

Flavours of Convolution for Unsupervised Aspect Extraction and Aspect-based Sentiment Analysis

Soumick Chatterjee^{1,2,3,*}, Sowmya Prakash² and Andreas Nürnberger^{2,3,4}

¹Genomics Research Centre, Human Technopole, Viale Rita Levi-Montalcini, 1, 20157 Milan, Italy

²Faculty of Computer Science, Otto von Guericke University Magdeburg, Universitätspl. 2, 39106 Magdeburg, Germany

³Data and Knowledge Engineering Group, Otto von Guericke University Magdeburg, Universitätspl. 2, 39106 Magdeburg, Germany

⁴Centre for Behavioural Brain Sciences, Magdeburg, Germany

Abstract

Sentiment analysis, the task of understanding the underlying sentiment from a given data, is frequently used in various fields, from market research to recommender systems. To analyse different sentiments for different aspects present in a given sentence, aspect-based sentiment analysis (ABSA) approaches have been proposed. The task of ABSA can be divided into two subtasks: aspect extraction (AE) and aspect-based polarity detection (ABPD). Most of the state-of-the-art approaches are based on some form of neural networks using convolutional layers. Recently, flavours of convolution, like extremely separated convolution layer (XSepConv) - which can reduce computational cost along with the parameter size of large kernels and depth-wise over-parameterised convolutional layer (DOConv) - which can improve the training efficiency, have been proposed. They have shown their superior capabilities when it comes to image-related tasks, but they have not been explored for textual tasks like ABSA. This paper performs unsupervised AE and weakly supervised ABSA using those specialised convolutional layers. It could be shown that using such layers instead of traditional convolutional layers can significantly improve the performance of the model.

Keywords

Aspect extraction, Aspect based sentiment analysis, Extremely Separated Convolution, Depthwise Overparameterised, Convolutional Layer

1. Introduction

Historically rooted in the intricacies of linguistic study, the domain of Natural Language Processing (NLP) has burgeoned over the years, assimilating the expertise of computer scientists, mathematicians, statisticians, and several other fields, transmogrified into a prominent offshoot of artificial intelligence, adept at navigating vast corpora. A typical NLP pipeline usually consists of data cleaning, data preprocessing, data modelling, and model evaluation based on various metrics¹. Aspect-based sentiment analysis (ABSA) is one of the rapidly progressing

NL4AI 2023: Seventh Workshop on Natural Language for Artificial Intelligence, November 6-7th, 2023, Rome, Italy [1]

*Corresponding author.

✉ contact@soumick.com (S. Chatterjee); andreas.nuernberger@ovgu.de (A. Nürnberger)

🌐 <https://www.soumick.com/> (S. Chatterjee)

🆔 0000-0001-7594-1188 (S. Chatterjee); 0000-0003-4311-0624 (A. Nürnberger)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

¹Intro to NLP: towardsdatascience.com (last access: 17.10.2023)

domains in NLP. ABSA is applicable across several domains to perform market research, such as e-Commerce, manufacturing, healthcare, etc. In a given statement: “I liked their food, but the ambience was awful”, the aspects in the statement are “food” and “ambience” expressed as positive and negative sentiment, respectively. Thus, the ABSA concentrates on sentiments with respect to aspects and, hence, provides fine-grained information.

There are two main tasks involved in the Aspect-Based Sentiment Analysis (ABSA):

1. **Aspect extraction (AE) [2]:** This method detects aspects from the unstructured text based on the context. AE mainly deals with exploring aspects of interest and grouping the extracted words into predefined aspect terms in the text data.
2. **Aspect-based polarity detection (ABPD) [3]:** Once the aspects are extracted from a sentence, ABPD methods are applied to detect the sentiments of the extracted aspects. There are numerous methods that perform aspect-based sentiment analysis. The focus of this work is on using word embedding, which is a weakly supervised method. For example, consider the review “Masanielli is a great place to meet, they serve delicious pizze”. Here, humans can recognise the aspect terms, but it is difficult for a neural network to identify them correctly without annotations. The current work learns aspect and sentiment embedding together to recognise both aspect terms and the corresponding opinions simultaneously.

1.1. Related Work

Various approaches have been proposed in recent times for both AE and ABPD - using supervised, unsupervised, as well as weakly supervised learning.

