

HiGen: Hierarchy-Aware Sequence Generation for Hierarchical Text Classification

Vidit Jain^{*♣}, Mukund Rungta^{†*♣◇}, Yuchen Zhuang[♣], Yue Yu[♣], Zeyu Wang[♣], Mu Gao[♣], Jeffrey Skolnick[♣], Chao Zhang[♣]

[♣]Georgia Institute of Technology, [♣]Yale University

[◇]Microsoft, Cambridge, USA

{vjain312,mrungta8,yczhuang,yueyu,mu.gao,skolnick,chaozhang}@gatech.edu,
zeyu.wang.zw389@yale.edu

Abstract

Hierarchical text classification (HTC) is a complex subtask under multi-label text classification, characterized by a hierarchical label taxonomy and data imbalance. The best-performing models aim to learn a static representation by combining document and hierarchical label information. However, the relevance of document sections can vary based on the hierarchy level, necessitating a dynamic document representation. To address this, we propose HiGen, a text-generation-based framework utilizing language models to encode dynamic text representations. We introduce a level-guided loss function to capture the relationship between text and label name semantics. Our approach incorporates a task-specific pretraining strategy, adapting the language model to in-domain knowledge and significantly enhancing performance for classes with limited examples. Furthermore, we present a new and valuable dataset called ENZYME, designed for HTC, which comprises articles from PubMed with the goal of predicting Enzyme Commission (EC) numbers. Through extensive experiments on the ENZYME dataset and the widely recognized WOS and NYT datasets, our methodology demonstrates superior performance, surpassing existing approaches while efficiently handling data and mitigating class imbalance. We release our code and dataset here: <https://github.com/viditjain99/HiGen>.

1 Introduction

Hierarchical text classification (HTC) is a task that involves categorizing text data into predefined categories organized in a hierarchical structure (Banerjee et al., 2019; Wang et al., 2022a; Zhou et al., 2020). It holds great importance in various text mining applications, including scientific paper recommendation (Zhang et al., 2020; Xu et al., 2023),

semantic indexing (Li et al., 2019), and online advertising (Agrawal et al., 2013). HTC poses unique challenges when compared to traditional text classification, as it deals with imbalanced data distributions and complex dependencies between multiple levels of categories within the hierarchy. The hierarchical structure is represented as a directed acyclic graph (DAG), which must be encoded in the predictive model along with the text to generate the final hierarchical label. However, the data imbalance becomes more pronounced as we move down the levels, presenting a significant challenge for HTC.

Existing methods for hierarchical text classification (HTC) can be categorized into three groups: global (Zhou et al., 2020; Chen et al., 2021), local (Shimura et al., 2018; Banerjee et al., 2019; Wehrmann et al., 2018) and generative (Risch et al., 2020; Yu et al., 2022; Huang et al., 2022). Local approaches predict each hierarchical level using independent classifiers. Global models, on the other hand, use a single classifier and incorporate hierarchical information into the loss function. Finally, generative approaches use text generation frameworks to model the hierarchical structure. Some prior works flatten the label structure, leading to an exponential increase in the number of classes and the loss of hierarchical dependency. Additionally, some models fail to capture the correlation between text and label name semantics. As a result, these approaches struggle to perform well on long-tailed classes with limited training data.

Harnessing the power of large pretrained language models (PLMs) to capture text-label correlation, we employ a transformer-based sequence-to-sequence (seq2seq) framework (Lewis et al., 2020). These models are pretrained on extensive text data, enabling them to encode transferable linguistic features across tasks (Liu et al., 2019). Our model transforms HTC into a text generation problem, generating labels conditioned on the input text and previously generated labels while utilizing a hier-

*Equal contribution

† Work done by the author while at Georgia Institute of Technology

archical level-guided semantics framework. While [Chen et al. \(2021\)](#) introduces a matching loss for hierarchy-aware text-label relationships, it neglects the level-dependency between the transformed text and label representations. Our approach incorporates a margin loss, aligning text semantics with positive labels at each taxonomic level while pushing negative semantics apart. Furthermore, we propose a token constraint loss during training to discourage undesired token generation.

Inspired by the success of pretraining tasks ([Gururangan et al., 2020](#)), we propose a task-specific pretraining strategy for our generative model. Our approach presents a notably simplified design in contrast to previous works which use complex encoder architectures or training paradigms ([Zhou et al., 2020](#); [Deng et al., 2021](#); [Jiang et al., 2022](#)). We propose a pretraining step that takes advantage of weak supervision to jointly model text representations and hierarchical label information. The encoder encodes the input document and masked hierarchical label, while the decoder regenerates the original label. The pretraining step leverages weakly labeled data from the same domain, generated using an LLM. This domain-specific knowledge, discussed in [Mueller et al. \(2022\)](#); [Gururangan et al. \(2020\)](#), proves advantageous for addressing the data imbalance challenge commonly encountered in HTC datasets.

The major contributions of this work are:

- We propose a generation-based HTC framework that effectively captures document-label dependencies across levels using a level-guided semantic loss.
- We devise an efficient pretraining strategy that leverages in-domain data to align the language model with the target task and domain.
- Our approach demonstrates remarkable performance on classes with limited examples, surpassing prior works even with minimal training instances.
- We present ENZYME, a dataset of 30,523 full-text PubMed articles with Enzyme Commission (EC) numbers. It features a four-level single-path hierarchy, making it larger than any existing datasets for biomedical HTC.

2 Related Work

HTC is a multi-label classification problem where the classification task is performed in a hierarchical

manner. Most prior work in HTC can be categorized into three broad categories: local, global, and generative. In the local approach, each hierarchical level has its own classifier tailored to the unique classes at that level. Conversely, the global approach employs a unified classifier that encompasses all classes across every hierarchy level. The generative approach, a recent advancement in hierarchical text classification, capitalizes on the generative capabilities of language models to predict text labels. More details on each of the categories are explained below.

Local HTC Local approaches in HTC use local classifiers at each level or class ([Banerjee et al., 2019](#); [Wehrmann et al., 2018](#); [Shimura et al., 2018](#); [Peng et al., 2018](#)). [Banerjee et al. \(2019\)](#) initialize binary classifiers at lower levels with parameters from the parent classifier. [Wehrmann et al. \(2018\)](#) combine local and global losses to encode information within and across hierarchical levels. [Shimura et al. \(2018\)](#) address data imbalance with parameter transfer techniques. [Peng et al. \(2018\)](#) employs a Graph-CNN-based model with recursive regularization for deep hierarchical representations.

