

Utilizing Local Hierarchy with Adversarial Training for Hierarchical Text Classification

Zihan Wang, Peiyi Wang, Houfeng Wang*

National Key Laboratory of Multimedia Information Processing
School of Computer Science, Peking University
{wangzh9969, wangpeiyi9979}@gmail.com, wanghf@pku.edu.cn

Abstract

Hierarchical text classification (HTC) is a challenging subtask of multi-label classification due to its complex taxonomic structure. Nearly all recent HTC works focus on how the labels are structured but ignore the sub-structure of ground-truth labels according to each input text which contains fruitful label co-occurrence information. In this work, we introduce this local hierarchy with an adversarial framework. We propose a HiAdv framework that can fit in nearly all HTC models and optimize them with the local hierarchy as auxiliary information. We test on two typical HTC models and find that HiAdv is effective in all scenarios and is adept at dealing with complex taxonomic hierarchies. Further experiments demonstrate that the promotion of our framework indeed comes from the local hierarchy and the local hierarchy is beneficial for rare classes which have insufficient training data.

Keywords: hierarchical text classification, adversarial training, local hierarchy

1. Introduction

Hierarchical text classification (HTC) aims to categorize a text sample into a set of labels that are organized as a structured hierarchy (Silla and Freitas, 2011). Various multi-class classification problems can be extended to HTC by giving a pre-defined label hierarchy, such as scientific documents classification (Lewis et al., 2004; Sadat and Caragea, 2022) or news categorization (Kowsari et al., 2018). As the main difference to ordinary multi-class classification problems, how to utilize the large-scale, imbalanced label hierarchy is the key challenge of HTC (Mao et al., 2019).

Nearly all of the recent HTC works modeled the label hierarchy with a graph encoder: they feed in either both text and label information for a mixture representation (Zhou et al., 2020; Wang et al., 2022b) or the hierarchy solely for a graph representation which is then fused with text representation (Zhou et al., 2020; Deng et al., 2021; Zhao et al., 2021). However, these works mainly focus on the constant global hierarchy but ignore the sub-graph corresponding to each input text, which can contain structured label co-occurrence information (Jiang et al., 2022). This so-called local hierarchy is first introduced by Wang et al. (2022a) to generate positive samples for contrastive learning and they observe small improvements. However, this work considers the local hierarchy as flat but ignores its structure. Following this work, Jiang et al. (2022) takes further advantage of local hierarchy by a sequence-to-sequence approach. In their method, the local hierarchy is utilized through teacher forcing, which takes a ground-truth label

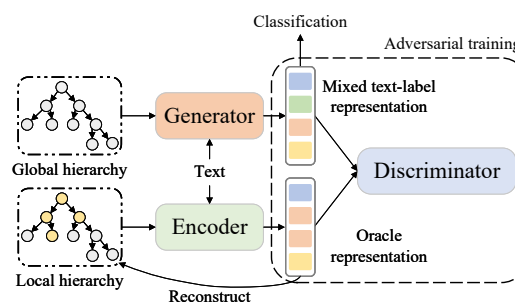


Figure 1: A demonstration of our adversarial framework. A generator and an encoder use global and local hierarchy as input respectively, and the output representations are trained adversarially.

sequence as input and predicts the next step of that sequence. However, the transition from a graph to a sequence still loses some hierarchical information. Besides, their method generates labels layer by layer, so this method requires multiple inference times compared to other works.

Is there a way that can directly model the local hierarchy as a graph? Stepping forward from the contrastive learning of Wang et al. (2022a), a simple idea will supply: we can train the original text representation to be similar to a local-hierarchy-incorporated representation. In contrast to contrastive learning which requires dissimilar example pairs and previous work ignores totally, in this paper, we propose a hierarchy-aware adversarial framework (HiAdv) to incorporate local hierarchy.

In detail, during training, our HiAdv learns two representations adversarially. As shown in Figure 1, a generator inputs the text and global hierarchy and outputs a mixed text-label representation. Besides,

*Corresponding author.

an encoder generates an oracle representation by giving the ground-truth labels as input. In order to obtain high-quality local-hierarchy-related representation, the encoder organizes as an autoencoder-like structure: it encodes the local hierarchy to a representation and then reconstructs the representation back to the local hierarchy. A discriminator is then asked to distinguish the raw representation from the one with the local hierarchy. By encouraging the generator to fool the discriminator, the raw representation should be similar enough to the one with the local hierarchy after training. Furthermore, this idea does not rely on specific model architecture: we can apply it to any HTC model that involves a graph encoder.

We summarize our contributions as follows:

- We propose an adversarial framework, HiAdv, for HTC to incorporate local hierarchy. This is the first attempt to incorporate the entire local hierarchy.
- The framework does not rely on specific architecture so it can adapt to most existing HTC models which rely on a graph encoder to encode label hierarchy.
- Experiments demonstrate that our framework can constantly improve the performance of basic models on three datasets and achieve new state-of-the-art with the latest architecture as a backbone. We release our code at <https://github.com/wzh9969/HiAdv>.

