

Enhanced Coherence-Aware Network with Hierarchical Disentanglement for Aspect-Category Sentiment Analysis

Jin Cui¹, Fumiyo Fukumoto^{1*}, Xinfeng Wang¹, Yoshimi Suzuki¹,
Jiyi Li¹, Noriko Tomuro², and Wanzeng Kong³

¹ University of Yamanashi, ² DePaul University, ³ Hangzhou Dianzi University
{g22dtsa5,fukumoto,g22dtsa7,ysuzuki,jyli}@yamanashi.ac.jp
tomuro@cs.depaul.edu, and kongwanzeng@hdu.edu.cn

Abstract

Aspect-category-based sentiment analysis (ACSA), which aims to identify aspect categories and predict their sentiments has been intensively studied due to its wide range of NLP applications. Most approaches mainly utilize intrasentential features. However, a review often includes multiple different aspect categories, and some of them do not explicitly appear in the review. Even in a sentence, there is more than one aspect category with its sentiments, and they are entangled intra-sentence, which makes the model fail to discriminately preserve all sentiment characteristics. In this paper, we propose an enhanced coherence-aware network with hierarchical disentanglement (ECAN) for ACSA tasks. Specifically, we explore coherence modeling to capture the contexts across the whole review and to help the implicit aspect and sentiment identification. To address the issue of multiple aspect categories and sentiment entanglement, we propose a hierarchical disentanglement module to extract distinct categories and sentiment features. Extensive experimental and visualization results show that our ECAN effectively decouples multiple categories and sentiments entangled in the coherence representations and achieves state-of-the-art (SOTA) performance. Our codes and data are available online: <https://github.com/cuijin-23/ECAN>.

Keywords: aspect category detection, sentiment analysis, coherence, hierarchical disentanglement

1. Introduction

With the rapid growth of Internet services, a large number of user-generated textual reviews have become available, which has drawn much attention to aspect-based sentiment analysis (ABSA) research (Zhang et al., 2022a). To date, many of its variants, including aspect-based sentiment classification (ABSC) and aspect term-based sentiment analysis (ATSA), have been studied (Zhang and Qian, 2020; Li et al., 2021; Zhang et al., 2022c; Ma et al., 2023). One such attempt is aspect-category-based sentiment analysis (ACSA) consisting of two sub-tasks: aspect category detection (ACD), which identifies aspect categories, and aspect-category sentiment classification (ACSC), which predicts their sentiments (Cai et al., 2020; Li et al., 2020c; Zhou and Law, 2022).

The primary approaches to this task mainly leverage intrasentential features to learn their models (Zhou and Law, 2022). However, two major issues remain in the way of state-of-the-art (SOTA) performance. One is that a review often includes multiple different aspect categories, and some of them do not explicitly appear in the review. Taking an example from the REST 15 dataset shown in Figure 1, there are several aspect categories, such as “service” and “restaurant.” “Raymond” in S_3 may be a bartender’s name, while we cannot identify that it is related to the “service” category with only S_3 , since “service” in S_3 shows an im-

*Corresponding author.

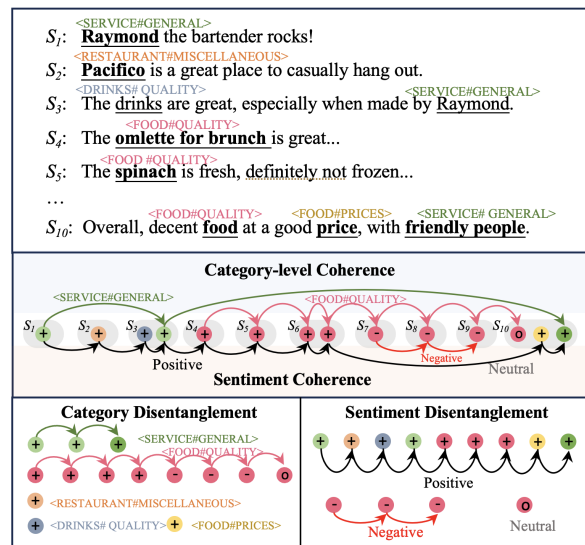


Figure 1: A review (ID: P#9) from SemEval-2015: Words with angle brackets and the colored circle show aspect categories. “+,” “-,” and “o” denote positive, negative, and neutral sentiments, respectively. Arcs indicate that they have an identical category or sentiment in a coherent review.

licit aspect category. A user often gives reviews with a consistent opinion. For example, “Raymond” in S_3 is strongly related to both “Raymond” in S_1 and “friendly people” in S_{10} which indicates “service” served by people. However, the aforementioned approaches cannot capture how sentences

are connected or how the entire review is organized to convey information to the reader.

Another issue is that there is more than one aspect category and its sentiment in a sentence, and they are entangled. A vanilla sentiment-level token fails to discriminately preserve all sentiment characteristics. For instance, S_{10} in Figure 1 contains a neutral opinion on the aspect of food, i.e., “decent food” and a positive opinion on the restaurant’s price and service, i.e., “good” and “friendly.” Therefore, “food” is likely to be incorrectly identified as positive. As illustrated in the middle box in Figure 1, i.e., category-level coherence, this indicates that both aspect categories and their sentiments are entangled intra-sentence as well as throughout sentences in a review.

Motivated by the issues mentioned above, we propose an enhanced coherence-aware network with hierarchical disentanglement (ECAN) for the ACSA task. Specifically, we leverage coherence modeling to capture contexts across the whole review. It also enables the model to learn explicit opinions from contexts to help the implicit aspect and sentiment identification. Furthermore, to address the issue that aspect categories and sentiments are entangled, we propose a hierarchical disentanglement module. The module applies a parallel attention mechanism on coherence-aware representation to both multiple categories and aspects to extract distinct categories and sentiment features. As illustrated in the bottom boxes of Figure 1, we disentangle (i) the category representations, and (ii) the words that represent sentiments according to their sentiment polarities by utilizing the coherence features. The sentence-level representation is thus divided into several groups with relative category and sentiment features.

In summary, the main contributions of our approach are as follows:

- (1) We propose an ECAN that leverages coherence modeling to learn the explicit opinions from contexts to help the implicit aspect and sentiment identification for the ACSA task;
- (2) We propose a hierarchical disentanglement of category and sentiment for mining more fine-grained features;
- (3) Extensive experiments and visualization results on four benchmark datasets have demonstrated that the ECAN outperforms SOTA methods in both ACD and ACSC tasks.

2. Related Work

Aspect Category Sentiment Analysis. Early studies on ACSA first detected the aspect category and then predicted the sentiment polarities for the detected categories (Xue and Li, 2018; Tay et al., 2018). Liang et al. (2019) utilized a sparse coeffi-

cient vector to guide the representations of aspect category and sentiment. As such, the representations of aspect category and sentiment are learned independently, which hampers the performance of the model. Several studies focus on the issue and explore the joint prediction of two associated elements (Wang et al., 2019; Li et al., 2020a; Wu et al., 2021; Lin et al., 2023), e.g., a multi-task learning approach (Li et al., 2020c) and a hierarchical graph convolutional network (Cai et al., 2020).