1.1.1. Aspect Extraction (AE)

Although supervised learning is a popular approach, recently several researchers have shown that unsupervised neural networks can also provide good results - without the need to have manually annotated data. The sequential rule mining approach proposed by Liu et al. [4] analyses the text review at different detail levels based on the characteristics of the product and the set of opinions. He et al. [5] proposed an unsupervised aspect extraction method using an attention model. Sokhin et al. [2] extended that idea to a multi-attention convolutional model (CMAM) for unsupervised aspect extraction which can identify aspects based on the importance of terms and their related terms.

1.1.2. Aspect-Based Polarity Detection (ABPD)

Over the years, a variety of techniques have been proposed for sentiment prediction or polarity detection and can be grouped into two general categories: sentence-based [6, 7] and aspect-based [8, 9]. It is worth noting that aspect-based polarity detection offers a distinct advantage over its sentence-based counterpart. This is because a single sentence can contain varying polarities related to different aspects. Although many of the approaches are supervised, there have also been introductions of semi-supervised or weakly supervised techniques [10].

1.1.3. Aspect-Based Sentiment Analysis (ABSA) - as a single pipeline:

Many of the proposed approaches perform both AE and ABPD tasks in a single pipeline - using techniques such as multi-element join detection [11], multitask learning [12]. Huang et al. [3] proposed a weakly supervised ABSA model using joint aspect-sentiment topic embedding.

1.1.4. Flavours of Convolution:

One of the most common types of layer used in neural networks for aspect extraction and aspect-based polarity detection is the convolutional layer. Specialised types of convolution layers, such as Extremely Separated Convolution Layer (XSepConv) [13] and the Depth-Wide Overparameterised Convolutional Layer (DOConv) [14], which are usually used for image-related tasks, have been shown to outperform the conventional convolution layer. But they have not yet been used for text-based tasks.

The XSepConv layer combines spatially separable architecture into depth-wise convolution to reduce the parameter size of large kernels while also reducing the computational costs. This layer additionally uses an additional depth-wise convolution of size 2×2 with an advanced symmetric padding approach, which neutralises the impact from spatially separable convolutions. This architecture can be a cost-effective option in comparison to depth-wise convolution with larger kernel sizes. Experiments performed by the authors [13] on four benchmark datasets (CIFAR-10, SVHN, CIFAR-100, and Tiny-ImageNet) show better performance by replacing the depth-wise convolution kernel with XSepConv for image classification tasks, while decreasing the computation and size of parameters. The DOConv architecture overparameterises the convolution layer by adding depth-wise convolution that consists of a separate kernel for each input channel. The authors [14] have shown that this architecture's over-parameterisation speeds up the training process of convolution neural networks and performs better in many tasks like image classification, detection, and segmentation.

1.2. Contribution

Previous work on aspect-based sentiment analysis is mainly focused on supervised or unsupervised learning using regular 2-D Convolutions. In this research, the weakly supervised and unsupervised approaches are redesigned with specialised convolution layers. The current work proposes a novel extremely separable convolution-based architecture for unsupervised aspect extraction and depth-wise over-parameterised convolution-based architecture for weakly supervised aspect-based sentiment. Both architectures are built by replacing the regular 2D-Convolution layers in the neural network. The results from our experiment show statistically significant improvements in comparison to the state-of-the-art models.

2. Methodology

Two different types of tasks were performed in this research with the help of two different models: unsupervised aspect extraction (AE) and weakly supervised aspect-based sentiment analysis (ABSA). The latter performed AE and ABPD together inside a single model. Fig. 1 portrays the overall workflow of the two models.

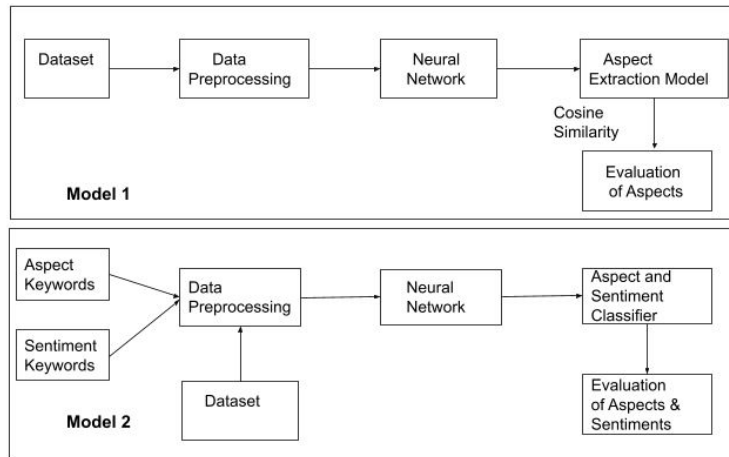


Figure 1: Workflow of the Aspect-Based Sentiment Analysis pipeline. Model 1: Unsupervised Aspect Extraction, Model 2: Weakly supervised Aspect-based Sentiment Analysis

