

Context-Guided BERT for Targeted Aspect-Based Sentiment Analysis

Zhengxuan Wu¹, Desmond C. Ong^{2,3}

¹ Symbolic Systems Program, Stanford University

² Department of Information Systems and Analytics, National University of Singapore

³ Institute of High Performance Computing, Agency for Science, Technology, and Research, Singapore
wuzhengx@stanford.edu, dco@comp.nus.edu.sg

Abstract

Aspect-based sentiment analysis (ABSA) and Targeted ABSA (TABSA) allow finer-grained inferences about sentiment to be drawn from the same text, depending on context. For example, a given text can have different targets (e.g., neighborhoods) and different aspects (e.g., price or safety), with different sentiment associated with each target-aspect pair. In this paper, we investigate whether adding context to self-attention models improves performance on (T)ABSA. We propose two variants of Context-Guided BERT (CG-BERT) that learn to distribute attention under different contexts. We first adapt a context-aware Transformer to produce a CG-BERT that uses context-guided *softmax*-attention. Next, we propose an improved Quasi-Attention CG-BERT model that learns a compositional attention that supports subtractive attention. We train both models with pretrained BERT on two (T)ABSA datasets: SentiHood and SemEval-2014 (Task 4). Both models achieve new state-of-the-art results with our QACG-BERT model having the best performance. Furthermore, we provide analyses of the impact of context in the our proposed models. Our work provides more evidence for the utility of adding context-dependencies to pretrained self-attention-based language models for context-based natural language tasks.

Introduction

People are living more of their lives online, both on social media and on e-commerce platforms, and this trend was exacerbated by the recent need for social distancing during the Covid-19 pandemic. Because people are using online review platforms like Yelp and delivery platforms more frequently, understanding the types of emotional content generated on such platforms could yield business insights or provide personalized recommendations (Kang, Yoo, and Han 2012). To this end, Sentiment Analysis techniques have been applied to understand the emotional content generated on microblogs (Kouloumpis, Wilson, and Moore 2011; Severyn and Moschitti 2015), online reviews (e.g., movie and restaurant reviews) (Socher et al. 2013; Kiritchenko et al. 2014), narratives (Wu et al. 2019; Ong et al. 2019) and other online social media (Lwin et al. 2020).

However, user-generated reviews contain more complex information than just a single overall sentiment. A review

of an upscale neighborhood (for potential renters or homebuyers) may praise the safety but express incredulity at the price. Identifying the different *aspects* (e.g., price, safety) embedded within a given text, and their associated sentiment, has been formalized as a task called Aspect-Based Sentiment Analysis (ABSA) (Pontiki et al. 2016; Poria et al. 2016). Targeted ABSA (TABSA) is a more general version of ABSA, when there are multiple targets in a review, each with their associated aspects. For example, given a review of neighborhoods: “*LOC1* area is more expensive but has a better selection of amenities than in *LOC2*” (where *LOC1* and *LOC2* are specially masked tokens), we note that the sentiment depends on the specific target (*LOC1* or *LOC2*) and their aspect. The sentiment towards the **price** of *LOC1* may be negative—and may be more important to a price-conscious student—but positive in terms of **convenience**, while the sentiment towards *LOC2*’s aspects are reversed.

Research using neural models for (T)ABSA has mainly focused on using deep neural networks such as RNNs or attention-gated networks to generate context-dependent sentence embeddings (Saeidi et al. 2016; Tang, Qin, and Liu 2016; Wang et al. 2016; Chen et al. 2017; Ma, Peng, and Cambria 2018; Liu, Cohn, and Baldwin 2018). Recently, with the advent of powerful self-attention models like the Transformer and BERT, Sun, Huang, and Qiu (2019) and Li et al. (2019) both applied pretrained BERT models to (T)ABSA, and showed promising performance improvements. However, these approaches simply used a pretrained BERT model as a blackbox: either via using BERT as an embedding layer or appending the aspect to the input sentence.

We propose to improve the BERT model architecture to be *context-aware*. A context-aware model (Yang et al. 2019) should distribute its attention weights appropriately under different contexts—in (T)ABSA, this means specific targets and aspects. Additionally, by incorporating context into the calculation of attention weights, we aim to enrich the learnt hidden representations of the models. Specifically, we propose two methods to integrate context into the BERT architecture: (1) a Context-Guided BERT (CG-BERT) model adapted from a recent context-aware self-attention network (Yang et al. 2019), which we apply to (T)ABSA; and (2) a novel Quasi-Attention Context-Guided BERT (QACG-BERT) model that learns quasi-attention weights—that could be negative—in a compositional man-

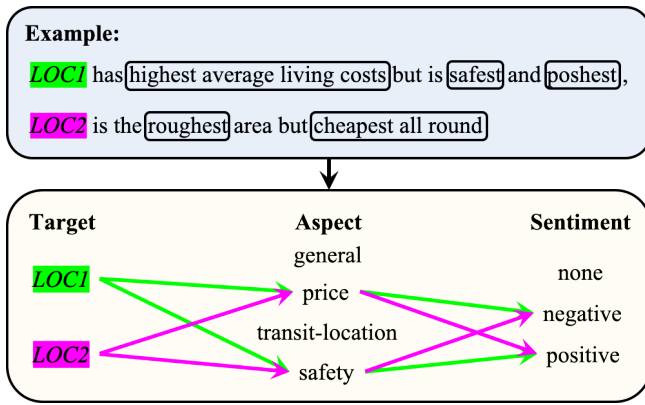


Figure 1: A labelled example from the SentiHood dataset, which compares the price and safety of two locations. The aspects not mentioned in the text are given the label *none*.

ner and which enables subtractive attention, that is lacking in softmax-attention (Tay et al. 2019). In particular, our contribution is three-fold¹:

1. We extend a recently-proposed context-aware self-attention network (Yang et al. 2019) to the (T)ABSA task by formulating a Context-Guided BERT model (CG-BERT).
2. We propose a new Quasi-Attention Context-Guided BERT model (QACG-BERT) that achieves new state-of-the-art (SOTA) results on two (T)ABSA datasets.
3. We analyze how context influences the self-attention and decisions of our models.