More recently, some attempts have been made to build sentiment triples or quadruples to simplify aspect-category tasks (Cai et al., 2021; Gao et al., 2022; Gou et al., 2023). However, most of these approaches ignore the effectiveness of disentangling sentence representations by categories and sentiments in sentences.

Coherence Modeling. With the success of deep learning techniques, coherence analysis has been widely studied. Several approaches are designed to model sentence-level coherence, including a transferable neural model (Xu et al., 2019), an entity-based neural local coherence model (Jeon and Strube, 2022), and a local attention mechanism-based hierarchical coherence model (Liao et al., 2021). Another approach is to learn coherence across a whole document, which includes a transformer-based architecture (Abhishek et al., 2021) and a multi-task framework that learns both word- and document-level representations (Frag and Yannakoudakis, 2019). Our approach lies across both sentence- and document-level coherence to provide a comprehension framework for sentiment coherence.

Disentangled Representation Learning. Disentangled representation learning (DRL) has been widely employed to decouple underlying factors in the observable data (Bao et al., 2019; Wang et al., 2022), as it is intractable to model explanatory factors in some scenarios (Locatello et al., 2019). Several DRL methods have been applied to NLP tasks, such as text-style transfer (Cheng et al., 2020; John et al., 2019; Nangi et al., 2021; Hu et al., 2023), multi-modal sentiment analysis (Zhang et al., 2022b), domain adaptation (Liu et al., 2018; Lee et al., 2021; Wang et al., 2023), and text generation (Chen et al., 2019; Thompson and Post, 2020).

Recent work on aspect-level sentiment analysis has primarily focused on the ATSA task. Silva and Marcacini (2021) attempted to disentangle the syntactic and semantic features by utilizing the DeBERTa model (He et al., 2021). Mei et al. (2023) proposed a graph-based model that extracts a specific linguistic property to help capture finer feature representations. However, to date, little research has focused on aspect categories and sentiment disentanglement for the ACSA task.




Figure 2: The main framework of our proposed method. It consists of four modules: (1) coherence-aware representation learning, (2) hierarchical disentanglement, (3) word-level syntactic learning, and (4) multi-task learning.

3. Task Definition

Given a review D with I sentences, $D = \{s_i\}_{i=1}^I$, where the i -th sentence $s_i = \{w_j\}_{j=1}^n$ consists of the number of n words. Let also $C = \{c_i\}_{i=1}^m$ be a set of m pre-defined aspect categories, and $p = \{positive, neutral, negative\}$ be a set of sentiment polarity labels in C . The goal of the ACSA task is to detect all the aspect categories appearing in s_i and classify the sentiment polarity for each detected category.

4. Approach

We present an ECAN model consisting of four components, (1) coherence-aware representation learning with XLNet (Yang et al., 2019), (2) hierarchical disentanglement, (3) word-level syntactic learning for enhancing sentiment contexts, and (4) multi-task learning. The overall architecture is illustrated in Figure 2.

4.1. Coherence-Aware Representation Learning with XLNet

Unlike previous works (Yang and Li, 2021; Zhang et al., 2019, 2022c) on ABSA which focus on local contexts within a sentence, ECAN mines both sentence- and document-level contexts from re-

views to detect aspect categories and their sentiment polarities. To mine document-level contexts, we utilize the pre-training model XLNet to model the coherence representation. We choose the XLNet as the backbone, due to its capacity to learn longer text sequences, which has been widely used to model the coherence representation (Cui et al., 2023; Jwalapuram et al., 2022; Jeon and Strube, 2020).

Specifically, for each document, we define an input sequence, s_1 [SEP] ... s_I [SEP] [CLS], padded with two special tokens, [SEP] and [CLS], at the end of input s_i . Here, [SEP] and [CLS] are the same tokens as those of BERT (Devlin et al., 2018). Then, we obtain the review embedding $e_d \in \mathbb{R}^{I \times d_m}$ via XLNet. Likewise, for each sentence, we create an input, s_i [SEP][CLS]; the word-level embedding is denoted as $e_{w_j} \in \mathbb{R}^{d_m}$, where d_m denotes the embedding size.

To learn robust coherence representations, we employ the sentence ordering contrastive learning (CL) task (Jwalapuram et al., 2022) as an auxiliary task. It enforces that the coherence score of the positive sample (original document) should be higher than that of the negative sample (disorder document). We randomly shuffle the sentences within the original review to generate the number of B negative samples and apply contrastive learning to align the coherent and incoherent representa-

tions. Let $f_\theta(\mathbf{e}_d)$ be a linear projection to convert coherent document embedding \mathbf{e}_d into coherence scores. The margin-based contrastive loss is given as follows:

$$\mathcal{L}_{cl} = - \sum_{d^+ \in \mathcal{D}} \log \left(\frac{\exp(f_\theta(\mathbf{e}_{d^+}))}{\exp(f_\theta(\mathbf{e}_{d^+})) + \sum_{j=1}^B \exp(f_\theta(\mathbf{e}_{d_j^-}) - \tau)} \right), \quad (1)$$

where d^+ refers to a positive sample, $f_\theta(\mathbf{e}_{d^+})$ indicates the score of the positive sample, $f_\theta(\mathbf{e}_{d_1^-}), \dots, f_\theta(\mathbf{e}_{d_B^-})$ denote the scores of B negative samples, and τ refers to the margin.

4.2. Hierarchical Disentanglement

Although coherence modeling with vanilla XLNet can capture contextual semantics, it fails to discriminate contextual features of different aspect categories and aspect sentiments within a sentence, as the coherence model captures global relationships among sentences within a whole document, which hinders the performance of more fine-grained aspect-based sentiment tasks, i.e., ACD and ACSC tasks. To address the issue, we simultaneously disentangle the categories and their sentiment representations within each sentence embedded by hierarchical document-level coherence.

4.2.1. Category Disentanglement

Inspired by the work of [Medina and Kalita \(2018\)](#), we adopt a parallel attention mechanism to obtain disentangled representation blocks indicating fine-grained category components. Formally, for the i -th attention channel ($i \in \{1, 2, \dots, d_c\}$), given the document-level representation \mathbf{e}_d , we employ a slicing operation to obtain the representation \mathbf{e}_s of the sentence s in \mathbf{e}_d as the $\mathbf{e}_i^{(0)}$ at the first layer, and then we stack transformer blocks to obtain the self-attention values in each channel:

$$\mathbf{e}_i^{(l+1)} = \text{softmax} \left(\sum_{j \in \mathcal{S}_i} \frac{\mathbf{e}_i^{(l)} \mathbf{W}_{*,i}^{q_1} (\mathbf{e}_j^{(l)} \mathbf{W}_{*,j}^{k_1})^\top}{\sqrt{d_k}} \right) \mathbf{e}_i^{(l)} \mathbf{W}_{*,i}^{v_1}, \quad (2)$$

where $\mathbf{W}^{q_1} \in \mathbb{R}^{d_m \times d_k}$, $\mathbf{W}^{k_1} \in \mathbb{R}^{d_m \times d_k}$ and $\mathbf{W}^{v_1} \in \mathbb{R}^{d_m \times d_k}$ are three weight matrices, and d_k is the dimension size. Subsequently, $\mathbf{e}_i^{(l)}$ at the last layer is fed into the feedforward network as follows:

