Yangqiu Song, and Ming Zhang. 2017. Document-Level Multi-Aspect Sentiment Classification as Machine Comprehension. In *Proceedings of the 2017 Conference on EMNLP*. 2044–2054.
[43] Hugo Zaragoza and Florence d'Alché Buc. 1998. Confidence measures for neural network classifiers. In *Proceedings of the Seventh Int. Conf. Information Processing and Management of Uncertainty in Knowlegde Based Systems*.
[44] Ziqian Zeng, Wenxuan Zhou, Xin Liu, and Yangqiu Song. 2019. A Variational Approach to Weakly Supervised Document-Level Multi-Aspect Sentiment Classification. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. 386–396.
[45] Qingxuan Zhang and Chongyang Shi. 2019. An Attentive Memory Network Integrated with Aspect Dependency for Document-Level Multi-Aspect Sentiment Classification. In *Asian Conference on Machine Learning*. 425–440.
[46] Xuchao Zhang, Fanglan Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2019. Mitigating Uncertainty in Document Classification. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. 3126–3136.