Background and Related Work

Self-attention Networks

Self-attention networks, exemplified in the Transformer (Vaswani et al. 2017), have become the *de facto* go-to neural models for a variety of NLP tasks including machine translation (Vaswani et al. 2017), language modeling (Liu and Lapata 2018; Dai et al. 2019) and sentiment analysis (Shen et al. 2018; Wu et al. 2019). BERT, a popular and successful variant of the Transformer, has successfully been applied to various NLP tasks (Reddy, Chen, and Manning 2019; Devlin et al. 2019). In addition, previous studies have also shown some evidence that self-attention weights may learn syntactic (Hewitt and Manning 2019) and semantic (Wu, Nguyen, and Ong 2020) information.

Formulation of *Quasi-attention*

Attention has been successfully applied to many NLP tasks. Various forms of attention are proposed including additive attention (Bahdanau, Cho, and Bengio 2015), dot-product attention (Luong, Pham, and Manning 2015) and scaled dot-product attention used in self-attention (Vaswani et al. 2017). Many of them rely on a *softmax* activation function to calculate attention weights for each position. As a

¹<https://github.com/frankaging/Quasi-Attention-ABSA>

result, the output vector is in the convex hull formed by all other hidden input vectors, preventing the attention gate from learning subtractive relations. Tay et al. (2019) recently proposed a way to overcome this limitation by allowing attention weights to be negative (“quasi” attention), which allows input vectors to add to (+1), not contribute to (0), and even subtract from (−1) the output vector.

Aspect-based Sentiment Analysis

Early research in ABSA introduced classic benchmark datasets and have proposed many baseline methods including lexicon-based and pre-neural classifiers (Pontiki et al. 2014; Kiritchenko et al. 2014; Pontiki et al. 2015, 2016). Since the debut of recurrent neural networks, various RNNs have been developed to generate aspect-aware sentence embeddings and sentiment labels (Tang et al. 2016; Chen et al. 2017; Li et al. 2018). Likewise, researchers have also adapted CNNs (Xue and Li 2018; Huang and Carley 2018), recursive neural networks (Nguyen and Shirai 2015), aspect-aware end-to-end memory networks (Tang, Qin, and Liu 2016) and cognitively inspired deep neural networks (Lei et al. 2019) to generate aspect-aware sentence embeddings.

Motivated by attention mechanisms in deep learning models, many recent ABSA papers have integrated attention into neural models such as RNNs (Wang et al. 2016; Chen et al. 2017; Liu and Zhang 2017; He et al. 2018), CNNs (Zhang, Li, and Song 2019), and memory networks (Ma et al. 2017; Majumder et al. 2018; Liu, Cohn, and Baldwin 2018) to learn different attention distributions for aspects and generate attention-based sentence embeddings. Most recently, self-attention-based models such as BERT have been applied to ABSA, by using BERT as the embedding layer (Song et al. 2019; Yu and Jiang 2019; Lin, Yang, and Lai 2019), or fine-tuning BERT-based models with an ABSA classification output layer (Xu et al. 2019). These papers show that using BERT brings significant performance gains in ABSA.

Targeted Aspect-based Sentiment Analysis

Building on ABSA, Saeidi et al. (2016) proposed a generalized TABSA task (with multiple potential targets) with a new benchmark dataset and LSTM-based baseline models. Various neural models have been proposed for TABSA such as a memory network with delayed context-aware updates (Liu, Cohn, and Baldwin 2018) and interaction-based embedding layers to generate context-aware embeddings (Liang et al. 2019). Researchers have also tried to integrate attention mechanism with LSTMs to predict sentiment for target-aspect pairs (Ma, Peng, and Cambria 2018). With the recent success of BERT-based models, various papers have used BERT to generate contextualized embeddings for input sentences, which are then used to classify sentiment for target-aspect pairs (Huang and Carley 2019; Hu et al. 2019). More recent papers have fine-tuned BERT for TABSA either by (i) constructing auxiliary sentences with different pairs of targets and aspects or (ii) modifying the top-most classification layer to also take in targets and aspects (Rietzler et al. 2020; Sun, Huang, and Qiu 2019; Li et al. 2019). To the best of our knowledge, no work has been

published on modifying the BERT architecture for TABSA tasks. Instead of keeping BERT as a blackbox, we enable BERT to be context-aware by modifying its neural architecture to account for context in its attention distributions.

Approach

We start by defining both TABSA and ABSA tasks, before we introduce our context-guided BERT models (Fig. 2). First, we describe the context-guided BERT model CG-BERT that uses *softmax*-attention, originally proposed by Yang et al. (2019), and the modifications we made to tailor it to the (T)ABSA task. Second, we propose a new neural architecture QACG-BERT that uses quasi-attention adapted from (Tay et al. 2019). Lastly, we describe our methods to formulate our context matrices, and to integrate with pre-trained BERT weights.

TABSA Task

We formulate the Sentihood dataset as a TABSA task. Given a sentence s with a sequence of words $\{w_1, w_2, \dots, w_n\}$ ², where some words are target pronouns $\{w_i, \dots, w_j\}$ from a fixed set T of k predefined targets $\{t_1, \dots, t_k\}$, the goal is to predict sentiment labels for each aspect associated with each unique target mentioned in the sentence. Following the setup in the original Sentihood paper (Saeidi et al. 2016), given a sentence s , a predefined target list T and a predefined aspect list $A = \{\text{general, price, transit-location, safety}\}$, the model predicts a sentiment label $y \in \{\text{none, negative, positive}\}$ for a given pair of $\{(t, a) : (t \in T, a \in A)\}$. Note that the model predicts a single sentiment label for each unique target-aspect pair in a sentence. We show an example of TABSA in Fig. 1.

ABSA Task

We use the SemEval-2014 Task 4 dataset (Pontiki et al. 2014) to formulate an ABSA task: Given a sentence, we predict a sentiment label $y \in \{\text{none, negative, neutral, positive, conflict}\}$ for each aspect $\{a : (a \in A)\}$ with a predefined aspect list $A = \{\text{price, anecdotes, food, ambiance, service}\}$.