2.1. Data

For the AE models, the Citysearch data [15] are used as the training corpus, and for validation purposes, the Semeval-2016 restaurant data [16] are used as domain experts manually annotate them. Following He et al. [5], specific gold labels such as food, staff, and ambience from the restaurant domain were used to evaluate the performance of the AE models in the restaurant dataset.

For ABSA, two different benchmark datasets were used representing two different domains: the Yelp dataset [3] and Amazon reviews [17], for the restaurant and laptop domain, respectively. For evaluation purposes, the SemEval-2015 [18] and SemEval-2016 [16] datasets were used. Text reviews with more than one label or text with no labels are not considered for model evaluation. As the ABSA approach used here is primarily weakly supervised, it is necessary to define a few keywords to detect topics related to aspects and sentiments. The aspects and sentiment keywords used in this research were taken from the baseline model [3].

2.2. Data Preprocessing

The data preprocessing step is necessary to remove punctuation and other unwanted text from the input. The procedure followed for the pre-processing of the text was similar to the baseline articles for both AE [2] and ABSA [3], and was performed using the NLTK library [19].

In AE, the training data was split into individual sentences. Then, removal of stopwords and punctuation was performed to create a clean file. This pre-processed file was supplied to the word2vec model [20, 21] to produce word embeddings. From the test dataset: food, staff and atmosphere aspects were used for validation purposes, which were also used in the baseline paper [5].

For ABSA, initial train and test data were tokenised and then word2vec embedding model was used to generate word vectors for aspects, sentiments, and joint topics vectors for aspects




Figure 2: Comparison of the network architectures for unsupervised aspect extraction: (a) Baseline architecture [2], (b) modified architecture based on DOConv and (c) modified architecture based on XSepConv. Inside the CMAM Attention Layers, n signifies the size/sequence length of the embeddings. In the cases of (a) and (b), the input embeddings are provided to each of the four sub-layers, to obtain features with different attentions (with different receptive fields), and the outputs are concatenated before supplying to the fully connected layer via dropout regularisation layer. In (c), three sets of sub-blocks were created by combining one of each - $conv2_2$ and $conv1_n$, with one of $conv3_1$, $conv5_1$, and $conv7_1$. The input embeddings were supplied to each of these three to finally obtain four sets of outputs with different receptive fields, which were then concatenated and supplied to the dropout and fully connected layers.

with sentiments.

2.3. Neural Networks

In this research, the state-of-the-art network architecture proposed by He et al. [5] was used. To evaluate the performance of the flavours of convolutions discussed above, the baseline model was modified by incorporating DOConv and XSepConv. This was done by replacing the 2D convolutional layers with 2D DOConv layers and 2D XSepConv layers - which is a combination of 2x2 and 1x200 convolutions. Fig. 2 shows the comparison of these three network architectures.

Weakly supervised ABSA was performed using the state-of-the-art network architecture proposed by Huang et al. [3]. Similar to the modified models for AE, the ABSA model was also