2. Related Work

2.1. Hierarchy text classification

Hierarchical text classification (HTC) is a challenging task due to its large-scale, imbalanced, and structured label hierarchy (Mao et al., 2019). Recent works for HTC focus on global approaches, which build only one classifier for the entire graph (Zhou et al., 2020). The early global approaches neglect the hierarchical structure of labels and view the problem as a flat multi-label classification (Johnson and Zhang, 2015). After various attempts to coalesce the label structure (Wu et al., 2019; Mao et al., 2019; Zhang et al., 2021), Zhou et al. (2020) demonstrate that encoding the holistic label structure directly by a structure encoder can most improve performance. Following this research, a bunch of models try to study how the structure encoder should interact with the text (Chen et al., 2020, 2021; Deng et al., 2021; Zhao et al., 2021; Zhu et al., 2023). Besides the direct text-label interaction, Wang et al. (2022b) adopts prompt tuning for HTC to better utilize pretrained language models, and they predict labels layer-wisely to manually

incorporate the depth information of the label hierarchy. Despite a different training strategy, this work still follows the principle of integrating the output of the structure encoder into the text representation. Ji et al. (2023) studies the problem under few-shot settings.

While previous works focus on the global hierarchy (i.e., the predefined label hierarchy), a few recent works claim the importance of local hierarchy (i.e., the hierarchy corresponding to each input text). Wang et al. (2022a) adopts contrastive learning to directly inject hierarchical knowledge into the text encoder. In this work, ground-truth labels serve as the guidance to select classification-related tokens, which are then constructed as positive samples for contrastive learning. Jiang et al. (2022) views the problem as sequence generation: labels at different hierarchy layers are generated step by step. From this perspective, the local hierarchy is introduced naturally: only the labels in the local hierarchy should be generated. However, this method gets rid of the structure encoder entirely, which has been proven to be effective in all of the previous works.

2.2. Adversarial Training

Adversarial training is originally proposed in the Generative Adversarial Network (GAN) for image generation (Goodfellow et al., 2020). It adversarially trains a discriminator against a generator: the discriminator aims to distinguish real images from generated ones while the generator struggles to fool the discriminator. In NLP, adversarial networks are also applied in some generation tasks. For sequence generation, adversarial training serves as an alternative to step-by-step generation to avoid exposure bias issues (Scialom et al., 2020; Liu et al., 2020; Chai et al., 2023). For NLU tasks, recent studies show that adversarial training can be well-deployed on pre-trained language models (Zhu et al., 2020; Jiang et al., 2020; Wang et al., 2021; Wu et al., 2023). Some works of neural topic modeling introduce adversarial networks to learn topic distributions (Wang et al., 2019, 2020; Hu et al., 2020). Both Wu et al. (2021) and Wu et al. (2022) adopt adversarial training for multi-domain text classification to extract shared features across domains. However, the usage of adversarial training in their works focuses on multi-domain data and has little relation to the classification task itself. As a result, these methods cannot be adopted for general text classification or the hierarchical text classification that we are studying.

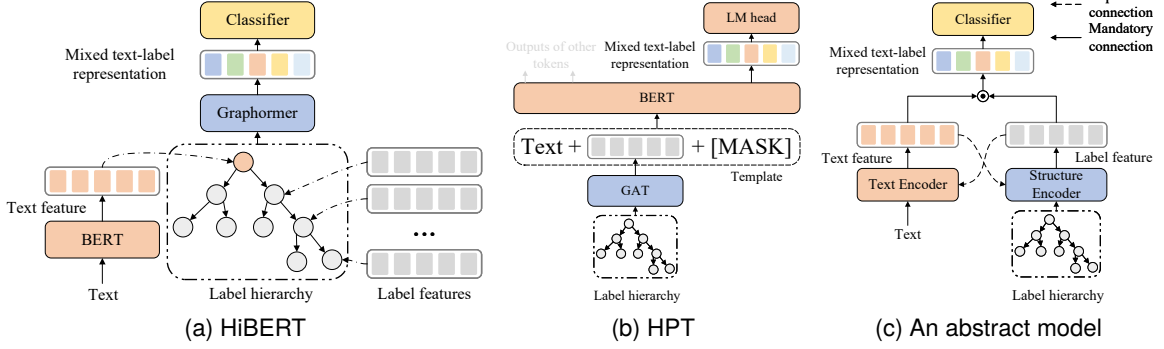


Figure 2: Two HTC models and an abstract HTC model. (a) HiBERT. Feed BERT output text representation into a graph encoder. (b) HPT. A prompt tuning model with a hierarchy-aware template. (c) An abstract model. Two encoders dealing with text and structure respectively and a mixture mechanism for a mixed representation.

3. Preliminary

3.1. Problem Definition

For each hierarchical text classification (HTC) dataset, we have a predefined label hierarchy (i.e., global hierarchy) $\mathcal{H} = (\mathcal{Y}, E)$, where \mathcal{Y} is the label set (also the node set of \mathcal{H}) and E is the edge set. \mathcal{H} is organized as a Directed Acyclic Graph (DAG) but we focus on a setting where every node except the root has one and only one father so that the hierarchy can be simplified as a tree-like structure. Given an input text \mathbf{x} , the models aim to categorize it into a label set (i.e., local hierarchy) $Y \subseteq \mathcal{Y}$. The predicted label set Y corresponds to one or more paths in \mathcal{H} . Each path starts from the root node and ends at any node on the tree.

3.2. Basic Models

Our framework is model-unrelated, so it is necessary to introduce two basic architectures we adopt in this work.

3.2.1. HiBERT

A text encoder with an auxiliary structure encoder is one of the simplest HTC models. Following previous work (Wang et al., 2022a), we adopt BERT (Devlin et al., 2019) as the text encoder and Graphormer (Ying et al., 2021) as the structure encoder. We name this architecture Hierarchy-aware BERT (HiBERT) for simplicity.