Global HTC Global approaches to HTC utilize a single classifier ([Gopal and Yang, 2013](#); [Wu et al., 2019](#); [Mao et al., 2019](#); [Peng et al., 2021](#)) to predict labels at different hierarchy levels. Early works consider parent-child dependencies ([Gopal and Yang, 2013](#); [Wu et al., 2019](#); [Mao et al., 2019](#); [Peng et al., 2021](#)), while recent approaches focus on global label structure ([Wang et al., 2021](#)), disjoint features ([Zhang et al., 2022](#)), label imbalance ([Deng et al., 2021](#)), prior hierarchy knowledge ([Zhou et al., 2020](#)), and semantic matching ([Chen et al., 2021](#)). These works learn text and hierarchy semantics separately, fusing them later. However, ([Wang et al., 2022a](#)) proposes a global approach using contrastive learning to learn a shared representation. Similarly, ([Jiang et al., 2022](#)) also aims for common representations but incorporates both local and global hierarchies. Lastly, ([Wang et al., 2022b](#)) uses prompt-tuning and multilabel MLM to learn shared semantics.

Generative HTC Early works on text generation for HTC ([Yang et al., 2018](#); [Risch et al., 2020](#)) use RNN and Transformer-based seq2seq models, with dynamic document representations outperforming static encoder methods. Recent approaches ([Yu et al., 2022](#); [Huang et al., 2022](#)) propose T5-based models ([Raffel et al., 2020](#)). The former addresses label inconsistency with DFS-based linearization

and constrained decoding, while the latter captures dependencies using BFS-based linearization and hierarchical path-based attention. These methods overlook the dependencies between the label name from different levels and the document text. We overcome this by capturing the document-label-name dependency across different levels using the proposed level-guided semantic loss.

In this work, we explore HTC under the umbrella of a sequence generation framework using a BART-based model (Lewis et al., 2020). We harness the benefits of denoising autoencoder pretraining to imbue the model with a strong understanding of the hierarchical structure. We also employ a well-designed objective function during supervised training for improved learning.

3 Problem Definition

Given an input text $X_i = \{x_1, x_2, \dots, x_n\}$, HTC aims to classify the text into a subset Y_i of label set Y . The label set Y is arranged as a Directed Acyclic Graph (DAG), denoted by $H = (L, E)$. L represents the set of nodes and E represents the edges indicating the nodes' parent-child relations. All the labels in Y constitute the nodes in the above graph i.e. L . We use BFS (Bundy and Wallen, 1984) to flatten the hierarchical labels into a multi-level sequential string of label nodes. We define a special set of symbols $S = \{/, \langle root \rangle\}$ to demarcate special relations of the label hierarchy in the flattened sequence. To classify X_i , the proposed sequence generation model generates $Y_i = \{y_1, y_2, \dots, y_k\}$, where $y_j \in \{Y + S\}$. BFS-based linearization helps each Y_i to correspond to one or more paths in root to leaf paths in H .

4 Methodology

This section presents the proposed generative framework based on the BART model for HTC. We first describe the pretraining regime aimed at learning robust joint text-label representations (Section 4.1). Following this, supervised training is performed to generate hierarchical label sequences for input documents (Section 4.2). We also explain the proposed objective function used (Section 4.3) for this phase that aims to capture text and label semantics jointly.

4.1 Pretraining Strategy

Pretraining a domain-specific language model has shown significant performance improvement

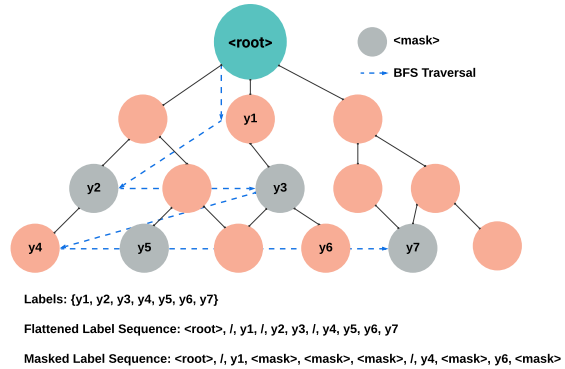


Figure 1: BFS-based label flattening and random token & span masking employed during pretraining.

on downstream tasks (Gururangan et al., 2020; Mueller et al., 2022). We utilize the BART model which has achieved state-of-the-art results on text generation tasks, as the backbone of our seq2seq approach. Its autoregressive nature and availability as a pre-trained model alleviate the reliance on extensive labeled data. The BART model comprises a transformer-based encoder and an autoregressive decoder.

To design our model, we adopt a BART-style denoising auto-encoder inspired by prior work (Aghajanyan et al., 2021). First, we transform a label set Y_i into a multi-level sequential label Y_i^{seq} using BFS, where $\langle root \rangle$ represents the root node and $/$ indicates the change of level. Our approach involves randomly masking certain levels in the hierarchical label and encoding them with the input text. Figure 1 provides an illustration of the label flattening and random masking techniques. The model is then trained to reconstruct the original hierarchical label. Formally, given a document X_i and masked label $Y_i^{masked_seq}$, we create an input sequence as follows:

$$input = [X_i \langle /s \rangle Y_i^{masked_seq}] \quad (1)$$

where $\langle /s \rangle$ is a special token used as a separator and $Y_i^{masked_seq}$ is the masked label sequence resulting from the masking process as illustrated in Figure 1. We encode both the text and label using the same encoder, allowing the model to learn a joint embedding in the text-label space and capture correlations between them. During training, the model generates an output for the masked input sequence, aiming to fill in the masked positions of the label:

$$output = \hat{Y}_i^{seq} \quad (2)$$

We use the cross-entropy loss, commonly used in masked language modeling, to train the model:

$$loss = CE(Y_i^{seq}, \hat{Y}_i^{seq}) = - \sum_i Y_i^{seq} \log(\hat{Y}_i^{seq})$$

where CE represents the cross-entropy loss, \hat{Y}_i^{seq} represents the model’s output probability distribution over the vocabulary.

The pretraining step aims to learn a joint representation for text and labels, capturing their interdependencies. The model is exposed to document-label pairs with partially masked hierarchical labels to develop a robust understanding of the hierarchy structure. Importantly, the model develops robust representations for all levels of the hierarchy, regardless of the number or levels of the masked nodes, as the extent of the masked label nodes is not known. The pretraining dataset is further described in Section 5.1.1.