$$\mathbf{v}_i^{(l)} = \text{ReLU}(\mathbf{e}_i^{(l)} \mathbf{W}_{*,i}^{FC_1} + \mathbf{b}_1) \mathbf{W}_{*,i}^{FC_2} + \mathbf{b}_2, \quad (3)$$

where $\mathbf{W}^{FC_1} \in \mathbb{R}^{d_m \times d_k}$, $\mathbf{W}^{FC_2} \in \mathbb{R}^{d_m \times d_k}$ are two linear transformation matrices, \mathbf{b}_1 and \mathbf{b}_2 are learnable parameters. After the number of d_c parallel channels, we obtain the representations of d_c disentangled blocks, i.e., $\mathbf{v}_1^{(l)}, \mathbf{v}_2^{(l)}, \dots, \mathbf{v}_{d_c}^{(l)}$. These disentangled representations are concatenated to further

select representative features using the attention pooling layer by the following equation:

$$M_j = \tanh(\mathbf{W}_M^\top [\mathbf{v}_1^{(l)}; \mathbf{v}_2^{(l)}; \dots; \mathbf{v}_{d_c}^{(l)}] + \mathbf{b}_M), \quad \alpha_M = \sum_{j \in d_j} \text{softmax}(\mathbf{W}_j^\top M_j), \quad (4)$$

where $\mathbf{W}_M \in \mathbb{R}^{d_c \cdot d_m \times d_m}$, $\mathbf{W}_j \in \mathbb{R}^{d_m \times d_m}$ and $\mathbf{b}_M \in \mathbb{R}^{d_c \cdot d_m}$ are learnable parameters, and α_M is the attention weight vector. The final category representation r_c is denoted as:

$$r_c = [\mathbf{v}_1^{(l)}; \mathbf{v}_2^{(l)}; \dots; \mathbf{v}_{d_c}^{(l)}] \alpha_M. \quad (5)$$

We note that the category representation r_c is partially optimized by the ACD task, whose gradients are based on ground-truth category information.

4.2.2. Sentiment Disentanglement

Similar to category disentanglement, we obtain the disentangled sentiment representation blocks from the document-level coherence representation \mathbf{e}_d via a parallel attention mechanism, where the number of parallel channels is defined as d_s . The difference is that the disentangled sentiment representations are optimized via the ACSC task only, containing ground-truth sentiment information. We note that the gradient update of disentangled sentiment representations is completely trained by the ACSC task; although, whether it will be trained or not is activated by the aspect category information.

Specifically, in the i -th channel, the self-attention results are computed layer by layer as follows:

$$\mathbf{u}_i^{(l+1)} = \text{softmax} \left(\sum_{j \in \mathcal{S}_i} \frac{\mathbf{u}_i^{(l)} \mathbf{W}_{*,i}^{q_2} (\mathbf{u}_j^{(l)} \mathbf{W}_{*,j}^{k_2})^\top}{\sqrt{d_k}} \right) \mathbf{u}_i^{(l)} \mathbf{W}_{*,i}^{v_2}, \quad (6)$$

where \mathbf{W}^{q_2} , \mathbf{W}^{k_2} and \mathbf{W}^{v_2} are three weight matrices. Empirically, we found that a simple linear transformation works better than an attention-pooling operation.

$$\mathbf{U} = [\mathbf{u}_1^{(l)}, \mathbf{u}_2^{(l)}, \dots, \mathbf{u}_{d_s}^{(l)}] \mathbf{W}_U. \quad (7)$$

Thereafter, the output \mathbf{U} is then fed into a position-wise feed-forward neural network, generating hidden disentangled representations \mathbf{u}_s of the sentiment representation:

$$\mathbf{u}_s = \text{ReLU}(\mathbf{U} \mathbf{W}^{FC_1} + \mathbf{b}_1) \mathbf{W}^{FC_2} + \mathbf{b}_2, \quad (8)$$

where \mathbf{W}^{FC_1} , \mathbf{W}^{FC_2} are two weight matrices.

4.3. Word-Level Syntactic Learning for Enhancing Sentiment Semantics

To capture the local contexts within a sentence, we use the Stanford parser¹ to obtain a word-level dependency tree of the target sentence and apply graph convolution operations to learn the local sentiment correlations between sentiment and aspect categories from the dependency tree. The syntactic adjacency matrix in the dependency tree of the sentence is defined by $\mathbf{A} \in \mathbb{R}^{n \times n}$, and the syntactic embedding nodes are denoted by $\mathbf{g}^{(0)} = [\mathbf{e}_{w_1}, \mathbf{e}_{w_2}, \dots, \mathbf{e}_{w_n}]$. Then, the normalized adjacency matrix can be obtained by $\hat{\mathbf{A}} = \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$, where $\mathbf{D} \in \mathbb{R}^{N \times N}$ is a diagonal matrix in which each entry, $d_{ij} = \sum_{j=1}^n \mathbf{a}_{ij}$ denotes the number of nonzero entries in the i -th row vector of the adjacency matrix \mathbf{A} . The node representations are updated as follows:

$$\begin{aligned} \tilde{\mathbf{g}}_i^{(l)} &= \sum_{j=1}^n \hat{\mathbf{a}}_{ij} \mathbf{W}^{(l)} \mathbf{g}_j^{(l-1)}, \\ \mathbf{g}_i^{(l)} &= \text{ReLU}(\tilde{\mathbf{g}}_i^{(l)} + \mathbf{b}^{(l)}), \end{aligned} \quad (9)$$

where $\mathbf{g}_i^{(l-1)} \in \mathbb{R}^{d_m}$ denotes the i -th word representation obtained from the graph convolution network (GCN) layer, and $\mathbf{g}_i^{(l)}$ refers to the i -th word representation of the current GCN layer. The weights $\mathbf{W}^{(l)}$ and bias $\mathbf{b}^{(l)}$ are learnable parameters. We feed the obtained local sentiment representation $\mathbf{g}_i^{(l)}$ into a linear layer, which can be denoted as $\hat{\mathbf{g}} \in \mathbb{R}^{d_m}$:

$$\hat{\mathbf{g}} = [\mathbf{g}_1^{(l)}, \dots, \mathbf{g}_n^{(l)}] \times \mathbf{W}_o^T + \mathbf{b}_o, \quad (10)$$

where W_o and b_o are the learnable weight and bias, respectively.

4.4. Multi-Task Learning

We use a multi-task framework to jointly learn to identify document-level coherence with sentence ordering contrastive learning, detect the underlying aspect categories, and identify their sentiments.