Specifically, as in Figure 2a, given a text input \mathbf{x} , BERT encodes it to a text representation,

$$\mathbf{h}_{\text{text}} = \text{BERT}(\mathbf{x}) \quad (1)$$

The structure encoder then takes the label hierarchy and text representation as input and outputs a mixed representation. To utilize label information,

we use the average of BERT token embedding of the label name as its label embedding \mathbf{l}_i . For the root node which has no particular meaning, we use the text representation as label embedding. We use the output graph representation of the root node as a mixed representation of text and labels:

$$\mathbf{h}_{\text{mix}} = \text{Graphormer}([\mathbf{h}_{\text{text}}, \mathbf{L}], \mathcal{H})_{\text{root}} \quad (2)$$

Finally, a linear classifier is used to calculate the probability of each class:

$$P = \text{sigmoid}(\mathbf{W}_c \mathbf{h}_{\text{mix}} + \mathbf{b}_c) \quad (3)$$

where \mathbf{W}_c and \mathbf{b}_c are weight matrix and bias.

3.2.2. Hierarchy Prompt Tuning

Hierarchy Prompt Tuning (HPT) (Wang et al., 2022b) is the state-of-the-art model which regards the problem as a masked language model problem. As shown in Figure 2b, HPT first encodes label hierarchy into prompting features \mathbf{T} by a GAT:

$$\mathbf{T} = \text{GAT}(\mathbf{L}, \mathcal{H}) \quad (4)$$

where $\mathbf{L} = \{\mathbf{l}_i\}$ is the label embedding that contains node information. Then, the input text \mathbf{x} and prompting features \mathbf{T} are filled into a template, where a [MASK] position is reserved as in the masked language model (MLM) task. BERT then encodes the template but we focus only on the output representation of the [MASK] token. By calculating the probability of filling each label embedding into the [MASK] position, HPT predicts in an MLM manner. We define the BERT output of the [MASK] position as mixed text-label representation \mathbf{h}_{mix} in accordance with HiBERT. Please refer to the original paper for more details.

4. Methodology

In this section, we introduce HiAdv. The whole framework is illustrated in Figure 3. Our framework can be applied to any HTC architecture that involves a text encoder and a graph encoder.

4.1. Abstraction of HTC Models

Despite differing in architecture details, both preceding models as well as most other architectures HTC works selected can be abstracted to a text encoder E_{text} and a structural encoder $E_{\text{structure}}$:

$$\begin{aligned} \mathbf{h}_{\text{text}} &= E_{\text{text}}(\mathbf{x}, \mathbf{h}_{\text{label}}) \\ \mathbf{h}_{\text{label}} &= E_{\text{structure}}(\mathbf{L}, \mathcal{H}, \mathbf{h}_{\text{text}}) \end{aligned} \quad (5)$$

As in Equation 5 and Figure 2c, the text encoder takes text \mathbf{x} as input while the structural encoder is built according to the label hierarchy \mathcal{H} and takes label embedding \mathbf{L} as input. Besides, one encoder may utilize the output representation of the other. Each encoder outputs a representation in R^d .

Finally, a mixture mechanism coalesces two representations:

$$\mathbf{h}_{\text{mix}} = \mathbf{h}_{\text{text}} \odot \mathbf{h}_{\text{label}} \quad (6)$$

and a multi-label classifier C predicts probability distributions:

$$P = \text{sigmoid}(C(\mathbf{h}_{\text{mix}})) \quad (7)$$

During training, a multi-label classification loss such as binary cross entropy is adopted:

$$L_C = \text{Loss}(P, Y) \quad (8)$$

This abstract model generates a representation \mathbf{h}_{mix} according to a given input \mathbf{x} , so we can define it as a generator under the perspective of adversarial training.

4.2. Encoder Network

In the context of adversarial training, an encoder network encodes a real input $\hat{\mathbf{x}}$ into a real distribution. In our scenario, the encoder network attempts to generate an oracle representation that takes local hierarchy as input and can perform the best classification. We organize the encoder as an autoencoder: it encodes local hierarchy along with text into the oracle representation and then reconstructs the local hierarchy.

To incorporate local hierarchy, we modify the label embedding in Equation 5. For label y_i , we add an oracle label embedding to indicate whether this label is in the local hierarchy Y :

$$\hat{\mathbf{l}}_i = \mathbf{l}_i + \begin{cases} \mathbf{e}_1, & y_i \in Y \\ \mathbf{e}_0, & y_i \notin Y \end{cases} \quad (9)$$

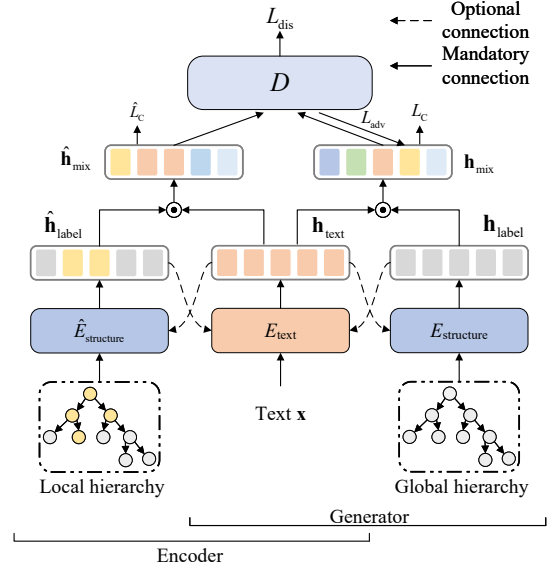


Figure 3: A demonstration of our adversarial framework. The generator and the encoder share the same text encoder. We omit the classifier for clarity, which takes \mathbf{h}_{mix} and $\hat{\mathbf{h}}_{\text{mix}}$ for input to classify and generates classification losses L_C and \hat{L}_C during training.