4.2 Sequence to Sequence Modeling

The pretraining enhances the model’s understanding of the label hierarchy and domain-specific knowledge. In the supervised training phase, we utilize the seq2seq framework for HTC, as shown in Figure 2. We use the BFS label linearization to obtain multi-level label sequences. The input text document $X_i = \{x_1, x_2, x_3, \dots, x_m\}$ is encoded to obtain the encoder output hidden representation h^e .

$$h^e = \text{Encoder}(X_i) \quad (3)$$

The hidden representations from the encoder are used to initialize the decoder. The decoder then generates the hierarchical label \hat{Y}_i step-by-step autoregressively. The autoregressive process followed by the decoder can be represented as:

$$p(\hat{Y}_i | X_i) = \prod_{k=1}^n p(\hat{Y}_i^k | X_i, \hat{Y}_i^{<k}) \quad (4)$$

where \hat{Y}_i^k denotes the prediction from the decoder for level k of the hierarchy, and n represents the depth of the hierarchy. At every time step k , the output \hat{Y}_i^{k-1} and the hidden state h_{k-1}^d from the previous time step $k-1$ is given as input to the decoder to generate the next hidden state h_k^d and a prediction \hat{Y}_i^k for the current time step.

$$h_k^d, \hat{Y}_i^k = \text{Decoder}(h^e, h_{k-1}^d, \hat{Y}_i^{k-1}) \quad (5)$$

The decoder leverages information from the encoded document and the label from the previous

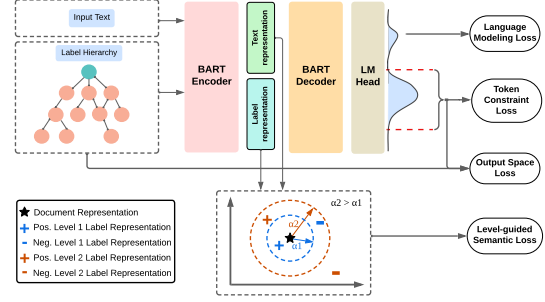


Figure 2: The proposed model’s architecture consists of an encoder that takes the document and its corresponding label name as input. The decoder generates a hierarchical label output with 2 levels. To calculate the first three losses, the LM Head predicts a distribution over the vocabulary, and the hierarchy edges are considered. For the semantic loss, the text and label name representations from the encoder are projected onto a shared embedding space. Positive document and label semantics are pulled together, while negatives are pushed apart. The margins α_1 and α_2 control the attraction between levels 1 and 2, with $\alpha_1 > \alpha_2$.

level to predict the current level in the hierarchy. This allows our model to mimic the HTC process and learn important aspects of the hierarchy structure, including label relations, valid root-to-leaf node paths, and the overall label structure (Risch et al., 2020).

4.3 Training Objective

This section presents the proposed objective function we use to train our generative framework for the proposed HTC task. First, we provide a description of each loss function being employed and then define the final training objective used to train our model.

Language Modeling Loss. The predictions from each time step of the decoder \hat{Y}_i^k are concatenated together to form the final prediction of the model \hat{Y}_i . The ground truth, Y_i comprises the original flattened label. The language modeling loss for HiGen can be expressed as:

$$\mathcal{L}_{LM} = \text{crossentropy}(Y_i, \hat{Y}_i) \quad (6)$$

Output Space Loss. The label hierarchy is generally represented as a DAG, where each edge signifies a parent-to-child relationship. In HTC, imposing this unidirectionality during training helps the model comprehend these hierarchical relations. To this end, we use the formulation proposed by Zhang et al. (2021) inspired by the distributional inclusion

hypothesis (DIH) (Geffet and Dagan, 2005).

$$\mathcal{L}_O = \sum_N \sum_l \sum_n \max(0, \pi_c - \pi_p) \quad (7)$$

where N is the batch size, l is the number of predicted node labels and n is the number of hierarchy edges. π_p and π_c represent the predicted probabilities for a parent label node p and its child label node c in the output distribution over the vocabulary.

Under the DIH framework, the loss term can be precisely articulated as follows: if a document d belongs to a child class c with probability π_c , then it must belong to the parent class p with a probability no less than π_c . For instance, if a document has a 75% chance of being labeled with "Football", the likelihood of it being assigned to the parent category "Sports" should be 75% or higher. Note that the loss term exhibits asymmetry, being non-zero when $\pi_c > \pi_p$ but zero when $\pi_p > \pi_c$. Referring to Figure 2, loss calculation involves aggregating predicted probabilities for valid tokens (indicated by red dotted lines) from the LM Head and establishing parent-child pairs using the hierarchy edges. Through the imposition of a penalty when $\pi_c > \pi_p$, the objective function guides the model to learn the correct sequence of tokens, consequently learning the unidirectionality and orientation of the hierarchy edges.

Token Constraint Loss. Using a model trained for open-ended text generation in HTC can lead to irrelevant node labels due to the generation of stray tokens (non-existent in a given hierarchy). To address this and avoid restricting BART's generation capabilities due to a fixed vocabulary, instead, we introduce a loss function that penalizes predictions outside of a designated vocabulary represented by the red dotted lines in Figure 2. This vocabulary is constructed using the label hierarchy, ensuring alignment and enhancing overall results. Putting it formally,

$$\begin{aligned} V^{H'} &= \{l \in V | l \notin V^H\} \\ \mathcal{L}_T &= \sum_{l \in V^{H'}} \pi_l^i \end{aligned} \quad (8)$$

where l is a token, π^i represents a vector of predicted token probabilities for a training example by the LM Head, V is the entire vocabulary, V^H is the desired vocabulary for H , and $V^{H'}$ contains the remaining tokens. The loss function guides the model to learn the desired vocabulary's contents,

discouraging the generation of irrelevant labels and enhancing classification performance.

Level-guided Semantic Loss. Based on the previous loss functions, the model learns correlations between text semantics and the hierarchical structure. Label names serve as descriptions encoding paths in the hierarchy, providing valuable semantic information for HTC. For any classification task, the text and label semantics for positive pairs should be closer than unrelated pairs in an embedding space. Our training objective imposes this constraint, guided by label nodes at each hierarchical level.

In a traditional classification setting, document similarity can be determined by matching their labels. However, in HTC, some documents might be related to each other for some initial levels of the hierarchy but then diverge as we move down and vice versa. Capturing this nuance is crucial for HTC and is the prime distinction from traditional classification. We aim to encode this information using our level-guided semantic loss function.

We utilize two independent fully connected networks and project the text and label semantics onto a common embedding space.

$$\begin{aligned} E_t &= FC_t(h_t^e) \\ E_l &= FC_l(h_l^e) \end{aligned} \quad (9)$$

where FC_t & FC_l are two fully-connected networks, h_t^e & h_l^e are the encoder hidden representation for the document and label name. $E_t, E_l \in \mathcal{R}^{N \times d}$ represent the text and label name representations in the joint space and N is the batch size.