Context-Guided *Softmax*-attention

Our Context-Guided BERT (CG-BERT) model is based on the context-aware Transformer model proposed by Yang et al. (2019), which we adapted to the (T)ABSA task. Multi-headed self-attention (Vaswani et al. 2017) is formulated as:

$$\mathbf{A}_{\text{Self-Attn}}^h = \text{softmax} \left(\frac{\mathbf{Q}^h \mathbf{K}^{hT}}{\sqrt{d_h}} \right) \quad (1)$$

where $\mathbf{Q}^h \in \mathbb{R}^{n \times d}$ and $\mathbf{K}^h \in \mathbb{R}^{n \times d}$ are query and key matrices indexed by head h , and $\sqrt{d_h}$ is a scaling factor. We integrate context into BERT by modifying \mathbf{Q} and \mathbf{K} matrices of the original BERT model (Devlin et al. 2019):

$$\begin{bmatrix} \hat{\mathbf{Q}}^h \\ \hat{\mathbf{K}}^h \end{bmatrix} = \left(1 - \begin{bmatrix} \lambda_Q^h \\ \lambda_K^h \end{bmatrix} \right) \begin{bmatrix} \mathbf{Q}^h \\ \mathbf{K}^h \end{bmatrix} + \begin{bmatrix} \lambda_Q^h \\ \lambda_K^h \end{bmatrix} \left(\mathbf{C}^h \begin{bmatrix} \mathbf{U}_Q \\ \mathbf{U}_K \end{bmatrix} \right) \quad (2)$$

²We append a classifier token (i.e., [CLS]) in the beginning of each input sentence as in the BERT model (Devlin et al. 2019).

where $\mathbf{C}^h \in \mathbb{R}^{n \times d_c}$ is the context matrix for each head, $\{\lambda_Q^h, \lambda_K^h\} \in \mathbb{R}^{n \times 1}$ are learnt context weights, and $\{\mathbf{U}_Q, \mathbf{U}_K\} \in \mathbb{R}^{d_c \times d_h}$ are weights of linear layers used to transform input context matrix \mathbf{C}_h . The modified $\hat{\mathbf{Q}}$ and $\hat{\mathbf{K}}$ are then used to calculate context-aware attention weights using the dot-product of both matrices. In contrast to the original implementation (Yang et al. 2019), here we allow both λ_Q^h and λ_K^h to differ across heads, which allows variance in how context is integrated in each head.

We use a zero-symmetric gating unit to learn context gating factors $\{\lambda_Q, \lambda_K\}$:

$$\begin{bmatrix} \lambda_Q^h \\ \lambda_K^h \end{bmatrix} = \tanh \left(\begin{bmatrix} \mathbf{Q}^h \\ \mathbf{K}^h \end{bmatrix} \begin{bmatrix} \mathbf{V}_Q^h \\ \mathbf{V}_K^h \end{bmatrix} + \mathbf{C}^h \begin{bmatrix} \mathbf{U}_Q \\ \mathbf{U}_K \end{bmatrix} \begin{bmatrix} \mathbf{V}_Q^C \\ \mathbf{V}_K^C \end{bmatrix} \right) \quad (3)$$

where $\{\mathbf{V}_Q^h, \mathbf{V}_K^h, \mathbf{V}_Q^C, \mathbf{V}_K^C\} \in \mathbb{R}^{d_h \times 1}$ are weights of linear layers to transform corresponding matrices. We chose to use *tanh* as our activation function as this allows the context to contribute to $\hat{\mathbf{Q}}^h, \hat{\mathbf{K}}^h$ both positively and negatively³. This enriches the representation space of both matrices, and the resulting attention distribution. We note that *tanh* may increase the magnitude of \mathbf{Q}, \mathbf{K} , and large magnitude of \mathbf{Q}, \mathbf{K} may push gradients to excessively small values, which may prevent model learning, as noted by Vaswani et al. (2017) and Britz et al. (2017). However, our results suggest that this did not negatively affect our model performance.

Context-Guided *Quasi*-Attention

Our second neural network model (QACG-BERT) architecture proposes using a Quasi Attention function for (T)ABSA. The value of self-attention weights $\mathbf{A}_{\text{Self-Attn}}^h$ in a vanilla implementation (using *softmax*), is bounded between $[0, 1]$. In other words, it only allows a convex weighted combination of hidden vectors at each position. This allows hidden states to contribute only additively, but not subtractively, to the attended vector. We include a *quasi*-attention calculation to enable learning of both additive as well as subtractive attention (Tay et al. 2019). Formally, we formulate our new attention matrix as a linear combination of a regular *softmax*-attention matrix and a *quasi*-attention matrix:

$$\hat{\mathbf{A}}^h = \mathbf{A}_{\text{Self-Attn}}^h + \lambda_A^h \mathbf{A}_{\text{Quasi-Attn}}^h \quad (4)$$

where λ_A^h is a scalar to represent the compositional factor to control the effect of context on attention calculation. $\mathbf{A}_{\text{Self-Attn}}^h$ is defined as in Eqn. 1. To derive the quasi-attention matrix, we first define two terms quasi-context query \mathbf{C}_Q^h and quasi-context key \mathbf{C}_K^h :

$$\begin{bmatrix} \mathbf{C}_Q^h \\ \mathbf{C}_K^h \end{bmatrix} = \mathbf{C}^h \begin{bmatrix} \mathbf{Z}_Q \\ \mathbf{Z}_K \end{bmatrix} \quad (5)$$

where $\{\mathbf{Z}_Q, \mathbf{Z}_K\} \in \mathbb{R}^{d_e \times d_h}$ are weights of linear layers to transform the raw context matrix, and \mathbf{C}^h is the same context matrix in Eqn. 2. Next, we define the quasi-attention matrix as:

$$\mathbf{A}_{\text{Quasi-Attn}}^h = \alpha \cdot \text{sigmoid} \left(\frac{f_\psi(\mathbf{C}_Q^h, \mathbf{C}_K^h)}{\sqrt{d_h}} \right) \quad (6)$$

³The original implementation (Yang et al. 2019) used *sigmoid*.