where \mathbf{e}_0 and \mathbf{e}_1 are two learnable embeddings that indicate whether y_i is in Y . As a result, the label embedding matrix $\hat{\mathbf{L}} = \{\hat{\mathbf{l}}_i\}$ contains the local hierarchy as a prior. As in Figure 3, we use an independent structure encoder to deal with the local hierarchy while other components of the encoder share with the generator:

$$\begin{aligned} \hat{\mathbf{h}}_{\text{label}} &= \hat{E}_{\text{structure}}(\hat{\mathbf{L}}, \mathcal{H}, \mathbf{h}_{\text{text}}) \\ \hat{\mathbf{h}}_{\text{mix}} &= \mathbf{h}_{\text{text}} \odot \hat{\mathbf{h}}_{\text{label}} \end{aligned} \quad (10)$$

where $\hat{\mathbf{h}}_{\text{mix}}$ is the demanded oracle representation which takes the local hierarchy as a prior.

The oracle representation $\hat{\mathbf{h}}_{\text{mix}}$ is then reconstructed to the local hierarchy Y with the classifier C , which is similar to an autoencoder that aims to find the best representation for reconstructing the local hierarchy. During training, $\hat{\mathbf{h}}_{\text{mix}}$ is guided by the classification loss the same way as \mathbf{h}_{mix} in the generator:

$$\hat{L}_C = \text{Loss}(\text{sigmoid}(C(\hat{\mathbf{h}}_{\text{mix}})), Y) \quad (11)$$

4.3. Discriminator

With a generated representation \mathbf{h}_{mix} and an oracle representation $\hat{\mathbf{h}}_{\text{mix}}$, a discriminator D attempts to distinguish whether a representation \mathbf{h} is generated with or without the local hierarchy. This is a binary classification problem so we use a simple two-layer linear classifier as the discriminator:

$$p = \text{Sigmoid}(\mathbf{W}_2 \text{ReLU}(\mathbf{W}_1 \mathbf{h} + \mathbf{b}_1) + b_2) \quad (12)$$

Algorithm 1: Framework of HiAdv.

Input: Training dataset $T_{\text{train}} = \{(\mathbf{x}, Y)\}$,
global hierarchy \mathcal{H} , test dataset
 $T_{\text{test}} = \{\mathbf{x}\}$
Output: Model prediction $\{\hat{Y}\}$
// Initialization
1 Initialize E_{text} , $E_{\text{structure}}$, $\hat{E}_{\text{structure}}$, C , and
 D randomly;
// Training
2 **while** not converged **do**
3 Sample a data point (\mathbf{x}, Y) from T_{train} ;
4 Calculate \mathbf{h}_{mix} and $\hat{\mathbf{h}}_{\text{mix}}$ by Equation 6
 and 10;
 // Update D
5 Fill \mathbf{h}_{mix} and $\hat{\mathbf{h}}_{\text{mix}}$ successively into
 Equation 12 to obtain p and \hat{p} ;
6 Calculate L_{dis} and \hat{L}_{dis} by Equation 13;
7 Back-propagate $L_{\text{dis}} + \hat{L}_{\text{dis}}$ and update
 parameters of the discriminator;
 // Update the rest
8 Calculate classification loss L_C and \hat{L}_C
 by Equation 8 and 11;
9 Calculate adversarial loss L_{adv} by
 Equation 15;
10 Back-propagate $L_C + \hat{L}_C + L_{\text{adv}}$ and
 update parameters of all components
 except the discriminator;
11 **end**
 // Testing
12 **foreach** Text $\mathbf{x} \in T_{\text{test}}$ **do**
13 Calculate P by Equation 7 using E_{text} ,
 $E_{\text{structure}}$, and C ;
14 Select $\hat{Y} = \{P_i > \tau \forall i\}$;
15 **end**

where $\mathbf{W}_1 \in R^{d \times d}$, $\mathbf{W}_2 \in R^{1 \times d}$, $\mathbf{b}_1 \in R^d$, and
 $b_2 \in R$ are weights and bias.

The discriminator is guided by binary cross-
entropy loss:

$$L_{\text{dis}} = -(I \log p + (1 - I) \log(1 - p)) \quad (13)$$

where I is an indicator

$$I = \begin{cases} 0, & \mathbf{h} = \mathbf{h}_{\text{mix}} \\ 1, & \mathbf{h} = \hat{\mathbf{h}}_{\text{mix}} \end{cases} \quad (14)$$

4.4. Adversarial Framework

The generator, the encoder, and the discriminator
train adversarially. As mentioned, the discriminator
aims to distinguish the representation while the
generator attempts to fool the discriminator. As
shown in Figure 3, besides the classification loss,
the generator is guided by an adversarial loss:

$$L_{\text{adv}} = -(1 - I) \log p \quad (15)$$

Dataset	$ Y $	Depth	Avg($ y_i $)	Train	Dev	Test
WOS	141	2	2.0	30,070	7,518	9,397
NYT	166	8	7.6	23,345	5,834	7,292
RCV1-V2	103	4	3.24	20,833	2,316	781,265

Table 1: Data statistics. $|Y|$ is the number of
classes. Depth is the maximum level of hierar-
chy. Avg($|y_i|$) is the average number of classes
per sample.

Notice that L_{adv} only updates the parameters of
the generator but does not affect the parameters
of the encoder or that of the discriminator so that
we only consider $\mathbf{h} = \mathbf{h}_{\text{mix}}$ in Equation 15.