To combine the hierarchical information with text and label semantics, we construct document-label name pairs for every level of the hierarchy. So for a batch of data and a particular level k from the hierarchy, the pairing process involves associating each document embedding with every label name embedding in the batch. Pairs receive a positive label (1) when the document and label name embeddings correspond to the same level label; otherwise, a negative label (0) is assigned. Consequently, a positive document-label name embedding pair is denoted as $\{E_{t_i}^+, E_{l_j}^+\}$, while a negative pair is denoted as $\{E_{t_i}^-, E_{l_j}^-\}$. Intra-pair distances are calculated amongst the positive and negative document-label pairs using L2-normalized Euclidean distance. We use the mean of intra-document-label pair distances for both the positive and negative variants

for further calculations. Mathematically,

$$\begin{aligned}\gamma_k^+ &= \frac{1}{N^+} \sum_i \sum_j \|E_{t_i}^+ - E_{l_j}^+\| \\ \gamma_k^- &= \frac{1}{N^-} \sum_i \sum_j \|E_{t_i}^- - E_{l_j}^-\|\end{aligned}\quad (10)$$

To bring the positive document and label name semantics closer, a margin loss is applied with a level-specific margin parameter α_k . So for the level k , the loss is defined as,

$$\mathcal{L}_S^k = \max(0, \gamma_k^+ - \gamma_k^- + \alpha_k) \quad (11)$$

At every hierarchy level, we replicate the pairing and loss calculation using the outlined strategy. Notably, with level transitions, documents initially labeled as negatives may now share labels and vice versa, creating a fresh set of labeled pairs. The level-specific margin α_k progressively increases as we descend the hierarchy, to incorporate the increasing granularity of labels. This ensures the model aligns text and label semantics effectively. For a label hierarchy with m levels, the final loss becomes,

$$\mathcal{L}_S = \sum_{k=1}^m \mathcal{L}_S^k \quad (12)$$

In Figure 2, the lower part illustrates an example of a two-level hierarchy, with different colors representing the hierarchy levels. For level 1, positive label semantics are attracted to document semantics, while negative label semantics are pushed apart by at least α_1 (blue-dotted circle). The same process is applied to the next level, with increased separation α_2 to accommodate the higher semantic granularity (orange-dotted circle).

Based on the language modeling loss and the three proposed loss functions, we use the following objective function to learn the parameters of our model:

$$\min \mathcal{L}_{\text{HiGen}} = \mathcal{L}_{LM} + \lambda_1 \mathcal{L}_O + \lambda_2 \mathcal{L}_T + \lambda_3 \mathcal{L}_S$$

where, λ_1 , λ_2 and λ_3 are hyperparameters.

5 Experiments

5.1 Datasets

We introduce the **ENZYME** dataset, containing biomedical scientific literature and corresponding Enzyme Commission (EC) numbers. We provide an overview of this dataset, including the training,

Dataset	$ L $	Depth	$Avg(L_i)$	Train	Val	Test
ENZYME	4566	4	4	17422	8741	4360
WOS	141	2	2	30070	7518	9397
NYT	166	8	7.6	23345	5834	7292

Table 1: Statistics of datasets used. $|L|$: total number of target classes, Depth: levels of the hierarchy, $Avg(|L_i|)$: Average number of classes per example.

validation, and test splits, in Section 5.1.2. We also conduct experiments on benchmark datasets: Web-of-Science (WOS) (Kowsari et al., 2017) and NYT (Sandhaus, 2008), following the preprocessing and data splits proposed by Zhou et al. (2020). WOS and ENZYME focus on single-path HTC, while NYT incorporates multi-path taxonomic labels. Detailed statistical information can be found in Table 1. Experimental results are evaluated using Macro-F1 and Micro-F1 metrics, commonly used in prior literature.

5.1.1 Pretraining Datasets

For the ENZYME dataset, we use the articles extracted from PubMed¹. We randomly sample 200,000 articles along with their Enzyme Commission (EC) numbers that follow a hierarchical structure. These articles are loosely labeled as they do not reflect human-annotated EC numbers. WOS hosts a comprehensive collection of scientific articles spanning various domains of science. To produce meaningful and diverse abstracts, we employ the powerful ChatGPT model and generate ~ 3000 abstracts for pretraining. More details are mentioned in Appendix E. As for NYT, we have access to a vast repository of articles that were not assigned to specific training, testing, or validation sets by Zhou et al. (2020). This invaluable resource enables us to utilize these uncategorized articles for our proposed pretraining task. During the process of consolidating the pretraining dataset, we made sure to prevent any overlap of the pretraining data with the training, validation and testing sets.

5.1.2 ENZYME Dataset

We introduce a new dataset called ENZYME, which contains curated full-text biomedical articles from PubMed along with Enzyme Commission (EC) numbers (see Section A.1) and enzyme names. It consists of 30,523 articles in both PDF and parsed formats, making it unique in providing full-text biomedical documents with corresponding enzyme identification numbers. The EC numbers

¹<https://pubmed.ncbi.nlm.nih.gov/advanced/>

Models	ENZYME		WOS		NYT	
	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1
BERT	82.16	16.29	85.63*	79.07*	78.24*	65.62*
HiAGM (Zhou et al., 2020)	80.52	49.21	85.82	80.28	74.97	60.83
HiMatch (Chen et al., 2021)	72.37	38.19	86.20	80.53	74.62	59.28
HiMatch + BERT	75.24	40.67	86.70*	81.06*	76.79*	63.89*
HTCInfoMax (Deng et al., 2021)	70.56	37.24	85.58	80.05	-	-
HTCInfoMax + BERT	73.45	39.83	86.30*	79.97*	78.75*	67.31*
HGCLR (Wang et al., 2022a)	90.81	71.03	87.11	81.20	78.86	67.96
HPT (Wang et al., 2022b)	91.04	78.27	87.16	81.93	80.42	70.42
HBGL (Jiang et al., 2022)	91.10	83.05	87.36	82.00	80.47	70.19
Vanilla BART	88.11	67.76	86.26	79.34	80.08	69.3
SGM-T5 [†] (from Yu et al., 2022)	-	-	85.83	80.79	-	-
Seq2Tree-T5 [†] (Yu et al., 2022)	-	-	87.20	82.50	-	-
PAAM-HiA-T5 [†] (Huang et al., 2022)	-	-	90.36	81.64	77.52	65.97
HiGen (ours)	92.61	84.15	87.39	81.45	80.89	72.41

Table 2: Experimental results of our proposed approach on all datasets. †: Implementation not available, *: results from Wang et al. (2022a).

follow a hierarchical structure², with detailed statistics per level shown in Table 4. The dataset has a hierarchical taxonomy depth of 4, resulting in a complex structure with a large number of fine-grained classes, making this an unique and challenging dataset. This dataset is highly imbalanced with $\sim 50\%$ of the 4,566 classes having less than 2 examples (Figure 5). For our experiments, we focus on classes with 5 or more examples. Further details on the enzyme classification system and dataset construction are available in Appendix A.