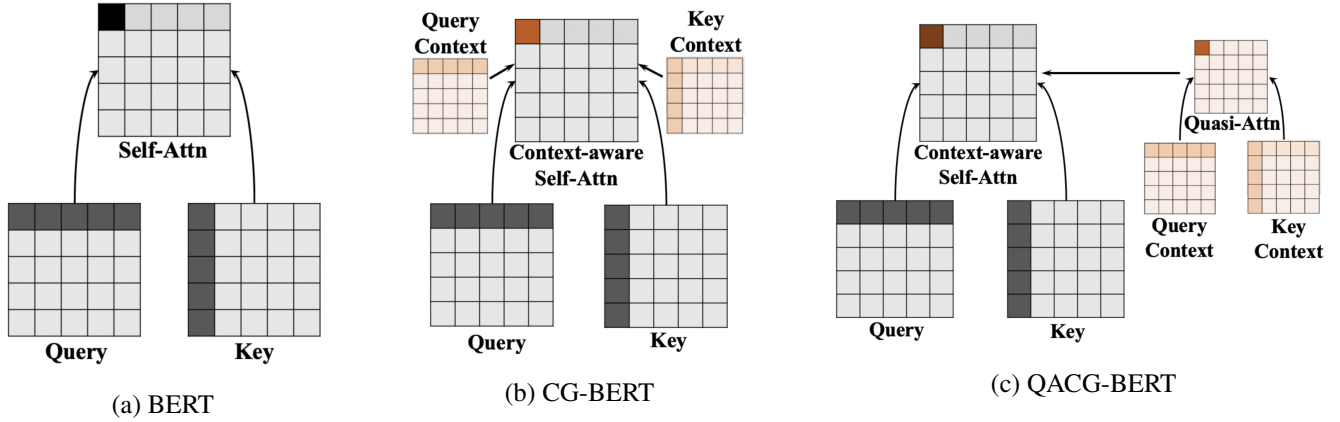


Figure 2: Illustration of the proposed models. (a) The vanilla self-attention network (e.g., BERT) calculates attention weights using the query and key matrices without considering context. (b) The CG-BERT model modifies query and key matrices using context, and then calculate attention weights as in (a). (c) The QACG-BERT model calculates attention weights by combining vanilla attention weights as in (a) with quasi-attention weights calculated using a separate pair of query and key matrices for context. Colors in the grids illustrate matrix operations.

where α is a scaling factor and $f_\psi(\cdot)$ is a similarity measurement to capture similarities between \mathbf{C}_Q^h and \mathbf{C}_K^h . For simplicity, we use dot-product to parameterize f_ψ , and set α to be 1.0. Other f_ψ that have been used include negative $L-1$ distance (Tay et al. 2019). As a result, our $\mathbf{A}_{\text{Quasi-Attn}}$ is bounded between $[0, 1]$. We then derive our *bidirectional* gating factor λ_A as:

$$\begin{bmatrix} \lambda_Q^h \\ \lambda_K^h \end{bmatrix} = \text{sigmoid} \left(\begin{bmatrix} \mathbf{Q}^h \\ \mathbf{K}^h \end{bmatrix} \begin{bmatrix} \mathbf{V}_Q^h \\ \mathbf{V}_K^h \end{bmatrix} + \begin{bmatrix} \mathbf{C}_Q^h \\ \mathbf{C}_K^h \end{bmatrix} \begin{bmatrix} \mathbf{V}_Q^C \\ \mathbf{V}_K^C \end{bmatrix} \right) \quad (7)$$

$$\lambda_A^h = 1 - (\beta \cdot \lambda_Q^h + \gamma \cdot \lambda_K^h) \quad (8)$$

where $\{\beta, \gamma\}$ are scalars that control the composition weightings. For simplicity, we set $\{\beta, \gamma\} = 1.0$. We formulate the gating factor to be *bidirectional*, meaning it takes on both positive and negative values, and the output is bounded between $\{-1, 1\}$. Our intuition is that the context-based *quasi*-attention may contribute either positively or negatively to the final attention weights. Consider Eqn. 4: as the first term $\mathbf{A}_{\text{Self-Attn}}^h$ is in $\{0, 1\}$, and the second term is made up of a term (λ_A) that is in $[-1, 1]$ and another ($\mathbf{A}_{\text{Quasi-Attn}}$) that is in $[0, 1]$, hence the final attention $\hat{\mathbf{A}}$ lies in $[-1, 2]$. That is to say, the final attention weights can take values representing compositional operations including subtraction (-1), deletion ($\times 0$), inclusion/addition ($+1/+2$) among hidden vectors across positions. We hypothesize that the *quasi*-attention provides a richer method to integrate context into the calculation of attention.

Classification

We use the final hidden state (the output of the final layer of the BERT model) of the first classifier token (i.e., [CLS]) as the input to the final classification layer for a C -class classification. This is similar to previous studies (Sun, Huang, and Qiu 2019). For a given input sentence, we denote this

vector as $e_{\text{CLS}} \in \mathbb{R}^{1 \times d}$. Then, the probability of each sentiment class y is given by $y = \text{softmax}(e_{\text{CLS}} \mathbf{W}_{\text{CLS}}^T)$ where $\mathbf{W}_{\text{CLS}} \in \mathbb{R}^{C \times d}$ are the weights of the classification layer, and $y \in \mathbb{R}^{1 \times C}$. The label with highest probability will be selected as the final prediction.

Context Matrix

We use a single integer to represent a context associated with an aspect and a target in any (T)ABSA task, and only an aspect in the ABSA task. We transform these integers into embeddings via a trainable embedding layer, which derives \mathbf{C}^h in CG-BERT, and \mathbf{C}_Q^h and \mathbf{C}_K^h in QACG-BERT. For example, given $|t|$ targets and $|a|$ aspects for any (T)ABSA task, the total number of possible embeddings is $|t| \cdot |a|$. We then concatenate the context embedding with the hidden vector for each position $\mathbf{E} \in \mathbb{R}^{n \times d}$, and pass them into a feed-forward linear layer with a residual connection to formulate the context matrix $\mathbf{C}^h = [\mathbf{E}_c, \mathbf{E}] \mathbf{W}_c^T$ where $\mathbf{E}_c \in \mathbb{R}^{n \times d}$ is the context embedding and $\mathbf{W}_c \in \mathbb{R}^{d \times 2d}$ are the learnt weights for this feed-forward layer.

Integration with Pretrained BERT

Previous studies show that fine-tuning pretrained BERT models increases performance significantly in many NLP tasks (Sun, Huang, and Qiu 2019; Rietzler et al. 2020). Since our models share most of the layers with a standard BERT model, we import weights from pretrained BERT models for these overlapping layers. The weights of the newly added layers are initialized to be small⁴ random numbers drawn from a normal distribution $\mathcal{N}(0, \sigma^2)$ with $\sigma = e^{-3}$. As a result, the gating factors in Eqn. 2 and Eqn. 1 start at values close to zero. This initialization enables the task-specific

⁴We also tried initializing the weights in the newly added layers with larger variance, and found similar performance.