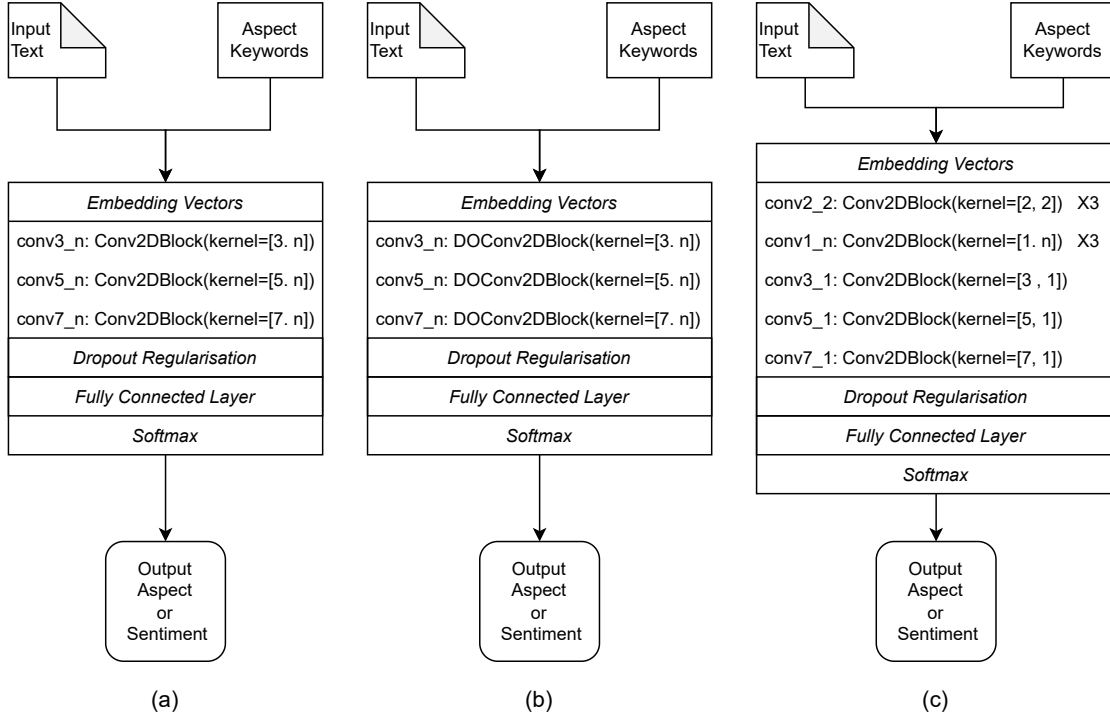


Figure 3: Comparison of the network architectures for weakly supervised aspect-based sentiment analysis: (a) Baseline architecture [3], (b) modified architecture based on DOConv and (c) modified architecture based on XSepConv. Each convolutional block contains a convolution layer (or DO-Conv layer in the case of b), followed by a ReLU activation and 1D maxpool operation. In the cases of (a) and (b), the input embeddings are provided to each of the three sub-layers, to obtain features with different receptive fields, and the outputs are averaged before supplying to the fully connected layer. In (c), three sets of sub-layers were created by combining one of each - *conv2_2* and *conv1_n*, with one of *conv3_1*, *conv5_1*, and *conv7_1*. The input embeddings were supplied to each of these three sets and to *conv1_n* - to finally obtain four sets of outputs, which were then averaged and supplied to the fully connected layer.

modified by replacing the convolutional layers with the DOConv and XSepConv layers. The comparison of the models can be seen in Fig. 3.

2.4. Implementation and Network Training

The method was implemented using PyTorch, and the code is openly available on GitHub². For the unsupervised AE, the method was completely implemented from scratch, while for weakly supervised ABSA, the codebase available from the original authors was modified in this research.

²Code of this research on GitHub: <https://github.com/soumickmj/DeepSentiment>

2.4.1. Unsupervised Aspect Extraction

The goal of this model is to estimate the aspects based on the closest words in the embeddings. The Citysearch dataset was used as a training dataset and the aspect embeddings were initialised according to K-Means centroids. The embedding space was explored to determine the aspects, by matching word embeddings with aspect embedding. The attention mechanism filters out the aspect words based on relevance with the aspect embeddings. The embeddings are then modified on the basis of multiple convolution layers of the neural network model with varying kernel sizes that analyse word embeddings at different levels. Then, sentence embeddings were constructed by combining the word embeddings retained by the attention layer. This dimensionality reduction technique tries to maintain aspect-based embeddings with minimal distortion.

Network training aims to reduce the error caused by sentence reconstruction. For network training, the embedding size, window size, and negative sample size were set to 200, 10, and 5, respectively. The number of aspect embeddings was set to 14 for the restaurant dataset, following Brody et al. [22]. The embedding model used was Word2Vec and, for training the neural network, Triplet Margin [23] was used as the loss function and optimised for 15 epochs with a batch size of 50 using Adam optimiser [24] with a learning rate of 0.001. The maximum length of sentences was set to 20 and the vocabulary size was 9000. The aspects extracted from the trained model were mapped to the gold labels with the help of cosine similarity. The aspect representation words obtained are averaged on the basis of word embeddings. Then, the cosine similarity between the gold labels and the aspect words was computed, and if the similarity value is greater than the threshold of 0.2, the aspects assigned to one of the labels, else, were ignored. Let us consider an example, 'Pizzeria is always a fun place', the aspect extracted for this review is 'Ambience', which is the raw output from the model.

