

CoPHE: A Count-Preserving Hierarchical Evaluation Metric in Large-Scale Multi-Label Text Classification

Matúš Falis¹, Hang Dong², Alexandra Birch¹, and Beatrice Alex^{1,3,4}

¹Institute for Language, Cognition and Computation, University of Edinburgh

²Centre for Medical Informatics, Usher Institute, University of Edinburgh

³School of Literatures, Languages and Cultures, University of Edinburgh

⁴Edinburgh Futures Institute, University of Edinburgh

{s1206296, Hang.Dong, a.birch, balex}@ed.ac.uk

Abstract

Large-Scale Multi-Label Text Classification (LMTC) includes tasks with hierarchical label spaces, such as automatic assignment of ICD-9 codes to discharge summaries. Performance of models in prior art is evaluated with standard precision, recall, and F_1 measures without regard for the rich hierarchical structure. In this work we argue for hierarchical evaluation of the predictions of neural LMTC models. With the example of the ICD-9 ontology we describe a structural issue in the representation of the structured label space in prior art, and propose an alternative representation based on the depth of the ontology. We propose a set of metrics for hierarchical evaluation using the depth-based representation. We compare the evaluation scores from the proposed metrics with previously used metrics on prior art LMTC models for ICD-9 coding in MIMIC-III. We also propose further avenues of research involving the proposed ontological representation.

1 Introduction

The Large-Scale Multi-Label Text Classification (LMTC) is a multi-label document classification task with a large (in the order of thousands) and potentially structured label space. LMTC tasks with structured label spaces can be found in the domain of medicine, legislation, product categorisation, *etc.* Individual label spaces may follow different labelling rules, *e.g.*, the *EURLEX* (Chalkidis et al., 2019) dataset’s labels can be any node from the EUROVOC tree of concepts, while *AMAZON13K* (McAuley and Leskovec, 2013) assumes a dense assignment, where if a predicted (leaf) node’s truth value is positive, so is the truth value of all its ancestors – also known as the *true path rule*. A representative dataset for medical LMTC are *MIMIC-III* (Johnson et al., 2016) discharge summaries. These are weakly labelled with leaf nodes of the ICD-9

ontology¹. ICD-9 is a tree-structured ontology of medical conditions and procedures.

Since the release of the MIMIC-III dataset there have been several attempts at training neural models for automated coding of medical documents (Mullenbach et al., 2018; Rios and Kavuluru, 2018; Falis et al., 2019; Chalkidis et al., 2020; Dong et al., 2021). While some prior art has made use of ontological structure (Rios and Kavuluru, 2018; Falis et al., 2019; Manginas et al., 2020), the task has mostly been treated as a flat prediction of the assigned leaves. This is reflected in the evaluation metrics that are used across previous work – precision, recall and F_1 score on flat predictions. These metrics applied to flat vectors disregard the rich ontological structure. A notable exception is Manginas et al. (2020) – they fine-tune a BERT (Devlin et al., 2018) model, such that different layers within the model learn to represent the different depths of the target label space. Each layer undergoes flat evaluation with respect to labels on its depth.

The aim of this study is to review the use of ontological structure within the prior LMTC art with a focus on automated ICD-9 coding, point out issues in hierarchy representation in prior approaches, describe methods more suited for addressing the structured label space including a hierarchical evaluation metric, and propose further avenues of research utilising structured label spaces. Our implementation of evaluation and the representation of ICD-9’s graph are available to the community².

2 Background

Within a structured label space some labels are inherently closer to one another than to others - *e.g.*, in ICD-9 425.0 *Endomyocardial fibrosis* is closer to 425.3 *Endocardial fibroelastosis* than to 305.1 *Tobacco use disorder*. Flat prediction and standard

¹<https://www.cdc.gov/nchs/icd/icd9cm.htm>

²<https://github.com/modr00cka/CoPHE>

precision/recall/ F_1 score (from hereon referred to as *standard metrics*) of individual prediction level (leaf) codes treat all mispredictions equally – *e.g.*, having 425.0 be mispredicted as 425.3 is penalised the same way as mispredicting it as 305.1. This phenomenon has been addressed in information extraction (IE) by Maynard et al. (2008) through the use of distance metrics. The IE setting assumes both the gold standard and predictions to be associated with specified spans within the input text. This means an individual prediction can be associated with a true label, allowing direct comparison between them.