Model	Aspect Categorization			Sentiment	
	Strict Accuracy 1%	Macro-F1 1%	AUC 1%	Accuracy 1%	AUC 1%
LR (Saeidi et al. 2016)	-	39.3	92.4	87.5	90.5
LSTM-Final (Saeidi et al. 2016)	-	68.9	89.8	82.0	85.4
LSTM-Loc (Saeidi et al. 2016)	-	69.3	89.7	81.9	83.9
SenticLSTM (Ma, Peng, and Cambria 2018)	67.4	78.2	-	89.3	-
Dmu-Entnet (Liu, Cohn, and Baldwin 2018)	73.5	78.5	94.4	91.0	94.8
BERT-single (Sun, Huang, and Qiu 2019)	73.7	81.0	96.4	85.5	84.2
BERT-pair (Sun, Huang, and Qiu 2019)	79.8	87.9	97.5	93.6	97.0
CG-BERT (adapted from Yang et al. 2019)	79.7 (.3)	87.1 (.2)	97.5 (.2)	93.7 (.2)	97.2 (.2)
CG-BERT (auxiliary sentences)	80.1 (.2)	88.1 (.1)	97.6 (.1)	93.1 (.3)	97.0 (.2)
QACG-BERT	79.9 (.4)	88.6 (.2)	97.3 (.2)	93.8 (.2)	97.8 (.2)
QACG-BERT (auxiliary sentences)	80.9 (.3)	89.7 (.2)	97.8 (.2)	93.7 (.3)	97.1 (.2)

Table 1: Model performance on SentiHood TABSA dataset with best performances bolded. Top: Best results reported for previous models. “-” indicates not reported in the original paper. Bottom: Means and standard deviations across runs for our models. We also report our models trained with appended auxiliary sentences, the same inputs as Sun, Huang, and Qiu (2019).

weights to start from the pretrained weights and slowly diverge during training.

Experiments

Datasets

We evaluate our models with two datasets in English. For the TABSA task, we used the Sentihood dataset⁵ which was built by questions and answers from Yahoo! with location names of London, UK. It consists of 5,215 sentences, with 3,862 sentences containing a single target and 1,353 sentences containing multiple targets. For each sentence, we predict sentiment label y for each target-aspect pair (t, a) . For the ABSA task, we used the dataset from SemEval 2014, Task 4⁶, which contains 3,044 sentences from restaurant reviews. For each sentence, we predict the sentiment label y for each aspect a . Each dataset is partitioned to train, development and test sets as in its original paper.

As in previous studies (Pontiki et al. 2014; Saeidi et al. 2016) we define two subtasks for each dataset: (1) aspect categorization and 2) aspect-based sentiment classification. For aspect categorization, the problem is to detect whether a aspect a is mentioned (i.e., *none* means not mentioned) in the input sentence for a target t if it is a TABSA task. For aspect-based sentiment classification, we give the model aspects that present (i.e., ignoring *none*’s) and have the model predict the valence of the sentiment (i.e., potential labels include *negative* and *positive* for Sentihood, and *negative*, *neutral*, *positive*, *conflicting sentiment* for Semeval Task 4).

Experiment Settings

As in the original BERT-base model (Devlin et al. 2019), our models consists of 12 heads and 12 layers, with hidden layer size 768. The total number of parameters for both of our models increased slightly due to the additional linear layers added comparing to previous BERT-based models for ABSA

tasks (Sun, Huang, and Qiu 2019) which consists of about 110M parameters. The CG-BERT and QACG-BERT consists of about 124M parameters. We trained for 25 epochs with a dropout probability of 0.1. The initial learning rate is $2e^{-5}$ for all layers, with a batch size of 24. We used the pretrained weights from the uncased BERT-base model⁷.

We used a single Standard NC6 instance on Microsoft Azure, which is equipped with a single NVIDIA Tesla K80 GPU with 12G Memory. Training both models across two datasets took approximately 11 hours.

Exp-I: TABSA

For the TABSA task, we compared the performance of our models with previous models in Table 1. Following Ma, Peng, and Cambria (2018) and Sun, Huang, and Qiu (2019), for aspect categorization (is a given aspect present in the sentence? If aspect is not present, the label is by definition *none*), we report strict accuracy (model needs to correctly identify all aspects for a given target in the sentence to be counted as accurate), Macro-F1 (the harmonic mean of the Macro-precision and Macro-recall of all targets.) and AUC. For sentiment classification (given an aspect present in the sentence, is the valence *negative* or *positive*?), we report accuracy and AUC.

Results Our results showed that modifying BERT to be context-aware resulted in increased performance, surpassing the previous state-of-the-art. Across multiple evaluation metrics, our proposed quasi-attention model with pretrained BERT weights (QACG-BERT) performed the best. Additionally, we also evaluated our models trained with inputs appended with auxiliary sentences as in (Sun, Huang, and Qiu 2019) for comparison. Specifically, we append the target and aspect, for example, “[SEP] - location - 1 - price”, to the input. Our results showed that for some metrics, using auxiliary sentences improves performance but not in others.

⁵<https://github.com/uclnlp/jack/tree/master/data/sentihood>

⁶<http://alt.qcri.org/semeval2014/task4/>

⁷https://storage.googleapis.com/bert_models/2020_02_20/uncased.L-12_H-768_A-12.zip

Model	Aspect Categorization			Sentiment		
	Precision /%	Recall /%	F1 /%	Binary /%	3-class /%	4-class /%
XRCE (Brun, Popa, and Roux 2014)	83.23	81.37	82.29	-	-	78.1
NRC-Canada (Kiritchenko et al. 2014)	91.04	86.24	88.58	-	-	82.9
BERT-single (Sun, Huang, and Qiu 2019)	92.78	89.07	90.89	93.3	86.9	83.7
BERT-pair (Sun, Huang, and Qiu 2019)	93.57	90.83	92.18	95.6	89.9	85.9
CG-BERT (adapted from Yang et al. 2019)	93.02 (.27)	90.00 (.33)	91.49 (-)	94.3 (.3)	89.9 (.2)	85.6 (.4)
CG-BERT (auxiliary sentences)	93.12 (.37)	90.17 (.23)	91.62 (-)	94.7 (.3)	90.1 (.4)	85.7 (.5)
QACG-BERT	94.38 (.31)	90.97 (.28)	92.64 (-)	95.6 (.4)	90.1 (.3)	86.8 (.8)
QACG-BERT (auxiliary sentences)	94.27 (.39)	90.12 (.34)	92.14 (-)	95.8 (.3)	90.4 (.5)	86.9 (.7)

Table 2: Model performance on the Semeval-2014 Task 4 ABSA dataset, with best performances bolded. Aspect categorization and sentiment classification corresponds to Subtask 3 and Subtask 4 respectively. Top: Best results reported for previous models. “-” indicates not reported in the original paper. Bottom: Means and standard deviations across runs for our models.