During training, both the generator and the en-
coder take the same text \mathbf{x} as input and outputs
 \mathbf{h}_{mix} and $\hat{\mathbf{h}}_{\text{mix}}$ respectively. The discriminator then
calculates p so that every loss can be computed.
After training, the framework degrades to a sin-
gle generator, which can predict independently as
Equation 7. Since the problem is a multi-label clas-
sification, we select labels that have probabilities
greater than a threshold τ as model predictions.
The whole process is illustrated in Algorithm 1.

5. Experiments

5.1. Experiment Setup

Datasets and Evaluation Metrics We experi-
ment on Web-of-Science (WOS) (Kowsari et al.,
2018), NYTimes (NYT) (Sandhaus, 2008), and
RCV1-V2 (Lewis et al., 2004) datasets for anal-
ysis. The statistic details are illustrated in Table
1. We follow the data processing of previous work
(Zhou et al., 2020) and measure the experimental
results with Macro-F1 and Micro-F1.

Baselines For systematic comparisons, we intro-
duce a variety of HTC baselines.

- **BERT** (Devlin et al., 2019). A widely used
pretrained language model that can serve as
a text encoder. All of the other baselines as
well as our method are based on BERT.
- **HiAGM** (Zhou et al., 2020), **HTCInfoMax**
(Deng et al., 2021), and **HiMatch** (Chen et al.,
2021). These three methods all encode text
and taxonomic hierarchy separately and pro-
pose different migration strategies.
- **HGCLR** (Wang et al., 2022a). HGCLR regu-
lates BERT representation by contrastive learn-
ing. It only introduces nodes of the local hier-
archy but ignores how they are connected.
- **HPT** (Wang et al., 2022b). HPT proposes a
new architecture with prompt tuning. The text

Model	WOS (Depth 2)		RCV1-V2 (Depth 4)		NYT (Depth 8)	
	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1
BERT (Wang et al., 2022a)	85.63	79.07	85.65	67.02	78.24	66.08
BERT+HiAGM(Wang et al., 2022a)	86.04	80.19	85.58	67.93	78.64	66.76
BERT+HTCInfoMax(Wang et al., 2022a)	86.30	79.97	85.53	67.09	78.75	67.31
BERT+HiMatch (Chen et al., 2021)	86.70	81.06	86.33	68.66	-	-
HGCLR (Wang et al., 2022a)	87.11	81.20	86.49	68.31	78.86	67.96
HPT (Wang et al., 2022b)	87.16	81.93	87.26	69.53	80.42	70.42
HBGL (Jiang et al., 2022)	87.36	82.00	87.23	71.07	80.47	70.19
HiBERT	85.77	80.10	86.49	68.82	79.49	68.40
HiBERT + HiAdv †	86.38	80.78	86.74	69.43	79.56	69.30
HPT*	87.08	81.59	86.96	69.25	80.21	70.14
HPT* + HiAdv †	87.20	81.62	87.36	69.62	80.83	70.78

Table 2: F1 scores on 3 datasets. The best results are in boldface. *In order to fit in HiAdv, we modify the predicting strategy of HPT so the result is slightly lower. Hereinafter we still denote it as HPT for simplicity. †Improvements are statistical significant with $p < 0.05$.

encoder of HPT encodes text and the output of the graph encoder simultaneously.

- **HBGL** (Jiang et al., 2022). HBGL introduces local hierarchy in a sequence-to-sequence manner, differing from other works which are discriminative models. The architecture of HBGL is a pure BERT so we cannot adopt it as a backbone.

Implement Details We implement our model using PyTorch. Following previous work (Chen et al., 2021), we use bert-base-uncased as our text encoder. We test on two architectures introduced in Section 3.2. For HPT, since the adversarial framework can only focus on one representation, we modify the predicting strategy so that the model predicts with one single representation instead of predicting layer-wisely. For HiBERT, we adopt Zero-bounded Log-sum-exp & Pairwise Rank-based (Su et al., 2022) that HPT uses, which has proven to be a more suitable loss function than binary cross-entropy for HTC. We use a batch size of 8 to fill one Nvidia RTX 3090 (24G) fully. All other hyperparameters follow HPT (Wang et al., 2022b). For more stable training, we train the model without the adversarial loss for the first epoch and stop training if Macro-F1 on the development set does not increase for 5 epochs. Following Wang et al. (2022b), we select the checkpoint with maximum Macro-F1 on the development set as the final model and report the best results among 5 individual runs.

5.2. Main Results

Table 2 illustrates our main results. Among baseline methods, HBGL and HPT perform similarly on all datasets except the Macro-F1 of RCV1-V2, which we believe is the advantage of sequence-to-sequence generation.

The proposed HiAdv steadily improves the performance of both backbone models on all datasets. WOS has a label hierarchy of two layers and the local hierarchy has only one path, making it the easiest dataset among all tested datasets. As a result, applying HiAdv on HPT has little improvement because it has already learned the hierarchy information well without further treatment. HiBERT is a weaker model than HPT, so it can still benefit from HiAdv on a simple dataset. As for NYT which has a hierarchy of 8 layers and has the most intricate label structure, existing methods may not fully learn that information so that our adversarial framework achieves new state-of-the-art.

As mentioned, we use a variant of HPT by removing the layer-wise prediction. Although information about the depth of hierarchy is lost to some extent, the extra local hierarchy provided by the adversarial framework outweighs the loss except on WOS. WOS is a special case where the hierarchy is too easy to learn, so a handmade rule to force the model to predict layer-wise is sufficient.

5.2.1. Discussion on Model Capacity

Model capacity can affect the performance of HiAdv. HPT is a stronger backbone than HiBERT and it leads to divergent behaviors between these two models.