Aspect-category detection. The final category representation is obtained by r_c ; thus the probability of the i -th category $p_j^c \in \mathbb{R}^m$ is given by:

$$p_j^c = p(y_j^c | r_j) = \text{sigmoid}(\mathbf{W}_j^c r_c + \mathbf{b}_j^c). \quad (11)$$

The loss function of ACD, i.e. the binary cross entropy loss function of ACD, is given by:

$$\mathcal{L}_{ACD} = - \sum_{i=1}^m y_i^c \log p_i^c + (1 - y_i^c) \log(1 - p_i^c). \quad (12)$$

¹<https://stanfordnlp.github.io/CoreNLP/>

	REST 15	REST 16	LAP 15	LAP 16
# Train sentences	1102	1680	1397	2037
# Test sentences	572	580	644	572
# Categories	6	6	22	22
# Positive	1458	1966	1641	2113
# Neutral	91	137	185	234
# Negative	663	848	1091	1353

Table 1: Statistics of four benchmark datasets, REST and LAP in the SemEval-2015 and 2016.

Aspect-category sentiment classification. The sentiment representations are obtained by $r_s = [\hat{\mathbf{g}}, \mathbf{u}_s]$. Thus, the sentiment probability $p_j^s \in \mathbb{R}^p$ corresponding to its category is given by:

$$p_j^s = p(y_j^s | y_j^c, r_j) = \text{softmax}(\mathbf{W}_j^s r_s + \mathbf{b}^s). \quad (13)$$

Following Cai et al. (2020), a hierarchical prediction strategy is utilized to obtain the final sentiments prediction on the i -th category. We utilize cross-entropy loss as the objective loss function of ACSC which is defined as follows:

$$\mathcal{L}_{ACSC} = - \sum_{i=1}^m \sum_{j=1}^p \mathbb{I}(y_{i,j}^s) \log p_{i,j}^s, \quad (14)$$

where $\mathbb{I}(\cdot)$ refers to an indicator function to adjust the output of the hierarchical prediction. The final loss is given by:

$$\mathcal{L}_{total} = \delta_1 \mathcal{L}_{cl} + \delta_2 \mathcal{L}_{ACD} + \delta_3 \mathcal{L}_{ACSC}, \quad (15)$$

where $\delta_1, \delta_2, \delta_3 \in [0, 1]$ are hyperparameters used to balance the three tasks, CL, ACD, and ACSC.

5. Experiments

5.1. Datasets and Evaluation Metrics

We evaluate our model on four benchmark datasets: REST 15 and LAP 15 from the SemEval-2015 task 12 (Pontiki et al., 2015), and REST 16 and LAP 16 from the SemEval-2016 task 5 (Pontiki et al., 2016). We choose these four datasets because the input of our ECAN is based not on a sentence unit but on a textual review to leverage document-level coherence. Each dataset consists of restaurant and laptop domains, and positive, neutral, and negative sentiment polarities. The statistics of datasets are displayed in Table 1. We use precision (P), recall (R), and macro-averaged F1 scores (F1) as metrics for evaluation.

5.2. Baselines

To examine the efficacy of ECAN in ACD and ACSC tasks, we compared it with the following nine SOTA baselines which are classified into four groups:

	REST 15			REST 16			LAP 15			LAP 16		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
CAER-BERT	85.19	76.42	79.02	84.81	80.08	82.14	83.77	62.86	72.20	81.09	64.91	72.67
SCAN	86.23	81.15	83.62	85.01	85.01	84.45	80.34	71.63	74.82	81.39	62.01	71.12
AC-MIMLLN	<u>88.64</u>	80.76	84.17	86.05	83.16	82.59	81.21	69.36	73.21	82.09	67.08	74.22
LC-BERT	87.62	79.28	82.92	89.11	<u>85.71</u>	85.40	85.86	65.25	74.15	82.57	63.17	70.85
EDU-Capsule	84.63	78.42	81.40	89.91	84.21	84.62	92.47	60.34	73.03	81.65	63.32	72.07
Hier-BERT	86.59	<u>80.86</u>	83.59	90.53	85.23	85.66	88.27	60.88	72.06	82.68	66.38	73.19
Hier-GCN-BERT	88.18	79.28	84.12	<u>91.16</u>	85.29	86.54	<u>92.53</u>	75.10	<u>86.81</u>	<u>82.88</u>	74.42	77.84
Hier-GCN-XInet	88.19	77.27	84.39	90.44	85.41	<u>87.26</u>	91.78	<u>75.29</u>	86.09	82.16	<u>75.99</u>	<u>78.52</u>
Ours(ECAN)	91.62	82.39	85.67	91.69	86.02	88.75	93.64	78.43	88.80	83.45	76.23	79.21

Table 2: The main results of the ACD task. Bold font and underline indicate the best, and the second best result, respectively.

- *Single task learning-based:*

- **CAER-BERT** (Liang et al., 2019) adopts a sparse coefficient vector to select highly correlated words from the sentences to adjust the representations of the aspect categories and sentiments for ACSA.

- *Multi-task learning-based:*

- **SCAN** (Li et al., 2020b) employs graph attention networks to aggregate high-order representations of the nodes in sentence constituency parse trees.
- **AC-MIMLLN** (Li et al., 2020c) utilizes a multi-instance multi-label learning network to obtain the sentiments of the sentence toward the aspect categories by aggregating the key instance sentiments.
- **LC-BERT** (Wu et al., 2021) proposes a two-stage strategy that first locates the aspect term and then takes it as the bridge to find the related sentiment words.
- **EDU-Capsule** (Lin et al., 2023) learns elementary discourse unit representations by capsule network within its sentential context.

- *Hierarchical learning-based:*

- **Hier-BERT** (Cai et al., 2020) introduces the hierarchy method for ACD and ACSC tasks with BERT as the sentence encoder.
- **Hier-GCN-BERT** (Cai et al., 2020) utilizes the GCN sub-layer to model the inner-relations between category and inter-relations between category and sentiment based on Hier-BERT.
- **Hier-GCN-XInet** replaces the BERT encoder with a XLNet model in Hier-GCN-BERT model for comparison.

- *Generative approach-based:*

- **MvP** (Gou et al., 2023) introduces element order prompts to guide the language model to generate multiple sentiment tuples.

5.3. Implementation Details

Following Cai et al. (2020), we randomly chose 10% of the training data and used it as the development data. The optimal hyper-parameters were as follows: The initial learning rate for coherence-aware representation learning was $8e-6$, and the others were $2e-5$. The weight decay was set to $1e-3$, and the dropout rate was 0.1. The number of negative samples B was 5, and the margin τ was 0.1. The balance coefficients δ_1 , δ_2 , and δ_3 were set to 0.1, 0.5, and 0.5, respectively. The number of graph convolutional layers was 3. The number of disentangled blocks was set to 4 for both category disentanglement d_c and sentiment disentanglement d_s . All hyperparameters were tuned using Optuna². We used AdamW (Loshchilov and Hutter, 2017) as the optimizer.

5.4. Main Results

The experimental results on the four datasets show that the ECAN exceeds all baselines in both ACD and ACSC tasks, indicating the effectiveness of disentanglement of category and sentiment for coherence-aware representations. Specifically, it surpasses baselines by a large margin in the ACSC task, while being slightly better than the second-best methods in the ACD task.