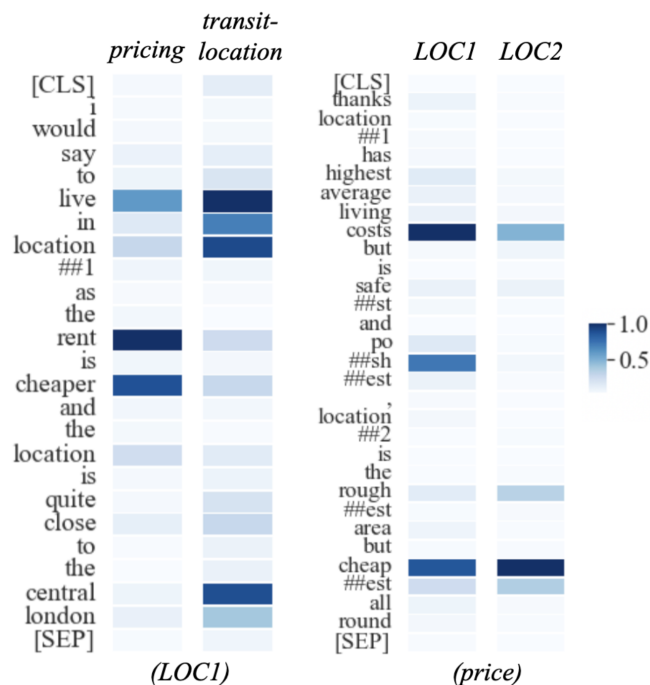


Figure 3: Examples of relevance scores via gradient sensitivity analysis for different aspects and locations, from our QACG-BERT model on the SentiHood TABSA test set. Left: Gradients associated with two aspects {pricing, transit-location} for target LOC1. Right: Gradients associated with the general aspect, for two targets LOC1 and LOC2. Values are normalized with respect to the maximum value within a sentence.

Exp-II: ABSA

For the ABSA task, we compared our models with multiple best performing models for the SemEval-2014 Task 4 dataset in Table 2. Following Pontiki et al. (2016), for aspect categorization, we report Precision, Recall, and F1. For aspect-based sentiment classification, we report accuracies for three different evaluations: binary classification (i.e.,

negative or positive), 3-class classification (negative, neutral or positive) and 4-class classification (negative, neutral, positive or conflict).

Results Consistent with Exp-I, our models improved over previous state-of-the-art models on the SemEval-2014 Task 4 dataset. For aspect categorization, our CG-BERT performed slightly worse than previous best performing models while QACG-BERT performed surpass the best model in F1 scores. For aspect sentiment classification, QACG-BERT performs better than CG-BERT and both models surpass previous SOTA performance. Similar to TABSA, our results show slight drops in some metrics while maintain same performance in others when trained with appended auxiliary sentences.

Feature Importance Analysis

To visualization importance of input tokens, we conduct gradient sensitivity analysis (Li et al. 2016; Arras et al. 2017). In Fig. 3, we visualize the gradient scores of our QACG-BERT model over two input sentences for which our model predicts sentiment labels correctly. For the first example, words like **rent** and **cheaper** are most relevant to the *price* aspect while words like **live**, **location**, **central** are more relevant to the *transit-location* aspect. Thus, our model learns how to pick out words that are important under different contexts. The second example concerns two targets (i.e., locations) for the same aspect, *price*, where the sentiment label for *LOC1* is *negative* while for *LOC2* is *positive*. For *LOC1*, the model correctly identifies **costs** (which in context refers to *LOC1*) and **posh**⁸. By contrast, the model identifies **cheap** when given *LOC2* as context, and assigns a greater gradient value, compared to the same word for *LOC1*. As a result, the model is able to identify words corresponding to different targets and different aspects.

⁸The BERT Tokenizer breaks up ‘poshest’ into po-sh-est. We note that the BERT vocabulary does not have ‘posh’ in it, but after fine-tuning, appears to learn that the word is price-relevant.

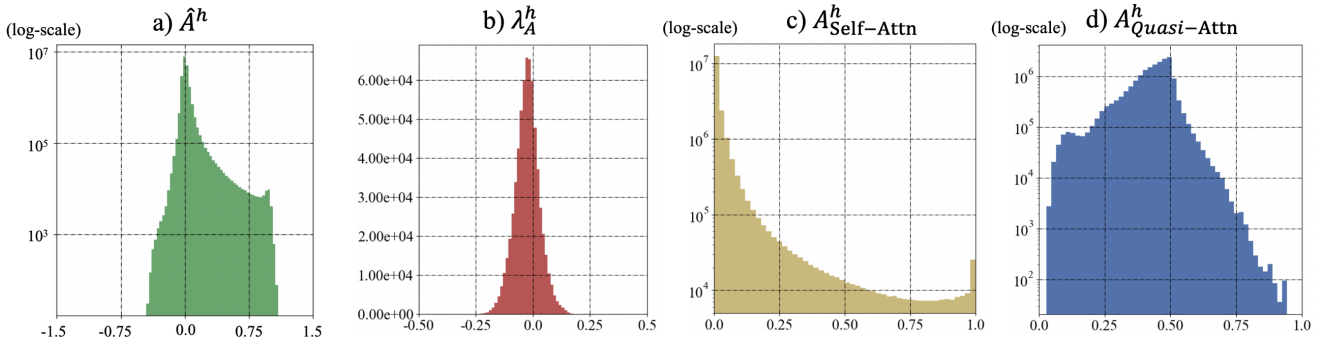


Figure 4: Histogram of QACG-BERT weights on $n = 200$ randomly selected examples from SentiHood test set where true label is either *negative* or *positive*. For the matrices, each value in the matrix is a data point in the histogram. We note that the vertical axes of (a), (c) and (d) are in log-scale.




Figure 5: Example of attention weights extracted from different matrices for an aspect-location pair, (*LOCI*, *price*), from the SentiHood test set. Values are normalized with respect to the absolute maximum value within a sentence. We note the high negative quasi-attention on **nice**, which is often a positive word, but not in the context of *price*.