5.1.3 Baselines

To compare the proposed method, we select a few recent baselines. HBGL (Jiang et al., 2022), SGM-T5 (Yang et al., 2018), PAAM-HiA-T5 (Huang et al., 2022), Seq2Tree-T5 (Yu et al., 2022), HGCLR (Wang et al., 2022a), HPT (Wang et al., 2022b), HiMatch (Chen et al., 2021) and HiAGM (Zhou et al., 2020), and HTCInfoMax (Deng et al., 2021). SGM-T5, PAAM-HiA-T5 and Seq2Tree-T5 have used a similar sequence generation approach for HTC. SGM uses a T5 encoder-decoder framework, while Seq2Tree-T5 proposes a tree-like framework with label linearization. HBGL leverages BERT’s large-scale parameters and language knowledge to model global and local hierarchies. HiAGM, HTCInfoMax and HiMatch incorporate fusion strategies to integrate text and hierarchy representations. HiAGM introduces hierarchy-aware multi-label attention, HTCInfoMax employs information maximization for modeling text-hierarchy interaction, and HiMatch matches text and label representations in a joint embedding space for clas-


sification. We also compare our approach with Vanilla-BART and a BERT-based HTC model.

5.1.4 Implementation Details

For HiGen, we use HuggingFace checkpoints to warm start the BART model, selecting between a PubMed-finetuned BART model (mse30/bart-base-finetuned-pubmed) for ENZYME and the base-BART model (facebook/bart-base) for WOS and NYT datasets. After an initial pretraining phase using the data in Section 5.1.1, we save the checkpoints and subsequently fine-tune them for all datasets. For WOS, we set semantic margins α_k to [0.05, 0.1], while for ENZYME, they are [0.02, 0.1, 0.15, 0.3]. Due to label sequence complexity, semantic loss is not applied to NYT at present. We use loss balancing factors λ_1 and λ_2 of [1e-3, 1e-6] for ENZYME, and [1e-3, 1e-5] for WOS and NYT, with λ_3 set to 1 for ENZYME and WOS. These values result from extensive hyperparameter tuning. Both pretraining and fine-tuning use a batch size of 12, the Adam optimizer, and a learning rate of 5e-5 with a linear schedule.

We employ consistent evaluation across all datasets by using the baseline implementations provided by their respective authors. For BERT, we train a multilabel classification model using the representation of the special [CLS] token. Like HiGen, BERT is warm started with a PubMed-finetuned checkpoint (microsoft/BiomedNLP-PubMedBERT-base-uncased-abstract-fulltext) for ENZYME, and base-BERT (bert-base-uncased) for WOS and NYT. All models are implemented in PyTorch.

²<https://www.enzyme-database.org/contents.php>



(a) Micro-F1 v/s Frequency of Classes (b) Macro-F1 v/s Frequency of Classes

Figure 3: Performance on ENZYME dataset for the long-tailed classes

5.2 Experimental Results


Table 2 presents the Micro and Macro F1 scores for the baselines and HiGen on the three datasets. The baselines are grouped into two categories: classification-based and generation-based. HiGen, our generation-based approach, achieves superior performance over both categories. Among the classification-based methods, HBGL stands out as the state-of-the-art approach, leveraging global and local hierarchies for text and label representation. To ensure a fair comparison on ENZYME, we initialize HiMatch + BERT, HGCLR, and HBGL with the PubMed BERT checkpoint. HiGen, utilizing dynamic text and label semantics, improves the Micro-F1 score by 1.51%, 0.03%, and 0.42% for ENZYME, WOS, and NYT respectively. Notably, HiGen exhibits exceptional performance on the ENZYME dataset for both Micro and Macro-F1 scores. A significant improvement over the Vanilla-BART establishes the power of proposed pertaining and loss functions. Larger improvement can be seen for dataset involving deeper hierarchy (ENZYME and NYT) compared to shallower hierarchy (WOS).

HiGen uses the BART model (140 million parameters) instead of the T5 model (220 million parameters), making it 37% more parameter-efficient. HiGen outperforms Seq2Tree-T5 for WOS. For NYT, HiGen surpasses PAAM-HiA-T5 by 3.47% and 6.44% on Micro and Macro-F1 scores, respectively. Performance figures on ENZYME are not available due to code unavailability.

In Appendix C.3, we juxtapose the training durations of HiGen and HBGL, revealing a 10-fold reduction in training time for HiGen, attributable to its simpler architecture, as evidenced in Table 8.

5.3 Performance on Long-Tailed Distribution

In our experiments, data imbalance is evident, with *long-tailed classes* having limited training and testing examples. These classes are crucial for evalu-



(a) Micro-F1 v/s Data proportion (b) Macro-F1 v/s Data proportion

Figure 4: Performance on ENZYME dataset on varying the training data proportion

ating the model’s performance, testing its ability to learn from sparse data, and assessing generalization capabilities. We select classes based on testing data frequency, grouping them into bins from 1 to 5.

Figure 3 presents results for the ENZYME dataset, displaying Micro-F1 (Figure 3a) and Macro-F1 (Figure 3b) scores against class frequency. We compare HiGen to baseline models BERT, HBGL, and Vanilla BART, with a focus on HBGL as the best-performing baseline. HiGen outperforms BERT and Vanilla BART significantly in both Micro and Macro-F1. While its Micro-F1 is similar to HBGL, HiGen excels in Macro-F1, particularly for "long-tailed classes" on the ENZYME dataset.

5.4 Data Efficiency

We assessed model robustness by training on different data proportions and evaluating on the original test set. We included baseline models like BERT, Vanilla BART, and HBGL for fair comparisons. Figure 4 displays results for the ENZYME dataset, known for its complex hierarchy and limited training examples. The curves represent model performance as training data increases.

Our model outperformed all baseline models with just 10% of the data, surpassing the second-best model, HBGL, thanks to pretraining knowledge. This synergy of pretraining and fine-tuning improved label hierarchy knowledge, enhancing performance. For more details, see Section 5.5.

5.5 Ablation Study

In this analysis, we evaluate the impact of individual components in our approach on the WOS dataset (Table 3). The importance of the pretraining stage is evident, with its removal leading to a significant drop in performance across both metrics. Even with just 3000 abstracts generated using

Model	Micro-F1	Macro-F1
Vanilla BART	86.26	79.34
HiGen	87.39	81.45
w/o pretraining	86.69	80.54
w/o \mathcal{L}_O	87.35	81.11
w/o \mathcal{L}_T	87.19	80.83
w/o \mathcal{L}_S	87.33	80.91

Table 3: Ablation study on HiGen for the WOS dataset.

LLMs like ChatGPT, our pretraining approach is highly effective, emphasizing their value for data-limited datasets. By refining prompts and adding more pretraining data, we expect further performance improvements. Both the semantic loss (\mathcal{L}_S) and token constraint loss (\mathcal{L}_T) significantly improve the Macro-F1 score by enhancing representation learning for minority classes. While the output space loss (\mathcal{L}_O) also contributes to Macro-F1 improvement, it has a comparatively smaller impact. In summary, all proposed components are crucial; removing any of them leads to a notable drop in performance, as evident when comparing HiGen to a Vanilla BART model. For further results and analyses, please refer to Appendix C.