5.4.1. Performance of Aspect Category Detection

Table 2 shows the results of aspect category detection. Overall, the ECAN attained an improvement over the second-best methods by 0.6~3.4% in precision, 0.3~4.2% in recall, and 0.9~2.3% in the F1-score. Table 2 also prompts the following observations and insights:

- Hierarchical learning, such as Hier-GCN-BERT and Hier-GCN-XLNet, are competitive among baselines, while the single task learning-based

²<https://github.com/pfnet/optuna>. The search ranges are reported in the appendix A.1

Methods	REST 15			REST 16			LAP 15			LAP 16		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
CAER-BERT	71.27	63.75	66.25	77.69	68.39	73.60	58.34	51.86	60.78	65.36	54.73	60.79
SCAN	71.25	67.05	69.09	73.87	76.25	75.05	71.01	45.49	55.45	61.06	45.42	52.08
AC-MIMLLN	71.45	65.18	68.17	75.66	74.89	75.27	59.06	63.78	55.29	68.08	55.89	61.39
LC-BERT	72.08	65.76	68.77	76.34	74.16	75.75	66.84	50.80	57.72	71.05	53.11	55.58
EDU-capsule	72.67	67.34	69.90	75.93	77.00	74.81	73.98	52.27	65.43	67.02	49.20	58.01
Hier-BERT	70.42	65.75	68.01	76.99	73.88	75.41	70.19	48.41	57.30	66.01	53.71	59.23
Hier-GCN-BERT	75.12	66.47	71.94	77.62	75.61	77.68	75.00	66.84	70.69	68.32	62.54	65.09
Hier-GCN-XInet	76.99	68.44	70.53	77.81	75.84	76.35	73.09	64.17	69.38	70.3	63.69	67.12
MvP*	67.80	68.63	68.21	73.76	75.49	74.62	-	-	-	-	-	-
Ours(ECAN)	84.38	74.43	79.09	84.92	79.65	82.20	83.56	75.34	79.24	76.11	65.45	70.45

Table 3: The main results of the ACSC task. “**” refers to the results based on sentiment tuple prediction according to the original paper (Gou et al., 2023).

	REST 15			REST 16			LAP 15			LAP 16		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
ACD task												
w/o Senti-dis	90.89	80.13	84.43	90.78	84.55	86.80	93.08	76.19	87.40	82.59	72.41	76.85
w/o Cate-dis	89.87	<i>79.68</i>	<i>83.01</i>	90.55	<i>84.27</i>	86.33	92.06	<i>75.57</i>	<i>86.76</i>	<i>81.79</i>	<i>70.12</i>	<i>75.32</i>
w/o W-syn	89.67	80.23	84.11	91.20	84.62	87.85	92.90	77.19	87.80	82.98	73.17	77.76
w/o cl	<i>87.95</i>	79.83	84.19	<i>88.77</i>	85.57	<i>86.01</i>	<i>91.96</i>	76.32	87.31	83.08	71.86	77.31
Ours(ECAN)	91.62	82.39	85.67	91.69	86.02	88.75	93.64	78.43	88.80	83.45	76.23	79.21
ACSC task												
w/o Senti-dis	79.71	<i>72.76</i>	<i>74.07</i>	83.17	77.34	79.12	82.18	<i>72.23</i>	<i>76.05</i>	73.83	63.12	68.89
w/o Cate-dis	82.59	74.15	75.14	84.49	79.37	81.84	83.35	72.60	76.42	74.91	<i>62.68</i>	69.09
w/o W-syn	<i>79.28</i>	73.30	75.63	82.54	77.78	80.58	<i>81.60</i>	72.35	77.76	<i>74.72</i>	63.72	68.74
w/o cl	80.92	72.86	75.25	<i>82.09</i>	<i>76.76</i>	<i>79.01</i>	82.27	72.86	77.28	<i>72.87</i>	63.39	<i>67.05</i>
Ours(ECAN)	84.38	74.43	79.09	84.92	79.65	82.20	83.56	75.34	79.24	76.11	65.45	70.45

Table 4: Ablation study. “w/o Senti-dis” refers to the result without the sentiment disentanglement part, “w/o Cate-dis” indicates the result without the category disentanglement part, “w/o W-syn” shows the result without word-level syntactic learning, and “w/o cl” stands for the result without the contrastive learning loss. The italic font value indicates the worst results.

method shows the worst results in all datasets. One reason is that hierarchical learning approaches build hierarchical category-sentiment graphs to aggregate the inter-relationships between categories and sentiments by co-occurring category pairs, which balance the effect of high-frequency and low-frequency categories.

- In addition to modeling the hierarchical relations between categories and sentiments like hierarchical learning, the ECAN also learns the document-level contexts to assist in the category detection task via coherence-aware contrastive learning, which helps to achieve superior performance.
- Interestingly, sentiment disentanglement is also beneficial for identifying the categories in the sentences. As we can see from Figure 3, the aspect term, e.g., “food,” is also highlighted when identifying the positive sentiment word, e.g., “excellent,” in the sentiment disentanglement part ②. It supports that the sentiment disentanglement process helps to connect the relations between sentiments and categories due to the shared representations optimized by multi-task learning.

5.4.2. Performance of Aspect Category Sentiment Classification

Table 3 shows the results of aspect category sentiment classification. Overall, the ECAN attained an improvement over the second-best methods by 7.1~11.4% in precision, 2.8~12.7% in recall, and 5.0~12.1% in the F1-score. In particular, it achieved remarkable results for REST 15 and LAP 15 compared with all baselines. The improvement of 7.8% on average through all metrics in the ACSC task is significantly better than that of 1.6% in the ACD task. In addition to similar observations in the ACD task, we have the following findings:

- The ECAN significantly outperforms all baseline models, highlighting its capability to achieve superior performance by disentangling sentiment representations related to aspect categories. Furthermore, coherence-aware representation learning and word-level syntactic learning play crucial roles in enhancing sentiment semantics.
- The Hier-BERT and LC-BERT work well in the ACD task, while they fail to effectively identify