Quasi-attention Visualization

We inspected the quasi-attention parameters learnt by our models. Specifically, we took the QACG-BERT model trained on the Sentihood dataset and extracted values for the following four variables: the final attention weights matrix \hat{A}^h , the *bidirectional* gating factor matrix λ_A^h , the vanilla self-attention matrix $A_{\text{Self-Attn}}^h$ and the context-guided quasi-attention matrix $A_{\text{Quasi-Attn}}^h$. Fig. 4 illustrates the histogram of values drawn from 200 examples from the test set.

We made several key observations. First, the behaviour of λ_A^h follows our intuition; it acts as a bidirectional control gate, with slightly more negative values, and determines whether context contributes to attention positively or nega-

tively. Second, the learnt weights in $A_{\text{Quasi-Attn}}^h$ are not near zero, with the mass of the distribution between .25 and .50, thus, it does contribute to attention. Lastly, the non-zero weights in the final matrix \hat{A}^h are mainly positive, but some of the weights take on negative values due to the bidirectional gating factor. This is important as it enables the model to attend to and “de-attend from” different parts of the input.

Finally, we turn to visualizing the quasi-attention. In Fig. 5, we visualize the weights of $A_{\text{Self-Attn}}^h$ and the product $\lambda_A^h A_{\text{Quasi-Attn}}^h$ extracted from one of the heads in the first self-attention layer, over an example input sentence from the test set. For this sentence, our model correctly predicted the sentiment label *negative* for the *price* aspect of *LOCI*. In this example, **nice** is a positive word for aspects like *general*. However, with respect to *price*, **nice** actually suggests a higher cost and thus is more negative. We see that while $A_{\text{Self-Attn}}^h$ on **nice** is high, $A_{\text{Quasi-Attn}}^h$ on **nice** is negative, which “subtracts” from the final attention. As a result, the sum derived by Eqn. 4 makes the final attention weight on the word **nice** less positive. We note that empirically, we find that the total attention \hat{A}^h is usually positive (Fig. 4(a)), i.e., quasi-attention, when negative, tends to be smaller in magnitude than self-attention.

Conclusion

We proposed two context-aware BERT-based models for ABSA and TABSA, which outperformed state-of-the-art results on two datasets. Our first CG-BERT model introduced a way of integrating context into pretrained BERT for ABSA. The second QACG-BERT model allowed quasi-attention, which enables compositional attention including subtraction (-1), deletion ($\times 0$), inclusion/addition ($+1/+2$). Our results and analyses show strong performance results, especially for our QACG-BERT model, in solving (T)ABSA tasks, and suggest potential success for such context-aware mechanisms for other context-based tasks in NLP.

Acknowledgements

This work was supported in part by a Singapore Ministry of Education Academic Research Fund Tier 1 grant to DCO.

References

- Arras, L.; Montavon, G.; Müller, K.-R.; and Samek, W. 2017. Explaining Recurrent Neural Network Predictions in Sentiment Analysis. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, 159–168.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the 4th International Conference on Learning Representations (ICLR)*.
- Britz, D.; Goldie, A.; Luong, M.-T.; and Le, Q. 2017. Massive Exploration of Neural Machine Translation Architectures. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 1442–1451.
- Brun, C.; Popa, D. N.; and Roux, C. 2014. XRCE: Hybrid Classification for Aspect-based Sentiment Analysis. In *SemEval@ COLING*, 838–842. Citeseer.
- Chen, P.; Sun, Z.; Bing, L.; and Yang, W. 2017. Recurrent attention network on memory for aspect sentiment analysis. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 452–461.
- Dai, Z.; Yang, Z.; Yang, Y.; Carbonell, J. G.; Le, Q.; and Salakhutdinov, R. 2019. Transformer-XL: Attentive Language Models beyond a Fixed-Length Context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2978–2988.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 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*.
- He, R.; Lee, W. S.; Ng, H. T.; and Dahlmeier, D. 2018. Exploiting Document Knowledge for Aspect-level Sentiment Classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 579–585.
- Hewitt, J.; and Manning, C. D. 2019. A structural probe for finding syntax in word representations. 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)*, 4129–4138.
- Hu, M.; Zhao, S.; Guo, H.; Cheng, R.; and Su, Z. 2019. Learning to Detect Opinion Snippet for Aspect-Based Sentiment Analysis. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, 970–979.
- Huang, B.; and Carley, K. M. 2018. Parameterized Convolutional Neural Networks for Aspect Level Sentiment Classification. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 1091–1096.
- Huang, B.; and Carley, K. M. 2019. Syntax-Aware Aspect Level Sentiment Classification with Graph Attention Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 5472–5480.
- Kang, H.; Yoo, S. J.; and Han, D. 2012. Senti-lexicon and improved Naïve Bayes algorithms for sentiment analysis of restaurant reviews. *Expert Systems with Applications* 39(5): 6000–6010.
- Kiritchenko, S.; Zhu, X.; Cherry, C.; and Mohammad, S. 2014. NRC-Canada-2014: Detecting aspects and sentiment in customer reviews. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, 437–442.
- Kouloumpis, E.; Wilson, T.; and Moore, J. 2011. Twitter sentiment analysis: The good the bad and the omg! In *Fifth International AAAI conference on weblogs and social media*.
- Lei, Z.; Yang, Y.; Yang, M.; Zhao, W.; Guo, J.; and Liu, Y. 2019. A human-like semantic cognition network for aspect-level sentiment classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 6650–6657.
- Li, J.; Chen, X.; Hovy, E.; and Jurafsky, D. 2016. Visualizing and understanding neural models in NLP. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Li, X.; Bing, L.; Lam, W.; and Shi, B. 2018. Transformation Networks for Target-Oriented Sentiment Classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 946–956.
- Li, X.; Bing, L.; Zhang, W.; and Lam, W. 2019. Exploiting BERT for End-to-End Aspect-based Sentiment Analysis. *W-NUT 2019* 34.
- Liang, B.; Du, J.; Xu, R.; Li, B.; and Huang, H. 2019. Context-aware Embedding for Targeted Aspect-based Sentiment Analysis. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4678–4683.
- Lin, P.; Yang, M.; and Lai, J. 2019. Deep Mask Memory Network with Semantic Dependency and Context Moment for Aspect Level Sentiment Classification. In *IJCAI*, 5088–5094.
- Liu, F.; Cohn, T.; and Baldwin, T. 2018. Recurrent Entity Networks with Delayed Memory Update for Targeted Aspect-Based Sentiment Analysis. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, 278–283.
- Liu, J.; and Zhang, Y. 2017. Attention modeling for targeted sentiment. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, 572–577.
- Liu, Y.; and Lapata, M. 2018. Learning structured text representations. *Transactions of the Association for Computational Linguistics* 6: 63–75.
- Luong, M.-T.; Pham, H.; and Manning, C. D. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 1412–1421.
- Lwin, M. O.; Lu, J.; Sheldenkar, A.; Schulz, P. J.; Shin, W.; Gupta, R.; and Yang, Y. 2020. Global sentiments surrounding the COVID-19 pandemic on Twitter: analysis of Twitter trends. *JMIR Public Health and Surveillance* 6(2): e19447.
- Ma, D.; Li, S.; Zhang, X.; and Wang, H. 2017. Interactive attention networks for aspect-level sentiment classification. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 4068–4074.
- Ma, Y.; Peng, H.; and Cambria, E. 2018. Targeted Aspect-Based Sentiment Analysis via Embedding Commonsense Knowledge into an Attentive LSTM. In *AAAI*, 5876–5883.