6 Conclusion

In this paper, we propose a sequence generation framework for HTC that captures the hierarchical nature of labels. The pretraining strategy enhances the model’s performance by adapting it to the specific domain and task. Moreover, we demonstrate the effectiveness of utilizing synthetic data generated through powerful language models like ChatGPT. The objective function during supervised training provides the model with additional contextual information about the hierarchical structure. Our proposed approach outperforms baseline models on three datasets, as measured by standard evaluation metrics. Leveraging the knowledge embedded in pretrained language models, our model performs exceptionally well on classes with limited examples and is data-efficient as well.

Limitations

While HiGen demonstrates performance improvement across three datasets, it is crucial to acknowledge certain limitations and identify potential areas for future improvement.

Firstly, ensuring the appropriate nature of the pretraining data is paramount. It is essential that the data aligns with the same domain as the original dataset, and more importantly, that the labels

in the pretraining data adhere to the same hierarchical structure as the original data. If such data is not readily available, alternative sources like large language models for obtaining weakly supervised pretraining data may need to be utilized.

Secondly, in the level-guided semantic loss, the current approach employs in-batch sampling. However, the selection of the positive and negative samples in a batch might not be optimal. Future works can explore enhancements by incorporating harder negative samples, which could potentially improve the model’s ability to learn more discriminative representations.

Lastly, the present approach involves several hyperparameters which can introduce additional complexity and computational overhead during the fine-tuning process. Future research efforts could focus on streamlining and simplifying these aspects to ensure a more efficient and user-friendly implementation.

Acknowledgement

This work was supported in part by NSF IIS-2008334, IIS-2106961, and CAREER IIS-2144338.

References

- Armen Aghajanyan, Dmytro Okhonko, Mike Lewis, Mandar Joshi, Hu Xu, Gargi Ghosh, and Luke Zettlemoyer. 2021. Htlm: Hyper-text pre-training and prompting of language models. *arXiv preprint arXiv:2107.06955*.
- Rahul Agrawal, Archit Gupta, Yashoteja Prabhu, and Manik Varma. 2013. Multi-label learning with millions of labels: Recommending advertiser bid phrases for web pages. In *Proceedings of the 22nd international conference on World Wide Web*, pages 13–24.
- Siddhartha Banerjee, Cem Akkaya, Francisco Perez-Sorrosal, and Kostas Tsioutsoulouklis. 2019. [Hierarchical transfer learning for multi-label text classification](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6295–6300, Florence, Italy. Association for Computational Linguistics.
- Alan Bundy and Lincoln Wallen. 1984. Breadth-first search. *Catalogue of artificial intelligence tools*, pages 13–13.
- A. Chang, L. Jeske, S. Ulbrich, J. Hofmann, J. Koblitz, I. Schomburg, M. Neumann-Schaal, D. Jahn, and D. Schomburg. 2021. [Brenda, the elixir core data resource in 2021: new developments and updates](#). *Nucleic Acids Res*, 49(D1):D498–d508.

- 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.
- Maayan Geffet and Ido Dagan. 2005. The distributional inclusion hypotheses and lexical entailment. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 107–114.
- Siddharth Gopal and Yiming Yang. 2013. [Recursive regularization for large-scale classification with hierarchical and graphical dependencies](#). In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '13*, page 257–265, New York, NY, USA. Association for Computing Machinery.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A Smith. 2020. Don't stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360.
- Wei Huang, Chen Liu, Bo Xiao, Yihua Zhao, Zhaoming Pan, Zhimin Zhang, Xinyun Yang, and Guiquan Liu. 2022. Exploring label hierarchy in a generative way for hierarchical text classification. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1116–1127.
- Ting Jiang, Deqing Wang, Leilei Sun, Zhongzhi Chen, Fuzhen Zhuang, and Qinghong Yang. 2022. Exploiting global and local hierarchies for hierarchical text classification. *arXiv preprint arXiv:2205.02613*.
- Kamran Kowsari, Donald E. Brown, Mojtaba Heidarysafa, Kiana Jafari Meimandi, Matthew S. Gerber, and Laura E. Barnes. 2017. [Hdltext: Hierarchical deep learning for text classification](#). In *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 364–371.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Keqian Li, Shiyang Li, Semih Yavuz, Hanwen Zha, Yu Su, and Xifeng Yan. 2019. Hiercon: Hierarchical organization of technical documents based on concepts. In *2019 IEEE International Conference on Data Mining (ICDM)*, pages 379–388. IEEE.
- Nelson F Liu, Matt Gardner, Yonatan Belinkov, Matthew E Peters, and Noah A Smith. 2019. Linguistic knowledge and transferability of contextual representations. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1073–1094.
- 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.
- Aaron Mueller, Jason Krone, Salvatore Romeo, Saab Mansour, Elman Mansimov, Yi Zhang, and Dan Roth. 2022. Label semantic aware pre-training for few-shot text classification. *arXiv preprint arXiv:2204.07128*.
- Hao Peng, Jianxin Li, Yu He, Yaopeng Liu, Mengjiao Bao, Lihong Wang, Yangqiu Song, and Qiang Yang. 2018. Large-scale hierarchical text classification with recursively regularized deep graph-cnn. In *Proceedings of the 2018 world wide web conference*, pages 1063–1072.
- Hao Peng, Jianxin Li, Senzhang Wang, Lihong Wang, Qiran Gong, Renyu Yang, Bo Li, Philip S. Yu, and Lifang He. 2021. [Hierarchical taxonomy-aware and attentional graph capsule rnnns for large-scale multi-label text classification](#). *IEEE Transactions on Knowledge and Data Engineering*, 33(6):2505–2519.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Julian Risch, Samuele Garda, and Ralf Krestel. 2020. Hierarchical document classification as a sequence generation task. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries in 2020*, pages 147–155.
- Evan Sandhaus. 2008. The new york times annotated corpus. *Linguistic Data Consortium, Philadelphia*, 6(12):e26752.
- Kazuya Shimura, Jiyi Li, and Fumiyo Fukumoto. 2018. [HFT-CNN: Learning hierarchical category structure](#)