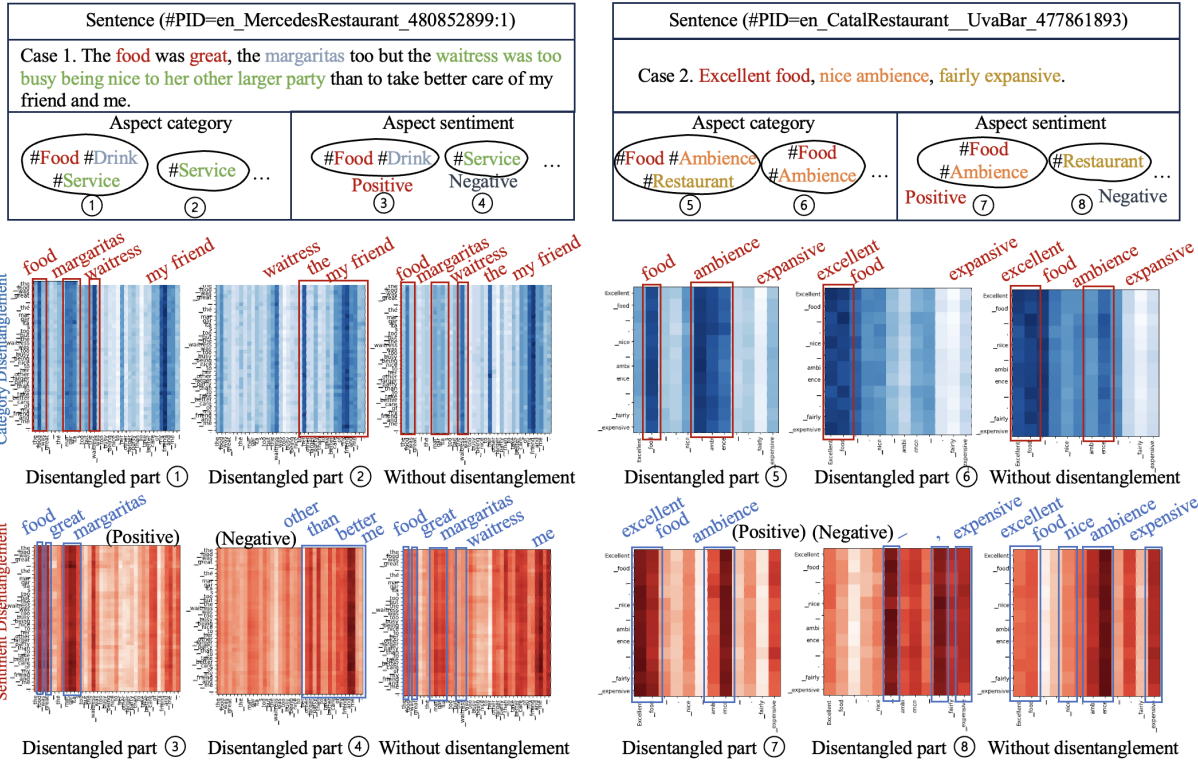


Figure 3: The visualization result of category and sentiment disentanglement. The top box on the left/right side indicates a sentence in a review with the source and disentangled categories and sentiments by the ECAN. The blue color and the red color denote the visualized correlations between categories and between sentiments in disentanglements. The warmer the tone colors, the higher the correlations.

the aspect-category sentiment, as they do not utilize GCN to learn many-to-many relatedness between aspect categories and sentiments for node embedding enhancement.

5.5. Ablation Study

We conducted ablation experiments to verify the effectiveness of each component in the ECAN. As shown in Table 4, the results in the ACD and ACSC tasks prompt the following observations:

- As expected, category disentanglement (i.e., w/o Cate-dis) contributes the most to the ACD task, and sentiment disentanglement (i.e., w/o Senti-dis) works better in the ACSC task. It also supports that without hierarchical disentanglement, the performance is fairly limited by the entangled sentiments within each sentence.
- Without coherence-aware representation by contrastive learning (i.e., w/o cl), the performance of the ECAN significantly drops in both the ACD and ACSC tasks, demonstrating that document-level coherence is a strong clue in assisting in learning coherence-aware representation.
- The underperformance of the ECAN without word-level syntactic learning (w/o W-syn) in the

ACSC task, indicates that syntactic dependencies help to derive enhanced sentiment semantic representations. We also found that syntactic analysis contributes to the ACSC of sentences containing conflicting sentiment expressions, such as the sentence “Even the chickpeas, which I normally find too dry, were good”.

5.6. Visualization of Hierarchical Disentanglement

To understand how the hierarchical disentanglement model assists in ACSA, we present two cases and provide performance comparisons through visualization. The comparisons are (i) the disentangled category and sentiment parts in reviews, and (ii) the counterparts without disentanglement.

Case 1. There are three aspect categories, “food,” “drink,” and “service,” and their sentiments are positive, positive, and negative, respectively. From the left part of Figure 3, we make the following observations:

- For category disentanglement, the disentangled part ① captures the aspect categories of “food,” “drink,” and “service,” and part ② disentangles “service” from others;

- For the sentiment disentanglement, part ③ detected positive sentiment of aspect categories, “food” and “drink” and part ④ highlighted the words expressing negative sentiments, as the review mentions the waitress put more efforts on her other larger party;
- Without disentanglement, the model could focus on only one part of the categories, i.e., it highlighted the negative sentiment on the waitress while reducing the impact of the sentiments on food and drink.

Case 2. Three aspect categories are mentioned in the review, i.e., “food,” “ambiance,” and “restaurant,” with positive, positive, and negative sentiment polarities, respectively. By observing the right part of Figure 3, we have the following findings:

- For category disentanglement, the disentangled part ⑤ recognizes the categories of “food” and “ambiance,” while part ⑥ concentrates on “food”;
- For the sentiment disentanglement, part ⑦ correctly identifies the positive sentiments as it focuses on “excellent food” and “ambiance.” In contrast, part ⑧ avoided the attention of positive words, such as “excellent” and “nice” and focused on “expensive” and other neutral tokens, such as “-” and “”.

5.7. Error Analysis

We conducted error analyses on the REST 16 and LAP 16 datasets. There are two major types of errors for the ACD task:

- Generalized categories, such as “restaurant# general” in the restaurant domain and “laptop# general” in the laptop domain, pose difficulties in our ECAN. For example, the category entity of “*great lunch spot.*” should be “restaurant,” while the ECAN predicts it to be “food” incorrectly.
- The special terms that appear with low frequency, such as names of places, persons, and products, are often ignored by our ECAN. This inspires us to inject more knowledge, e.g., by leveraging the large language model (Zhang et al., 2023). This remains a rich space for further exploration.

We also found two typical cases of errors for the ACSC task:

- In the case of a sentence with more than two transitions, for example, “*I liked the atmosphere very much but the food was not worth the price.*”, the ECAN incorrectly predicts positive polarity on the price of the food. It inspired us to explore special feature extractors for this to further improve the performance.
- Neutral sentiment is still challenging for the

ECAN, although we observed that it identified positive and negative sentiments well. One of the reasons is that the shortage of neutral training examples, which is shown in Table 1, complicates its identification. Analyzing neutral sentiments is also an interesting direction for future work.

6. Conclusion

We proposed an enhanced coherence-aware network (ECAN). It leverages coherence modeling to learn explicit opinions from contexts and a hierarchical disentanglement of category and sentiment to mine more fine-grained features for both ACD and ACSC tasks. Extensive experimental and visualization results show that our ECAN effectively decouples categories and sentiments entangled in the coherence representations and achieves SOTA performance.

Future work includes (1) seeking more interpretable solutions to disentangle the categories and sentiments; and (2) exploring stronger clues to guide the disentangled channel to learn independent features between channels, such as maximizing the variances of various disentangled representations.

Limitations

There is space for further improving the ability of ECAN to disentangle categories. The ECAN model adopts time-consuming modules, i.e., XLNet ($O(n^2)$) and GCN ($O(n^2)$) where n refers to the number of words, therefore its computational cost heavily relies on the length of textual reviews.

Ethics Statement

This paper does not involve the presentation of a new dataset, an NLP application, and the utilization of demographic or identity characteristics information.

Acknowledgements

We would like to thank anonymous reviewers for their thorough comments and suggestions. This work is supported by JKA and Kajima Foundation’s Support Program and the China Scholarship Council (No.202208330091).