2.4.2. Weakly Supervised Aspect-Based Sentiment Analysis

The author's implementation³ was used to train the baseline model and was modified with XSepConv and DOConv to experiment with the specialised convolutional layers. The model aims to get the prediction of aspects and the corresponding sentiment for the input sentence. To maintain sequential details for aspect-based sentiment classification, the CNN-based classifier is pre-trained on labels, which are assigned based on cosine similarity between embeddings from topic and embeddings from input text documents. Then, self-training on the unstructured text was performed to generalise the labels for aspect and sentiment prediction. The model was trained with an embedding size of 100, word size of 5, for 5 epochs using the SGD optimiser with a learning rate of 0.001 and batch size of 5. Based on the aspect and sentiment keywords, the representative terms obtained from cosine similarity were used to make meaningful words. Let us have a closer look using an example, 'Pizzeria restaurant serves delicious pizza', the model will output 'Food' and 'Positive', which are aspect and sentiment prediction, respectively.

³AE+ABPD: <https://github.com/teapot123/JASen>

2.5. Evaluation

The obtained results were evaluated with the help of precision and the F1 score. To verify statistical significance, the alpha value is chosen using the decision-theoretic method [25], which determines the optimal level of significance for different sample sizes⁴. The comparison is said to be statistically significant if the p-value obtained from the t-test of two related samples⁵ was less than the alpha value. The F1 scores for each label were considered to perform the t tests.

3. Results

3.1. Unsupervised Aspect Extraction

In comparison to the state of the art model, the XSepConv based network excels based on the precision values, while the DOConv architecture performs slightly better compared to the other architectures regarding F1 scores. Table 1 shows the results of the baseline model (scores obtained from the experiments and scores reported in the article [2]) for the Semeval restaurant dataset.

Restaurant Dataset Aspects	Model	Precision	F1 Score
Food	CMAM (Baseline, Implemented)	0.8548	0.8751
	CMAM (Baseline, from the article)	0.887	0.915
	XSepConv	0.9335	0.9092
	DOConv	0.8659	0.9077
Staff	CMAM (Baseline, Implemented)	0.8590	0.6860
	CMAM (Baseline, from the article)	0.804	0.735
	XSepConv	0.8322	0.7631
	DOConv	0.8926	0.7749
Ambience	CMAM (Baseline, Implemented)	0.6227	0.7036
	CMAM (Baseline, from the article)	0.763	0.760
	XSepConv	0.6314	0.7354
	DOConv	0.7737	0.7611
Weighted Average of all aspects	CMAM (Baseline, Implemented)	0.8167	0.8015
	XSepConv	0.8587	0.8454
	DOConv	0.8567	0.8517

Table 1

Unsupervised aspect extraction results for restaurant dataset

For the food aspect, the CMAM results reported in the original article came out as the winner based on the F1 score, however, if the implemented baseline is considered then the XSepConv is the winner. On the basis of the precision, XSepConv is the clear winner for this aspect. For the staff and ambience aspects, DOConv came on top concerning both precision and F1 score. When the weighted average of precision is considered, XSepConv turns out to be the overall

⁴OptSig: <https://rdrr.io/cran/OptSig/>

⁵SciPy ttest rel: https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.ttest_rel.html

winner, but according to the weighted average of F1-score, the DOConv based architecture performs better. To choose the final winning model, the statistical significance of the results was computed. The computed alpha value was 0.3651 for a sample size of 3 (food, staff, and ambience). Based on the p-values obtained from the t-tests which were less than the alpha value, it can be said that both XSepConv and DOConv performed better than the implemented CMAM baseline model with a statistical significance. The improvement observed by the DOConv over XSepConv was statistically insignificant.