- for multi-label short text categorization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 811–816, Brussels, Belgium. Association for Computational Linguistics.
- Boyan Wang, Xuegang Hu, Peipei Li, and S Yu Philip. 2021. Cognitive structure learning model for hierarchical multi-label text classification. *Knowledge-Based Systems*, 218:106876.
- 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, Binghui 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.
- Jonatas Wehrmann, Ricardo Cerri, and Rodrigo Barros. 2018. Hierarchical multi-label classification networks. In *International conference on machine learning*, pages 5075–5084. PMLR.
- 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.
- Ran Xu, Yue Yu, Joyce Ho, and Carl Yang. 2023. Weakly-supervised scientific document classification via retrieval-augmented multi-stage training. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2501–2505.
- Pengcheng Yang, Xu Sun, Wei Li, Shuming Ma, Wei Wu, and Houfeng Wang. 2018. Sgm: sequence generation model for multi-label classification. *arXiv preprint arXiv:1806.04822*.
- Chao Yu, Yi Shen, and Yue Mao. 2022. Constrained sequence-to-tree generation for hierarchical text classification. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1865–1869.
- Dong Zhang, Shu Zhao, Zhen Duan, Jie Chen, Yanping Zhang, and Jie Tang. 2020. A multi-label classification method using a hierarchical and transparent representation for paper-reviewer recommendation. *ACM Transactions on Information Systems (TOIS)*, 38(1):1–20.
- Xinyi Zhang, Jiahao Xu, Charlie Soh, and Lihui Chen. 2022. La-hcn: label-based attention for hierarchical multi-label text classification neural network. *Expert Systems with Applications*, 187:115922.
- Yu Zhang, Zhihong Shen, Yuxiao Dong, Kuansan Wang, and Jiawei Han. 2021. Match: Metadata-aware text classification in a large hierarchy. In *Proceedings of the Web Conference 2021*, pages 3246–3257.
- Jie Zhou, Chungping 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.

A ENZYME Dataset




Figure 5: ENZYME dataset statistics

A.1 Enzyme Classification (EC) System

The Enzyme Commission (EC) number is a system for classifying enzymes, based on the chemical reactions they catalyze. The number is made up of four digits separated by periods and shows the enzyme's class, subclass, sub-subclass, and serial number respectively. This number is used to identify enzymes uniquely in the Enzyme Nomenclature database managed by the International Union of Biochemistry and Molecular Biology (IUBMB). The first digit (the "class" level) defines one of seven major enzymatic reaction types, e.g., hydrolases that use water to break a chemical bond. The second (the "subclass" level) and third digits (the "sub-subclass" level) represent more specific reaction types within their parent (sub-)class. The fourth digit (the "serial number" level) is a unique identifier assigned to a specific enzyme, usually determined by a specific substrate. For example, the enzyme number EC 1.3.8.2 represents an enzyme in class 1 (Oxidoreductases), subclass 3 (Acting

on CH-CH group of donors), sub-subclass 8 (With a flavin as acceptor) with a unique number 2 (to identify a specific substrate 4,4'-diapophytoene).

Level	# classes
1	7
2	24
3	31
4	462
# classes in the dataset:	4566

Table 4: Number of classes per level in the ENZYME dataset

A.2 Construction

The BRENDA database (Chang et al., 2021) was parsed to collect enzymes with both EC number definitions and PubMed references. Following an inspection of retrieved entries, we downloaded the freely available PDFs for each of the articles from PubMed³. To extract well-formatted text from these files, we used an open-source tool called SciPDFParser⁴. For this work, we use the title, abstract and introduction sections from each paper.

B Data Efficiency

In Section 5.4, we initially examined the performance of a restricted training set on the ENZYME dataset. In this subsequent section, we broaden our scope to encompass two additional datasets, namely WOS and NYT. Table 5 shows the comparison between HBGL and HiGen for WOS and NYT for different proportions of training data. Within this experimental framework, our primary objective is to assess the performance of HiGen in comparison to the best-performing baseline model, HBGL. Our results show a noteworthy trend as the proportion of training data increases: there is a consistent and incremental improvement in performance for both HBGL and HiGen. However, what distinguishes our findings is the remarkable consistency with which HiGen outperforms HBGL across all fractions of training data, as reflected by superior Micro and Macro F1 scores.

C Performance Analysis

C.1 Ablation Study Results

The performance figures for the ENZYME and NYT datasets shown in Tables 6 & 7 respectively,

³These journal articles and preprints are accessible via Open Access that allows reuse.

⁴https://github.com/titipata/scipdf_parser

confirm the analysis presented in the previous section for the WOS dataset. The pretraining step is absolutely crucial as it gives the largest performance boost. Both the semantic (\mathcal{L}_S) and token (\mathcal{L}_T) losses significantly contribute to improving the model’s Macro-F1 score while the contribution of the output space loss (\mathcal{L}_O) is significant but to a lesser extent. Across all scenarios, the removal of pretraining has the most profound impact, while omitting the proposed loss functions significantly influences performance, although with a less pronounced effect as compared to pretraining (Tables 3, 6 & 7).

Considering the improvement in HiGen over the Vanilla BART model, for the WOS dataset, the Micro-F1 score improves by 1.13 points and the Macro-F1 score goes up by 2.11 points. For datasets with progressively larger hierarchies, NYT and ENZYME, we observe more pronounced improvements. In the case of NYT, which features a larger hierarchy than WOS with 166 unique labels and 8 levels, the Micro-F1 score improves by 0.81 points, and the Macro-F1 score shows a 3.11 point enhancement. The ENZYME dataset features a four-level deep hierarchy but with a notably higher count of unique labels (4566), surpassing both WOS and NYT. Here, the Micro-F1 score sees a substantial 4.5 points increase, while the Macro-F1 score remarkably jumps by 16.39 points.

Across all three datasets, a notable disparity is evident wherein Vanilla-BART exhibits subpar performance in comparison to the proposed HiGen model. This highlights the pivotal contribution of the proposed loss functions and the pretraining strategy in boosting the generative model’s performance. Specifically, in the Data Efficiency and Ablation Study analyses, we showcase the poor performance of Vanilla BART in contrast to HiGen.

C.2 Comparison with Baselines

For all cases, the addition of the proposed framework over Vanilla BART yields substantial performance improvements. While the generative backbone contributes to performance enhancement to a certain degree when compared to baselines, a comprehensive evaluation against more recent baselines (HBGL, HGCLR, and T5-based baselines; refer to Table 2) reveals that Vanilla BART alone falls short of surpassing them. Hence, the additional modifications proposed under HiGen are absolutely essential.

Data Proportion	WOS				NYT			
	HBGL		HiGen		HBGL		HiGen	
	Micro F1	Macro F1	Micro F1	Macro F1	Micro F1	Macro F1	Micro F1	Macro F1
0.1	82.35	75.01	82.92	73.83	74.90	57.96	74.95	54.96
0.3	85.35	79.06	85.65	79.19	77.84	64.88	77.67	65.24
0.5	86.11	80.30	86.63	79.96	78.75	66.75	79.57	69.36
0.7	86.46	80.85	86.48	80.86	79.72	68.36	79.89	69.34

Table 5: Performance on WOS and NYT datasets on varying the training data proportion for HBGL (left) and HiGen (right). We used HBGL because it is the best-performing model to compare against.