- Majumder, N.; Poria, S.; Gelbukh, A.; Akhtar, M. S.; Cambria, E.; and Ekbal, A. 2018. IARM: Inter-aspect relation modeling with memory networks in aspect-based sentiment analysis. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3402–3411.
- Nguyen, T. H.; and Shirai, K. 2015. Phrasernn: Phrase recursive neural network for aspect-based sentiment analysis. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 2509–2514.
- Ong, D.; Wu, Z.; Zhi-Xuan, T.; Reddan, M.; Kahlale, I.; Mattek, A.; and Zaki, J. 2019. Modeling emotion in complex stories: the Stanford Emotional Narratives Dataset. *IEEE Transactions on Affective Computing*.
- Pontiki, M.; Galanis, D.; Papageorgiou, H.; Androutsopoulos, I.; Manandhar, S.; Al-Smadi, M.; Al-Ayyoub, M.; Zhao, Y.; Qin, B.; De Clercq, O.; et al. 2016. Semeval-2016 task 5: Aspect based sentiment analysis. In *10th International Workshop on Semantic Evaluation (SemEval 2016)*.
- Pontiki, M.; Galanis, D.; Papageorgiou, H.; Manandhar, S.; and Androutsopoulos, I. 2015. Semeval-2015 task 12: Aspect based sentiment analysis. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, 486–495.
- Pontiki, M.; Galanis, D.; Pavlopoulos, J.; Papageorgiou, H.; Androutsopoulos, I.; and Manandhar, S. 2014. SemEval-2014 Task 4: Aspect Based Sentiment Analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, 27–35.
- Poria, S.; Cambria, E.; Hazarika, D.; and Vij, P. 2016. A Deeper Look into Sarcastic Tweets Using Deep Convolutional Neural Networks. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, 1601–1612.
- Reddy, S.; Chen, D.; and Manning, C. D. 2019. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics* 7: 249–266.
- Rietzler, A.; Stabinger, S.; Opitz, P.; and Engl, S. 2020. Adapt or Get Left Behind: Domain Adaptation through BERT Language Model Finetuning for Aspect-Target Sentiment Classification. In *Proceedings of The 12th Language Resources and Evaluation Conference*, 4933–4941.
- Saeidi, M.; Bouchard, G.; Liakata, M.; and Riedel, S. 2016. SentiHood: Targeted Aspect Based Sentiment Analysis Dataset for Urban Neighbourhoods. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, 1546–1556.
- Severyn, A.; and Moschitti, A. 2015. Twitter sentiment analysis with deep convolutional neural networks. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 959–962.
- Shen, T.; Zhou, T.; Long, G.; Jiang, J.; Pan, S.; and Zhang, C. 2018. DiSAN: Directional Self-Attention Network for RNN/CNN-free language understanding. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Socher, R.; Perelygin, A.; Wu, J.; Chuang, J.; Manning, C. D.; Ng, A. Y.; and Potts, C. 2013. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, 1631–1642.
- Song, Y.; Wang, J.; Jiang, T.; Liu, Z.; and Rao, Y. 2019. Attentional encoder network for targeted sentiment classification. *arXiv preprint arXiv:1902.09314*.
- Sun, C.; Huang, L.; and Qiu, X. 2019. Utilizing BERT for Aspect-Based Sentiment Analysis via Constructing Auxiliary Sentence. 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)*, 380–385.
- Tang, D.; Qin, B.; Feng, X.; and Liu, T. 2016. Effective LSTMs for Target-Dependent Sentiment Classification. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics*, 3298–3307.
- Tang, D.; Qin, B.; and Liu, T. 2016. Aspect Level Sentiment Classification with Deep Memory Network. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 214–224.
- Tay, Y.; Luu, A. T.; Zhang, A.; Wang, S.; and Hui, S. C. 2019. Compositional De-Attention Networks. In *Advances in Neural Information Processing Systems*, 6135–6145.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, 5998–6008.
- Wang, Y.; Huang, M.; Zhu, X.; and Zhao, L. 2016. Attention-based LSTM for aspect-level sentiment classification. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 606–615.
- Wu, Z.; Nguyen, T.-S.; and Ong, D. 2020. Structured Self-Attention Weights Encodes Semantics in Sentiment Analysis. In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, 255–264.
- Wu, Z.; Zhang, X.; Zhi-Xuan, T.; Zaki, J.; and Ong, D. C. 2019. Attending to emotional narratives. In *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*, 648–654. IEEE.
- Xu, H.; Liu, B.; Shu, L.; and Philip, S. Y. 2019. BERT Post-Training for Review Reading Comprehension and Aspect-based Sentiment Analysis. 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)*, 2324–2335.
- Xue, W.; and Li, T. 2018. Aspect Based Sentiment Analysis with Gated Convolutional Networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2514–2523.
- Yang, B.; Li, J.; Wong, D. F.; Chao, L. S.; Wang, X.; and Tu, Z. 2019. Context-aware self-attention networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 387–394.
- Yu, J.; and Jiang, J. 2019. Adapting BERT for target-oriented multimodal sentiment classification. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 5408–5414. AAAI Press.
- Zhang, C.; Li, Q.; and Song, D. 2019. Aspect-based Sentiment Classification with Aspect-specific Graph Convolutional Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4560–4570.