3.2. Weakly Supervised Aspect-Based Sentiment Analysis

The results of the baseline model and the modified models were compared for the Semeval restaurant dataset and the Amazon laptop dataset, in terms of aspect terms and its corresponding polarity detection. The results are presented from Tables 2 to 5. It was observed that the XSepConv model failed to identify all aspect labels and detects mostly the majority classes. Therefore, a modified version of XSepConv combined with 2D convolutional layers was used which performed comparatively better.

Restaurant Dataset Aspects	Model	Precision	F1 Score
Location	ABSA	0.1333	0.2
	XSepConv	-	-
	XSepConv + Conv2d	0.1429	0.2105
	DOConv	0.1667	0.2353
Drinks	ABSA	0.4348	0.5634
	XSepConv	-	-
	XSepConv + Conv2d	0.4118	0.5526
	DOConv	0.4348	0.5634
Food	ABSA	0.8757	0.887
	XSepConv	0.5416	0.7026
	XSepConv + Conv2d	0.8567	0.8785
	DOConv	0.9024	0.8796
Ambience	ABSA	0.8125	0.7939
	XSepConv	0.8333	0.137
	XSepConv + Conv2d	0.8393	0.7642
	DOConv	0.8571	0.8308
Service	ABSA	0.9146	0.8219
	XSepConv	0	0
	XSepConv + Conv2d	0.9182	0.8111
	DOConv	0.866	0.8506
Weighted Average of all aspects	ABSA	0.8584	0.839
	XSepConv	0.3774	0.3913
	XSepConv + Conv2d	0.8513	0.8277
	DOConv	0.8624	0.8482

Table 2
Weakly supervised aspect detection for restaurant dataset

Restaurant Dataset Sentiments	Model	Precision	F1 Score
Positive	ABSA	0.7978	0.8548
	XSepConv	0.6989	0.8117
	XSepConv + Conv2d	0.8207	0.852
	DOConv	0.803	0.8578
Negative	ABSA	0.8202	0.6986
	XSepConv	0.8471	0.4431
	XSepConv + Conv2d	0.7788	0.7232
	DOConv	0.8232	0.7078
Weighted Average of all sentiments	ABSA	0.8062	0.7965
	XSepConv	0.7542	0.6741
	XSepConv + Conv2d	0.8051	0.8039
	DOConv	0.8106	0.8018

Table 3

Weakly supervised sentiment classification for restaurant dataset

For most of the aspects detected from the restaurant dataset, including weighted averages, neither XSepConv nor XSepConv + Conv2D managed to outperform the baseline method. However, the DOConv model outperformed the baseline in all, except for the food aspect, including the weighted average of precision and F1 score. However, during polarity detection, both XSepConv + Conv2d and DOConv outperformed the baseline model - based on the weighted average F1 score, XSepConv + Conv2d came out as the winner, if the precision is considered, then DOConv came on top.

For the laptop dataset, a mixed trend can be observed in regard to different aspects. Based on the weighted average of precision and F1 score, DOConv outperformed the baseline model, as well as the XSepConv+Conv2d model. However, for the polarity detection task, XSepConv+Conv2d performed slightly better than the DOConv model.

Statistical significance tests were performed on the values obtained to choose the winner. The computed alpha values for restaurant and laptop datasets were 0.27 and 0.19, respectively, for the aspect detection task and 0.46 for polarity detection, which were then used to judge the statistical significance of the results. For aspect detection, DOConv achieved statistically significant improvement over the baseline method, as well as over the XSepConv+Conv2d - for both datasets. However, for polarity detection, DOConv achieved a statistically significant improvement over the baseline. The comparison of XSepConv+Conv2d and DOConv was statistically insignificant for both datasets. Therefore, the DOConv model can be chosen as the winner for weakly supervised sentiment analysis.

4. Conclusion

This paper studies the applicability of specialised convolution layers, namely XSepConv and DOConv, in the field of sentiment analysis. These layers were originally proposed for the processing of image data and had not yet been applied for text processing. Evaluations in the context of unsupervised aspect extraction and weakly supervised aspect-based sentiment