When applying HiAdv on HiBERT, the improvements of Micro-F1 descend according to the difficulty of label hierarchy while the improvements of Macro-F1 for HiBERT are steady. This phenomenon indicates that the local hierarchy benefits rare classes with insufficient training samples more than sufficiently trained classes when the model capacity is limited. However, for a stronger basic model HPT, besides the WOS dataset which the model has learned well enough, applying HiAdv can improve both Micro-F1 and Macro-F1, showing

Variants	WOS (Depth 2)		RCV1-V2 (Depth 4)		NYT (Depth 8)	
	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1
HiBERT	86.90	80.82	87.50	68.69	79.93	69.84
HiBERT + HiAdv	87.10	81.36	87.83	69.66	80.00	70.68
- w/o adversarial loss	86.57	80.61	87.46	68.42	79.33	69.23
HiBERT + Contrastive learning	86.49	79.46	87.09	68.40	77.46	65.18
HPT	87.68	81.72	88.04	69.41	80.49	71.18
HPT + HiAdv	87.84	82.07	88.40	69.80	80.99	71.73
- w/o adversarial loss	87.40	81.61	88.03	69.34	80.66	71.16
HPT + Contrastive Learning	87.74	81.91	88.13	69.28	80.84	71.69

Table 3: Results of different optimizing strategies on the development set. The best results are in boldface. We test two backbones with 1) no treatment; 2) the proposed HiAdv; 3) the proposed HiAdv but without the adversarial loss; 4) contrastive learning.

that our adversarial framework is adept at dealing with complex structures for powerful models.

5.2.2. Comparison with Other Local Hierarchy Methods

When comparing our method with HGCLR and HBGL which also involves local hierarchy, the results demonstrate that our adversarial framework can better utilize the local hierarchy.

If we use HPT, which ignores local hierarchy completely, as a baseline, both HGCLR and HBGL reveal better performance on WOS which has the simplest hierarchy but behaves relatively poorly on NYT which has the most complex hierarchy. These two methods view the local hierarchy as either flat or a sequence so that they lose more information about structure when the hierarchy is more complex. Besides, HGBL utilizes local hierarchy in a sequence-generation manner so there is a gap between the ground-truth local hierarchy input for training and model-predicted hierarchy input for testing. When the hierarchy is deep, this gap becomes more evident. On the contrary, the promotion of our HiAdv has a positive correlation to the complexity of hierarchy: we observe the most improvement at NYT but nearly no improvement at WOS, showing that HiAdv can better handle complex local hierarchy and learn more from it than previous methods.

5.3. Analysis

5.3.1. Effect of Adversarial Training

We select adversarial training to incorporate local hierarchy but it is not the only choice. Wang et al. (2022a) uses contrastive learning for a similar purpose so we compare HiAdv with it here. As in Table 3, we exhibit results on the development set of removing the adversarial loss and replacing HiAdv with contrastive learning.

After removing the adversarial loss, HiAdv degrades into a data augmentation method where the local-hierarchy-involved data generated by the encoder serve as augment data. Although random improvement can be observed, in most cases simple data augmentation is counteractive. Directly introducing local hierarchy will leak the ground-truth labels so that models tend to rely on the local hierarchy instead of learning from it without further guidance.

Contrastive learning produces promising results with HPT but fails to work with HiBERT. Contrastive learning directly optimizes the target representation and pulls it toward the oracle representation. Although robust models like HPT may gain from this optimization, weak models such as HiBERT experience countereffects. On the contrary, adversarial training is a more gentle approach that only requires two representations to be similar enough to fool a discriminator. As a result, HiAdv is a more robust framework that can fit in more backbones.

5.3.2. Effect of Local Hierarchy

We further study how the local hierarchy affects HiAdv on the NYT dataset due to its complex hierarchy. Table 4 enumerates the results of some variants on the development set of NYT. We study how the model reacts when providing a partial local hierarchy, no local hierarchy, and a wrong local hierarchy. We implement the partial local hierarchy by randomly dropping 15% of labels and the wrong local hierarchy by randomly selecting the same amount of labels.

As is shown, our framework indeed learns from the local hierarchy. With it corrupted, results of both HiBERT and HPT drop significantly, showing that the promotion of HiAdv indeed comes from the local hierarchy. As for removing the hierarchy partially or completely, variants with more local hierarchy have better performance. This observation is in accordance with the deduction from the main

Variants	Micro-F1	Macro-F1
HiBERT + HiAdv	80.00	70.68
- w/ partial local hierarchy	79.51	69.64
- w/ no local hierarchy	78.46	68.09
- w/ wrong local hierarchy	78.67	67.02
HPT + HiAdv	80.99	71.73
- w/ partial local hierarchy	80.86	71.50
- w/ no local hierarchy	77.82	64.14
- w/ wrong local hierarchy	77.09	61.87

Table 4: Results of modifying local hierarchy on the development set of NYT dataset. The best results are in boldface. We show the original results and the results when we input partial, empty, and wrong local hierarchy.

results that HiAdv can learn more from complex hierarchy than simple one. With no local hierarchy or wrong hierarchy, the performance of HPT drops dramatically, which demonstrates that HPT is more sensitive to local hierarchy and thus more suitable for our framework from another perspective.


Besides, changing the local hierarchy has a relatively lower impact on Micro-F1 than on Macro-F1, which we have observed similar behaviors in the main results. These results further demonstrate that the local hierarchy elevates the performance of rare classes more than normal classes.