Model	Micro-F1	Macro-F1
Vanilla BART	88.11	67.76
HiGen	92.61	84.15
w/o pretraining	88.16	69.05
w/o \mathcal{L}_O	92.33	83.00
w/o \mathcal{L}_T	92.44	82.38
w/o \mathcal{L}_S	92.36	82.43

Table 6: Ablation study on HiGen for the ENZYME dataset.

Model	Micro-F1	Macro-F1
Vanilla BART	80.08	69.30
HiGen	80.89	72.41
w/o pretraining	77.89	69.65
w/o \mathcal{L}_O	79.10	71.95
w/o \mathcal{L}_T	79.31	71.85
w/o \mathcal{L}_S	-	-

Table 7: Ablation study on HiGen for the NYT dataset.

HiGen outshines baselines by large margins when the hierarchy size and data imbalance increases. This is evident for the NYT and ENZYME datasets. HiMatch (Chen et al., 2021) suffers a notable performance degradation, particularly in the Macro-F1 scores, where HiGen w/o pretraining outperforms it (Table 2 and Tables 6 & 7 for HiMatch and HiGen respectively). Note that these models employ complex architectures to encode the hierarchical structure whereas HiGen w/o pretraining uses an encoder-decoder framework without any context of the hierarchical structure. To compete with the more recent baselines, HGCLR (Wang et al., 2022a) and HBGL (Jiang et al., 2022), we utilize pretraining to allow our model to grasp the hierarchical structure. Importantly, the pretraining approach we introduce is straightforward to implement and well-established. Simultaneously, it allows us to maintain a streamlined model architecture while achieving substantial performance gains.

C.3 Computational Cost

To emphasize the simplicity and efficacy of our proposed architecture, we conduct a comparison between HiGen and the best-performing baseline, HBGL, in terms of the training times and computational resource utilization. Table 8 provides an overview of the training times for both models across all datasets. Both models were trained on a single NVIDIA RTX A5000 (24G) GPU.

The reported times reveal a stark efficiency in our proposed approach, where HiGen is approximately 10 times faster than HBGL, notwithstanding the notably simpler architecture employed. Most importantly, despite the simpler design and reduced training times, HiGen achieves superior performance as compared to the aforementioned baseline.

Dataset	HiGen (hrs)	HBGL (hrs)
ENZYME	1.5	11.6
WOS	1.6	15.2
NYT	1.2	13.2

Table 8: Comparison of training times for HiGen and HBGL on all datasets.

D Hyperparameter Study

To study the influence of the loss-balancing factors λ_1 , λ_2 & λ_3 , we conduct a hyperparameter study for the WOS and ENZYME datasets. The results are reported in Tables 9 & 10. Our training and hyperparameter strategy was two-phased. Initially, we set the loss balancing factor for L_S to 1 based on preliminary experiments. Subsequently, an extensive search was conducted for λ_1 and λ_2 , with values ranging from 1 to $1e-8$ with a step factor of 0.1. From the results (Table 9), we observed that our model was quite robust to changes in middle ranges. Notably, a distinct performance peak emerged at $\lambda_1 = 1e-3$ and $\lambda_2 = 1e-5$.

To verify the generalizability of these findings, we performed a similar analysis for the ENZYME

λ_1	λ_2	λ_3	Micro-F1	Macro-F1
1e - 2	1e - 2	1	86.63	79.18
1e - 3	1e - 3	1	87.13	81.03
1e - 3	1e - 5	1	87.39	81.45
1e - 4	1e - 4	1	86.91	80.67
1e - 5	1e - 5	1	86.71	80.27

Table 9: Hyperparameter study on the WOS dataset.

λ_1	λ_2	λ_3	Micro-F1	Macro-F1
1e - 2	1e - 2	1	92.31	80.06
1e - 3	1e - 3	1	92.39	82.81
1e - 3	1e - 5	1	92.38	83.81
1e - 3	1e - 6	1	92.61	84.15
1e - 4	1e - 6	1	92.58	83.87
1e - 6	1e - 6	1	92.34	82.58

Table 10: Hyperparameter study on the ENZYME dataset.

dataset. The performance figures have been reported in Table 10. The performance peaks around the same hyperparameter settings as observed for the WOS dataset. This pattern similarly holds true for the NYT dataset, where the optimal hyperparameter settings align with those of WOS. This coherence in hyperparameter settings greatly simplified the process of tuning multiple parameters. Since almost the same settings work across all datasets, we argue that our model is quite robust to these values and future works could use this setting as a good starting point on other datasets as well.

E Generating Pretraining Data

Given the unavailability of similar datasets following the same label hierarchy for the WOS dataset, we employ the powerful ChatGPT API to generate the pretraining data. To capture the domain of science and the corresponding sub-category, we carefully designed the prompts. For instance, when focusing on the domain of "*Mechanical engineering*" and sub-category "*computer-aided design*", our curated prompt is as follows:

Write 20 different abstracts for scientific articles in the *Mechanical Engineering* domain and *computer-aided design*.

We modify the prompt accordingly for different domains and sub-categories. To ensure diversity and non-redundancy, we execute the API multiple times for each domain and sub-category combination, resulting in a well-balanced corpus of nearly 3000 abstracts. Refer to Table 11 for some examples of abstracts generated by ChatGPT.

To give an idea about the dissimilarity between the original test set in the WOS dataset and the generated pretraining data, we find the overlap between them. In order to do so, we preprocess the abstracts from both data sources by removing the stop words. Then, we compare the Jaccard Similarity between the two sets of data. Averaging the Jaccard Similarity across all categories we get the value to be 4.12%. This indicates that there is a minuscule overlap between the words (none of the overlaps exceeds 6%) from the test set of WOS and the pretraining data obtained from ChatGPT. This indicates that there is no data leakage during the pretraining step.

Category	Sub-domain	Text
Computer Science	Data Structures	In this study, we propose a new data structure called the Bloom filter tree. The Bloom filter tree combines the space efficiency of Bloom filters with the fast search capabilities of trees. We show that the Bloom filter tree outperforms other data structures in terms of space usage and query time.
Civil Engineering	Stealth Technology	This article presents an experimental study on the effect of moisture on the mechanical properties of asphalt concrete. The study was conducted through laboratory tests on asphalt concrete samples with varying moisture content. The results show that moisture has a significant impact on the stiffness and strength of the asphalt concrete.
Biochemistry	Genetics	This article discusses the emerging field of microRNA research and its implications for gene regulation and disease pathogenesis. The research team investigates the role of microRNAs in various biological processes, including development, cell differentiation, and immune response. The findings highlight the potential of microRNAs as therapeutic targets for various diseases.

Table 11: Examples of abstracts generated by ChatGPT.