Laptop Dataset Aspects	Model	Precision	F1 Score
Support	ABSA	0.8333	0.4348
	XSepConv	0	0
	XSepConv + Conv2d	0.625	0.4
	DOConv	0.75	0.48
OS	ABSA	0.6667	0.642
	XSepConv	0	0
	XSepConv + Conv2d	0.6429	0.6429
	DOConv	0.7073	0.6410
Display	ABSA	0.878	0.72
	XSepConv	0.5416	0.7026
	XSepConv + Conv2d	0.881	0.7327
	DOConv	0.875	0.7071
Battery	ABSA	0.6552	0.7917
	XSepConv	0.5926	0.4923
	XSepConv + Conv2d	0.7255	0.8315
	DOConv	0.74	0.8409
Company	ABSA	0.5574	0.6296
	XSepConv	0.1655	0.2831
	XSepConv + Conv2d	0.5614	0.6154
	DOConv	0.614	0.6731
Mouse	ABSA	0.7317	0.7895
	XSepConv	0	0
	XSepConv + Conv2d	0.6818	0.7595
	DOConv	0.7143	0.7792
Software	ABSA	0.56	0.6222
	XSepConv	0	0
	XSepConv + Conv2d	0.5385	0.6087
	DOConv	0.5143	0.6545
Keyboard	ABSA	0.7667	0.7419
	XSepConv	0	0
	XSepConv + Conv2d	0.8276	0.7869
	DOConv	0.7742	0.7619
Weighted Average of all aspects	ABSA	0.7185	0.6766
	XSepConv	0.1948	0.1106
	XSepConv + Conv2d	0.7013	0.6784
	DOConv	0.7275	0.6948

Table 4
Weakly supervised aspect detection for laptop dataset

analysis (integrated aspect extraction and polarity detection) have been presented here. For the unsupervised aspect extraction task, both layers outperformed the baseline model with the conventional convolutional layer with statistical significance. DOConv achieved a better F1 score than XSepConv, but the difference was statistically insignificant. For the weakly supervised aspect extraction, however, the XSepConv performed poorly, and a combination of XSepConv and the conventional convolutional layer performed better, but failed to outperform

Laptop Dataset Sentiments	Model	Precision	F1 Score
Positive	ABSA	0.6981	0.7184
	XSepConv	0.7248	0.61
	XSepConv + Conv2d	0.6977	0.7453
	DOConv	0.7081	0.7331
Negative	ABSA	0.7365	0.7148
	XSepConv	0.6414	0.7155
	XSepConv + Conv2d	0.7778	0.7192
	DOConv	0.7538	0.7261
Weighted Average of all sentiments	ABSA	0.7177	0.7166
	XSepConv	0.6821	0.664
	XSepConv + Conv2d	0.7386	0.732
	DOConv	0.7313	0.7295

Table 5
Weakly supervised sentiment classification for laptop dataset

the baseline model with the conventional convolutional layers on the restaurant dataset and only outperformed without any statistical significance on the laptop dataset. The model with DOConv on the other hand outperformed the baseline model with statistical significance on both datasets. When it comes to weakly supervised aspect-based polarity detection, DOConv achieved improvements over the baseline model, but was statistically insignificant. The model that used a combination of XSepConv and conventional convolutional layers achieved a statistically significant improvement over the baseline model.

Overall, DOConv performed the best among the three types of convolutions used in this work. Moreover, the experiments show that both of these specialised convolutional layers have great potential when it comes to textual data and should be further investigated. Research papers have shown that increasing the number of keywords for the weakly supervised aspect-based sentiment analysis can improve the performance of the model, which was not evaluated in the context of this research. The aspects extracted using the unsupervised aspect extraction model can be used in the weakly supervised ABSA model, and will be examined in future work. Furthermore, increasing the depth of the network and synonym-based data augmentation techniques might improve the performance. Both will also be evaluated in the future.

References

- [1] E. Bassignana, D. Brunato, M. Polignano, A. Ramponi, Preface to the Seventh Workshop on Natural Language for Artificial Intelligence (NL4AI), in: Proceedings of the Seventh Workshop on Natural Language for Artificial Intelligence (NL4AI 2023) co-located with 22th International Conference of the Italian Association for Artificial Intelligence (AI* IA 2023), 2023.
- [2] T. Sokhin, M. Khodorchenko, N. Butakov, Unsupervised neural aspect extraction with related terms, in: Conference on Artificial Intelligence and Natural Language, Springer, 2020, pp. 75–86.