7. Bibliographical References

Tushar Abhishek, Daksh Rawat, Manish Gupta, and Vasudeva Varma. 2021. Transformer models

- for text coherence assessment. *arXiv preprint arXiv:2109.02176*.
- Yu Bao, Hao Zhou, Shujian Huang, Lei Li, Lili Mou, Olga Vechtomova, Xinyu Dai, and Jiajun Chen. 2019. Generating sentences from disentangled syntactic and semantic spaces. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6008–6019.
- Hongjie Cai, Yaofeng Tu, Xiangsheng Zhou, Jianfei Yu, and Rui Xia. 2020. Aspect-category based sentiment analysis with hierarchical graph convolutional network. In *Proceedings of the 28th international conference on computational linguistics*, pages 833–843.
- Hongjie Cai, Rui Xia, and Jianfei Yu. 2021. Aspect-category-opinion-sentiment quadruple extraction with implicit aspects and opinions. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 340–350.
- Mingda Chen, Qingming Tang, Sam Wiseman, and Kevin Gimpel. 2019. A multi-task approach for disentangling syntax and semantics in sentence representations. *arXiv preprint arXiv:1904.01173*.
- Pengyu Cheng, Martin Renqiang Min, Dinghan Shen, Christopher Malon, Yizhe Zhang, Yitong Li, and Lawrence Carin. 2020. Improving disentangled text representation learning with information-theoretic guidance. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7530–7541.
- Jin Cui, Fumiyo Fukumoto, Xinfeng Wang, Yoshimi Suzuki, Jiyi Li, and Wanzeng Kong. 2023. Aspect-category enhanced learning with a neural coherence model for implicit sentiment analysis. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11345–11358.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Youmna Farag and Helen Yannakoudakis. 2019. Multi-task learning for coherence modeling. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 629–639.
- Tianhao Gao, Jun Fang, Hanyu Liu, Zhiyuan Liu, Chao Liu, Pengzhang Liu, Yongjun Bao, and Weipeng Yan. 2022. Lego-absa: A prompt-based task assemblable unified generative framework for multi-task aspect-based sentiment analysis. In *Proceedings of the 29th international conference on computational linguistics*, pages 7002–7012.
- Zhibin Gou, Qingyan Guo, and Yujiu Yang. 2023. MvP: Multi-view prompting improves aspect sentiment tuple prediction. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4380–4397, Toronto, Canada. Association for Computational Linguistics.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. Deberta: Decoding-enhanced bert with disentangled attention.
- Yahao Hu, Wei Tao, Yifei Xie, Yi Sun, and Zhisong Pan. 2023. Token-level disentanglement for unsupervised text style transfer. *Neurocomputing*, page 126823.
- Sungho Jeon and Michael Strube. 2020. Centering-based neural coherence modeling with hierarchical discourse segments. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7458–7472.
- Sungho Jeon and Michael Strube. 2022. Entity-based neural local coherence modeling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7787–7805.
- Vineet John, Lili Mou, Hareesh Bahuleyan, and Olga Vechtomova. 2019. Disentangled representation learning for non-parallel text style transfer. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 424–434.
- Prathyusha Jwalapuram, Shafiq Joty, and Xiang Lin. 2022. Rethinking self-supervision objectives for generalizable coherence modeling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6044–6059.
- Seunghun Lee, Sunghyun Cho, and Sunghoon Im. 2021. Dranet: Disentangling representation and adaptation networks for unsupervised cross-domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15252–15261.
- Ruifan Li, Hao Chen, Fangxiang Feng, Zhanyu Ma, Xiaojie Wang, and Eduard Hovy. 2021. Dual graph convolutional networks for aspect-based sentiment analysis. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint*

- Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6319–6329.
- Yuncong Li, Zhe Yang, Cunxiang Yin, Xu Pan, Lunan Cui, Qiang Huang, and Ting Wei. 2020a. A joint model for aspect-category sentiment analysis with shared sentiment prediction layer. In *Chinese Computational Linguistics: 19th China National Conference, CCL 2020, Hainan, China, October 30–November 1, 2020, Proceedings 19*, pages 388–400. Springer.
- Yuncong Li, Cunxiang Yin, and Sheng-hua Zhong. 2020b. Sentence constituent-aware aspect-category sentiment analysis with graph attention networks. In *Natural Language Processing and Chinese Computing: 9th CCF International Conference, NLPCC 2020, Zhengzhou, China, October 14–18, 2020, Proceedings, Part I 9*, pages 815–827. Springer.
- Yuncong Li, Cunxiang Yin, Sheng-hua Zhong, and Xu Pan. 2020c. Multi-instance multi-label learning networks for aspect-category sentiment analysis. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3550–3560.
- Bin Liang, Jiachen Du, Ruifeng Xu, Binyang Li, and Hejiao Huang. 2019. Context-aware embedding for targeted aspect-based sentiment analysis. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4678–4683.
- Dongliang Liao, Jin Xu, Gongfu Li, and Yiru Wang. 2021. Hierarchical coherence modeling for document quality assessment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13353–13361.
- Ting Lin, Aixin Sun, and Yequan Wang. 2023. Educapsule: aspect-based sentiment analysis at clause level. *Knowledge and Information Systems*, 65(2):517–541.
- Yen-Cheng Liu, Yu-Ying Yeh, Tzu-Chien Fu, Sheng-De Wang, Wei-Chen Chiu, and Yu-Chiang Frank Wang. 2018. Detach and adapt: Learning cross-domain disentangled deep representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8867–8876.
- Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. 2019. Challenging common assumptions in the unsupervised learning of disentangled representations. In *International conference on machine learning*, pages 4114–4124. PMLR.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Fukun Ma, Xuming Hu, Aiwei Liu, Yawen Yang, S Yu Philip, Lijie Wen, et al. 2023. Amr-based network for aspect-based sentiment analysis. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 322–337.
- Julian Richard Medina and Jugal Kalita. 2018. Parallel attention mechanisms in neural machine translation. In *2018 17th IEEE international conference on machine learning and applications (ICMLA)*, pages 547–552. IEEE.
- Xiaoyong Mei, Yougen Zhou, Chenjing Zhu, Mengting Wu, Ming Li, and Shirui Pan. 2023. A disentangled linguistic graph model for explainable aspect-based sentiment analysis. *Knowledge-Based Systems*, 260:110150.
- Sharmila Reddy Nangi, Niyati Chhaya, Sopan Khosla, Nikhil Kaushik, and Harshit Nyati. 2021. Counterfactuals to control latent disentangled text representations for style transfer. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 40–48.
- Emanuel Huber da Silva and Ricardo Marcondes Marcacini. 2021. Aspect-based sentiment analysis using bert with disentangled attention. In *Proceedings*.
- Yi Tay, Luu Anh Tuan, and Siu Cheung Hui. 2018. Learning to attend via word-aspect associative fusion for aspect-based sentiment analysis. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Brian Thompson and Matt Post. 2020. Paraphrase generation as zero-shot multilingual translation: Disentangling semantic similarity from lexical and syntactic diversity. In *Proceedings of the Fifth Conference on Machine Translation*, pages 561–570.
- Shanshan Wang, Yiyang Chen, Zhenwei He, Xun Yang, Mengzhu Wang, Quanzeng You, and Xingyi Zhang. 2023. Disentangled representation learning with causality for unsupervised domain adaptation. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 2918–2926.
- Xin Wang, Hong Chen, Si'ao Tang, Zihao Wu, and Wenwu Zhu. 2022. Disentangled representation learning. *arXiv preprint arXiv:2211.11695*.