5.3.3. Results on Imbalanced Hierarchy

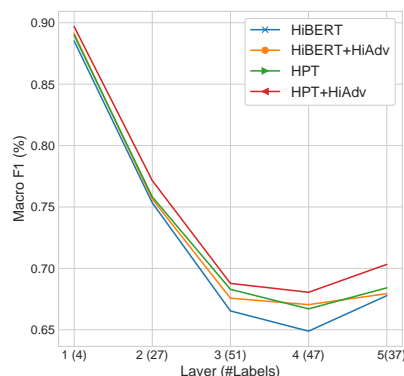
One of the key challenges of hierarchical text classification is the imbalanced label hierarchy. As mentioned, our method affects Macro-F1 more than Micro-F1, showing that it can alleviate the imbalance issues to some extent. In this section, we visualize how our model resolves the issue of imbalance on the development set of NYT.

The imbalance can be viewed from two perspectives in HTC. For one, the number of training instances of each class is different so some classes may not be fully trained. As shown in Figure 4a, we cluster labels into five bunches based on the number of training instances. HiAdv is effective to all clusters and can promote rare classes such as "<20%" cluster for HPT and "20%-40%" cluster for HiBERT.

For another, the number of labels at different hierarchy layers is different. As shown in Figure 4b, HiAdv mainly acts on deep layers (layers ≥ 4). The structure of label hierarchy is more complicated in the deep so this phenomenon demonstrates that our framework can utilize complex local hierarchy effectively.



(a)



(b)

Figure 4: Macro F1 scores of label clusters on the development set of NYT. (a) Label clusters grouped by the number of training samples. >80% means this cluster of labels has training instances of more than 80% of labels. The rest are arranged similarly. (b) Label clusters grouped by depth in the hierarchy.

5.3.4. Time Efficiency

HiAdv requires an additional graph encoder so it requires extra computational costs. Since our approach is a framework, the extra computational costs heavily depend on the backbone model. In this section, we study how HiAdv affects the computational costs of different backbones.

In Table 5 we exhibit the rough training time consumed by one epoch on the NYT dataset. The straightforward implementation of HiAdv requires roughly double the training time for both backbones. Compared to a single model that needs a one-time back-propagation, our framework requires two-time back-propagations for adversarial training. But from Figure 3 it can be observed that both passes of back-propagation share the same BERT encoder, which contains most of the parameters so an optimization algorithm is easy to find.

For HiBERT where BERT only takes text as input (Equation 1), the BERT encoder only needs one-time back-propagation exactly. As a result, the optimized HiAdv requires around 7% of extra

Variants	HiBERT	HPT
w/o HiAdv	560	690
w/ HiAdv, w/o optimization	1020	1380
w/ HiAdv, w/ optimization	600	810

Table 5: Rough seconds consumed on one epoch on the train set of NYT.

training time. For HPT where BERT takes graph information as input, the second pass is necessary. But the input sequence of BERT is still mostly the same on both passes so it is possible to merge those two sequences into one. With that optimization implemented, HiAdv only requires less than 20% more training time.

During inference, our framework degrades into a single backbone model so it does not affect inference time. In practice, applying our framework needs no extra computational cost after deployment.

6. Conclusion

In this paper, we propose a hierarchy-aware adversarial framework (HiAdv) for assisting existing HTC models to incorporate local hierarchy. For any HTC model that involves a text encoder and a graph encoder, HiAdv treats it as a generator while an extra encoder encodes the local hierarchy into an oracle representation. A discriminator tries to distinguish the original representation from the oracle representation while the generator attempts to fool the discriminator. Experiments show that the adversarial framework is adept at dealing with complex hierarchies or promoting weak models that cannot fully learn the hierarchy. Further experiments demonstrate that the effect of HiAdv comes from the local hierarchy and the local hierarchy is beneficial for classes with deficient training instances.

7. Acknowledgements

We thank all the anonymous reviewers for their constructive feedback. This work was supported by the National Natural Science Foundation of China (62036001) and the National Science and Technology Major Project (2022ZD0116308).

8. Bibliographical References

Yekun Chai, Qiyue Yin, and Junge Zhang. 2023. [Improved training of mixture-of-experts language gans](#). In *ICASSP 2023-2023 IEEE International*

Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 1–5. IEEE.

Boli Chen, Xin Huang, Lin Xiao, Zixin Cai, and Liping Jing. 2020. [Hyperbolic interaction model for hierarchical multi-label classification](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7496–7503.

Haibin Chen, Qianli Ma, Zhenxi Lin, and Jiangyue Yan. 2021. [Hierarchy-aware label semantics matching network for hierarchical text classification](#). 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 4370–4379, Online. Association for Computational Linguistics.

Zhongfen Deng, Hao Peng, Dongxiao He, Jianxin Li, and Philip Yu. 2021. [HTCInfoMax: A global model for hierarchical text classification via information maximization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3259–3265, Online. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. [Generative adversarial networks](#). *Communications of the ACM*, 63(11):139–144.

Xuemeng Hu, Rui Wang, Deyu Zhou, and Yuxuan Xiong. 2020. [Neural topic modeling with cycle-consistent adversarial training](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9018–9030, Online. Association for Computational Linguistics.

Ke Ji, Yixin Lian, Jingsheng Gao, and Baoyuan Wang. 2023. [Hierarchical verbalizer for few-shot hierarchical text classification](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2918–2933, Toronto, Canada. Association for Computational Linguistics.

- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2020. [SMART: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2177–2190, Online. Association for Computational Linguistics.
- Ting Jiang, Deqing Wang, Leilei Sun, Zhongzhi Chen, Fuzhen Zhuang, and Qinghong Yang. 2022. [Exploiting global and local hierarchies for hierarchical text classification](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4030–4039, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Rie Johnson and Tong Zhang. 2015. [Effective use of word order for text categorization with convolutional neural networks](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 103–112, Denver, Colorado. Association for Computational Linguistics.
- Zhiyue Liu, Jiahai Wang, and Zhiwei Liang. 2020. [Catgan: Category-aware generative adversarial networks with hierarchical evolutionary learning for category text generation](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8425–8432.
- Yuning Mao, Jingjing Tian, Jiawei Han, and Xiang Ren. 2019. [Hierarchical text classification with reinforced label assignment](#). 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 445–455, Hong Kong, China. Association for Computational Linguistics.
- Mobashir Sadat and Cornelia Caragea. 2022. [Hierarchical multi-label classification of scientific documents](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8923–8937, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. 2020. [Coldgans: Taming language gans with cautious sampling strategies](#). *Advances in Neural Information Processing Systems*, 33:18978–18989.
- Carlos N Silla and Alex A Freitas. 2011. [A survey of hierarchical classification across different application domains](#). *Data Mining and Knowledge Discovery*, 22(1):31–72.
- Jianlin Su, Mingren Zhu, Ahmed Murtadha, Shengfeng Pan, Bo Wen, and Yunfeng Liu. 2022. [Zlpr: A novel loss for multi-label classification](#). *arXiv preprint arXiv:2208.02955*.
- Boxin Wang, Shuohang Wang, Yu Cheng, Zhe Gan, Ruoxi Jia, Bo Li, and Jingjing Liu. 2021. [Infobert: Improving robustness of language models from an information theoretic perspective](#). In *International Conference on Learning Representations*.
- Rui Wang, Xueming Hu, Deyu Zhou, Yulan He, Yuxuan Xiong, Chenchen Ye, and Haiyang Xu. 2020. [Neural topic modeling with bidirectional adversarial training](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 340–350, Online. Association for Computational Linguistics.
- Rui Wang, Deyu Zhou, and Yulan He. 2019. [Atm: Adversarial-neural topic model](#). *Information Processing & Management*, 56(6):102098.
- Zihan Wang, Peiyi Wang, Lianzhe Huang, Xin Sun, and Houfeng Wang. 2022a. [Incorporating hierarchy into text encoder: a contrastive learning approach for hierarchical text classification](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7109–7119, Dublin, Ireland. Association for Computational Linguistics.
- Zihan Wang, Peiyi Wang, Tianyu Liu, Binghuai Lin, Yunbo Cao, Zhifang Sui, and Houfeng Wang. 2022b. [HPT: Hierarchy-aware prompt tuning for hierarchical text classification](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3740–3751, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Hongqiu Wu, Yongxiang Liu, Hanwen Shi, Min Zhang, et al. 2023. [Toward adversarial training on contextualized language representation](#). In *International Conference on Learning Representations*.
- Jiawei Wu, Wenhan Xiong, and William Yang Wang. 2019. [Learning to learn and predict: A meta-learning approach for multi-label classification](#). 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 4354–4364, Hong Kong, China. Association for Computational Linguistics.

Yuan Wu, Diana Inkpen, and Ahmed El-Roby. 2021. [Conditional adversarial networks for multi-domain text classification](#). In *Proceedings of the Second Workshop on Domain Adaptation for NLP*, pages 16–27, Kyiv, Ukraine. Association for Computational Linguistics.

Yuan Wu, Diana Inkpen, and Ahmed El-Roby. 2022. [Co-regularized adversarial learning for multi-domain text classification](#). In *International Conference on Artificial Intelligence and Statistics*, pages 6690–6701. PMLR.

Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. 2021. [Do transformers really perform badly for graph representation?](#) *Advances in Neural Information Processing Systems*, 34:28877–28888.

Xinyi Zhang, Jiahao Xu, Charlie Soh, and Lihui Chen. 2021. [La-hcn: Label-based attention for hierarchical multi-label text classification neural network](#). *Expert Systems with Applications*, page 115922.

Rui Zhao, Xiao Wei, Cong Ding, and Yongqi Chen. 2021. [Hierarchical multi-label text classification: Self-adaption semantic awareness network integrating text topic and label level information](#). In *International Conference on Knowledge Science, Engineering and Management*, pages 406–418. Springer.

Jie Zhou, Chunping Ma, Dingkun Long, Guangwei Xu, Ning Ding, Haoyu Zhang, Pengjun Xie, and Gongshen Liu. 2020. [Hierarchy-aware global model for hierarchical text classification](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1106–1117, Online. Association for Computational Linguistics.

Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu. 2020. [Freelb: Enhanced adversarial training for natural language understanding](#). In *International Conference on Learning Representations*.

He Zhu, Chong Zhang, Junjie Huang, Junran Wu, and Ke Xu. 2023. [HiTIN: Hierarchy-aware tree isomorphism network for hierarchical text classification](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7809–7821, Toronto, Canada. Association for Computational Linguistics.

9. Language Resource References

Kowsari, Kamran and Brown, Donald E and Heidarysafa, Mojtaba and Meimandi, Kiana Jafari and Gerber, Matthew S and Barnes, Laura E. 2018. *Web of Science Dataset*. Mendeley Data. PID <https://doi.org/10.17632/9rw3vkcfy4.6>.

Lewis, David D and Yang, Yiming and Russell-Rose, Tony and Li, Fan. 2004. *Rcv1: A new benchmark collection for text categorization research*. Goldsmiths, University of London. PID <https://trec.nist.gov/data/reuters/reuters.html>.

Evan Sandhaus. 2008. *The new york times annotated corpus*. Linguistic Data Consortium, Philadelphia, ISLRN 429-488-225-160-9.