- [3] J. Huang, Y. Meng, F. Guo, H. Ji, J. Han, Weakly-supervised aspect-based sentiment analysis via joint aspect-sentiment topic embedding, arXiv preprint arXiv:2010.06705 (2020).
- [4] B. Liu, M. Hu, J. Cheng, Opinion observer: analyzing and comparing opinions on the web, in: Proceedings of the 14th international conference on World Wide Web, 2005, pp. 342–351.
- [5] R. he, W. Lee, H. Ng, D. Dahlmeier, An unsupervised neural attention model for aspect extraction, 2017, pp. 388–397. doi:10.18653/v1/P17-1036.
- [6] D. Tang, F. Wei, N. Yang, M. Zhou, T. Liu, B. Qin, Learning sentiment-specific word embedding for twitter sentiment classification, in: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2014, pp. 1555–1565.
- [7] S. Chatterjee, P. G. Jose, D. Datta, Text classification using svm enhanced by multithreading and cuda, International Journal of Modern Education and Computer Science 12 (2019) 11.
- [8] N. Zainuddin, A. Selamat, R. Ibrahim, Hybrid sentiment classification on twitter aspect-based sentiment analysis, Applied Intelligence 48 (2018) 1218–1232.
- [9] H. Wan, Y. Yang, J. Du, Y. Liu, K. Qi, J. Z. Pan, Target-aspect-sentiment joint detection for aspect-based sentiment analysis, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 34, 2020, pp. 9122–9129.
- [10] M. Ruffolo, F. Visalli, A weak-supervision method for automating training set creation in multi-domain aspect sentiment classification., in: ICAART (2), 2020, pp. 249–256.
- [11] C. Wu, Q. Xiong, H. Yi, Y. Yu, Q. Zhu, M. Gao, J. Chen, Multiple-element joint detection for aspect-based sentiment analysis, Knowledge-Based Systems 223 (2021) 107073.
- [12] X. Wang, G. Xu, Z. Zhang, L. Jin, X. Sun, End-to-end aspect-based sentiment analysis with hierarchical multi-task learning, Neurocomputing 455 (2021) 178–188.
- [13] J. Chen, Z. Lu, J.-H. Xue, Q. Liao, Xsepconv: extremely separated convolution, arXiv preprint arXiv:2002.12046 (2020).
- [14] J. Cao, Y. Li, M. Sun, Y. Chen, D. Lischinski, D. Cohen-Or, B. Chen, C. Tu, Do-conv: Depthwise over-parameterized convolutional layer, arXiv preprint arXiv:2006.12030 (2020).
- [15] G. Ganu, N. Elhadad, A. Marian, Beyond the stars: improving rating predictions using review text content., in: WebDB, volume 9, Citeseer, 2009, pp. 1–6.
- [16] M. Pontiki, D. Galanis, H. Papageorgiou, I. Androutsopoulos, S. Manandhar, M. Al-Smadi, M. Al-Ayyoub, Y. Zhao, B. Qin, O. De Clercq, et al., Semeval-2016 task 5: Aspect based sentiment analysis, in: International workshop on semantic evaluation, 2016, pp. 19–30.
- [17] R. He, J. McAuley, Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering (2016) 507–517.
- [18] M. Pontiki, D. Galanis, H. Papageorgiou, S. Manandhar, I. Androutsopoulos, Semeval-2015 task 12: Aspect based sentiment analysis, in: Proceedings of the 9th international workshop on semantic evaluation (SemEval 2015), 2015, pp. 486–495.
- [19] S. Bird, E. Klein, E. Loper, Natural language processing with Python: analyzing text with the natural language toolkit, ” O’Reilly Media, Inc.”, 2009.
- [20] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, in: Y. Bengio, Y. LeCun (Eds.), 1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings, 2013. URL: <http://arxiv.org/abs/1301.3781>.

- [21] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, *Advances in neural information processing systems* 26 (2013).
- [22] S. Brody, N. Elhadad, An unsupervised aspect-sentiment model for online reviews, in: *Human language technologies: The 2010 annual conference of the North American chapter of the association for computational linguistics*, 2010, pp. 804–812.
- [23] V. Balntas, E. Riba, D. Ponsa, K. Mikolajczyk, Learning local feature descriptors with triplets and shallow convolutional neural networks., in: *Bmvc*, volume 1, 2016, p. 3.
- [24] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, *arXiv preprint arXiv:1412.6980* (2014).
- [25] J. H. Kim, I. Choi, Choosing the level of significance: A decision-theoretic approach, *Abacus* 57 (2021) 27–71.