- Yequan Wang, Aixin Sun, Minlie Huang, and Xiaoyan Zhu. 2019. Aspect-level sentiment analysis using as-capsules. In *The world wide web conference*, pages 2033–2044.
- Yang Wu, Zhenyu Zhang, Yanyan Zhao, and Bing Qin. 2021. Locate and combine: A two-stage framework for aspect-category sentiment analysis. In *Natural Language Processing and Chinese Computing: 10th CCF International Conference, NLPCC 2021, Qingdao, China, October 13–17, 2021, Proceedings, Part I 10*, pages 595–606. Springer.
- Peng Xu, Hamidreza Saghir, Jin Sung Kang, Teng Long, Avishek Joey Bose, Yanshuai Cao, and Jackie Chi Kit Cheung. 2019. A cross-domain transferable neural coherence model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 678–687.
- Wei Xue and Tao Li. 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)*, pages 2514–2523.
- Heng Yang and Ke Li. 2021. Improving implicit sentiment learning via local sentiment aggregation. *arXiv preprint arXiv:2110.08604*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pre-training for language understanding. *Advances in neural information processing systems*, 32.
- Chen Zhang, Qiuchi Li, and Dawei Song. 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)*, pages 4568–4578.
- Mi Zhang and Tiejun Qian. 2020. Convolution over hierarchical syntactic and lexical graphs for aspect level sentiment analysis. In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pages 3540–3549.
- Wenxuan Zhang, Yue Deng, Bing Liu, Sinno Jialin Pan, and Lidong Bing. 2023. Sentiment analysis in the era of large language models: A reality check. *arXiv preprint arXiv:2305.15005*.
- Wenxuan Zhang, Xin Li, Yang Deng, Lidong Bing, and Wai Lam. 2022a. A survey on aspect-based sentiment analysis: Tasks, methods, and challenges. *IEEE Transactions on Knowledge and Data Engineering*.
- Yuhao Zhang, Ying Zhang, Wenya Guo, Xiangrui Cai, and Xiaojie Yuan. 2022b. Learning disentangled representation for multimodal cross-domain sentiment analysis. *IEEE Transactions on Neural Networks and Learning Systems*.
- Zheng Zhang, Zili Zhou, and Yanna Wang. 2022c. Ssegcn: Syntactic and semantic enhanced graph convolutional network for aspect-based sentiment analysis. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4916–4925.
- Tao Zhou and Kris MY Law. 2022. Semantic relatedness enhanced graph network for aspect category sentiment analysis. *Expert Systems with Applications*, 195:116560.

8. Language Resource References

- Maria Pontiki, Dimitrios Galanis, Harris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, AL-Smadi Mohammad, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphee De Clercq, et al. 2016. Semeval-2016 task 5: Aspect based sentiment analysis. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 19–30.
- Maria Pontiki, Dimitrios Galanis, Harris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. Semeval-2015 task 12: Aspect based sentiment analysis. In *Proceedings of the 9th international workshop on semantic evaluation (SemEval 2015)*, pages 486–495.

A. Appendix

A.1. Implementation and hyperparameter setting

We implemented ECAN and had experimented with Pytorch on a single GPU: NVIDIA GeForce RTX 3090 (24GB memory). The search ranges of the hyperparameters used in our experiments are shown in Table 5.

A.2. Parameter Analysis

We examined how key hyperparameters, the number of category disentanglement blocks d_c , and the

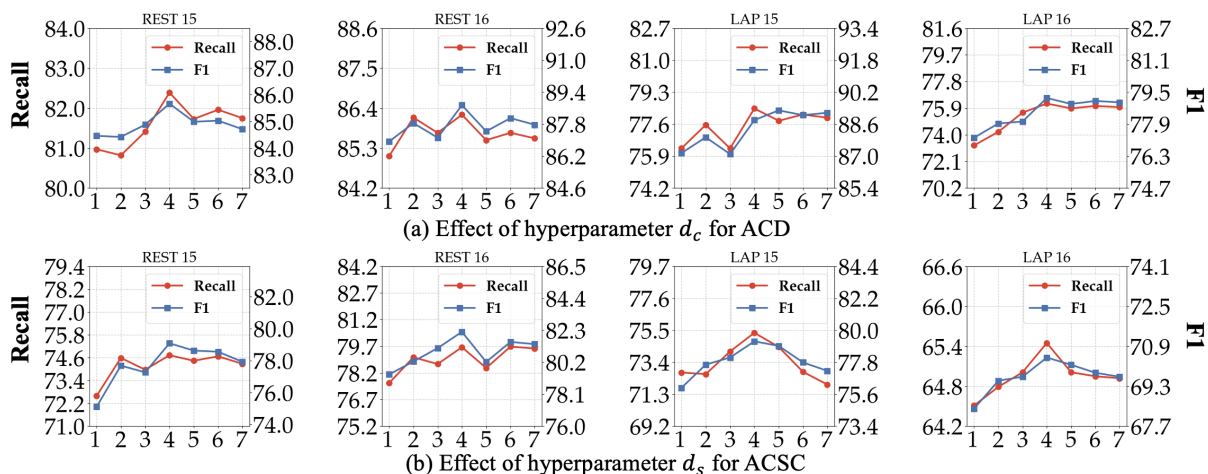


Figure 4: The effect of hyperparameters d_c and d_s .

Parameter	Range
LR of coherence modeling	$1e-6 \sim 1e-5$
LR of others	$1e-5 \sim 1e-4$
Weight decay	{ $1e-4, 1e-3, 1e-2$ }
Dropout rate	{0.1, 0.2, 0.3}
#Negative Samples B	{5, 6, 7, 8, 9, 10}
Margin τ	{0.05, 0.1, 0.15, 0.2}
δ_1	{0.05, 0.1, 0.15, 0.2}
δ_2	{0.3, 0.4, 0.5, 0.6, 0.7, 0.8}
δ_3	{0.3, 0.4, 0.5, 0.6, 0.7, 0.8}
#Block of GCN	{1, 2, 3}
# disentangled blocks	{1, 2, 3, 4, 5, 6, 7}

Table 5: Search range of each hyperparameter: LR refers to the learning rate. LR of coherence modeling indicates the learning rate of coherence-aware representation learning. LR of others shows hierarchical disentanglement and word-level syntactic learning.

number of sentiment disentanglement blocks d_s affected the performance of the ECAN. The results are given in Figure 4.

Effect of d_c . The highest value of d_c is four for all datasets. We can also observe that on the REST 15 and REST 16 datasets, the performance is better when d_c increases to four, while the performance decreases when d_c becomes higher than four. On LAP 15 and LAP 16, the performance does not improve. However, its computational cost increases with d_c being higher than four.

Effect of d_s . The best value of d_s is also four for all datasets. It is reasonable that the performance is getting worse with d_c being higher than four because there are only three types of sentiment polarities and a higher number of disentanglement parts could mislead the model.