The LMTC setting uses weak labels – predictions and true labels appear on the document level, without exact association to spans within the text. Due to the absence of information regarding associated spans, direct links between individual predictions and true labels do not exist. Label comparison is performed on full vectors representing multiple labels, hence the IE approach is not directly usable. Kosmopoulos et al. (2015) address hierarchical label spaces in document classification with *set-based* measures. The gold standard and prediction vectors are extended to include ancestor nodes within the hierarchical label-space and augmented according to its structure, and the true-path rule.

Let $X = \{x_i | 1, \dots, M\}$ and let $Y = \{y_i | 1, \dots, N\}$ represent the set of predicted and true codes for a certain document respectively. Assume we have access to an augmentation function $An_j(x)$ which returns the ancestors of x up to the j^{th} order.

Let $X_{aug} = X \cup \{An_j(x_i) | 1, \dots, M\}$ and let $Y_{aug} = Y \cup \{An_j(y_i) | 1, \dots, N\}$ represent the set of predicted and gold ancestor codes for X and Y respectively. Standard metrics can then be applied to the X_{aug} and Y_{aug} sets. A correct assignment of predicted lower-level (leaf) codes results in correct assignment of their ancestors. In the case of incorrect prediction-level assignments, the closer a mismatched predicted code is to the gold standard leaf within the ontology, the more matches will occur across the levels of the hierarchy.

On the leaf level each code appears at most once per document. Duplicates can occur when $An_j(x)$ produces the same ancestor for multiple codes. As X_{aug} is a set, duplicates are removed. Hence, the set-based approach captures whether an ancestor is present, but not how many of its descendants

were predicted. This results in loss of information regarding over-/under- predictions of classes on the ancestral level. Over- and under-prediction is a valuable phenomenon to track, particularly if the label-space includes inexplicit rules – for instance, for some nodes only a single descendant can be predicted at a time, as individual siblings are mutually exclusive (*e.g.*, a patient can be assigned at most one of codes 401.0, 401.1, and 401.9, which represent malignant, benign, and unspecified hypertension respectively – concepts that are mutually exclusive). Furthermore, retaining this numeric data on ancestral levels enables analyses on higher levels, *e.g.*, performance of a family of codes in the case of a semi-automated code-assignment application. For this reason we propose using a metric that retains the descendant counts for these ancestor codes.

To correctly define the augmentation function we need to ensure our representation of the hierarchy fits the setting. Previous work involving the ICD-9 hierarchy (Rios and Kavuluru, 2018; Falis et al., 2019; Chalkidis et al., 2020) represents it through the relation of direct ancestry considering parents and grandparents of the leaf nodes. As the ICD-9 has leaves at different depths, this representation results in structural issues, such as one code being both in the position of a parent and a grandparent for different leaves. For instance, code 364 has a parent relation to the leaf 364.3 and grandparent to leaf 364.11 (Figure 1). This poses an issue to aggregation and evaluation. We aim to address this issue by producing a representation of the hierarchy through the levels of the label space with each level representing all the nodes at a certain depth in the tree structure.

3 Method

3.1 Augmentation

To implement level-based augmentation, we first define the first three layers of the ICD-9 hierarchy on which leaves appear. An ICD-9 code (*e.g.*, 364.11) consists of a “category” (part of the code appearing prior to the decimal point, *e.g.*, 364) and “etiology” (appearing after the decimal point, *e.g.*, 11). The etiology can be represented with up to two digits. We define the basic levels of the hierarchy (encapsulating all the labels in MIMIC-III) as follows: codes with double digit etiology (e_2); codes of single digit etiology (e_1); codes described only with “category” (no etiology, e_0). Augmenta-

tion can be performed up to a higher user-defined level within the hierarchy by adding further layers representing chapters within the ontology.

The originally flat predicted and true labels are divided into their respective layers in the hierarchy. If a code appears in a level lower than the maximum level set by the user, the truth value of the code is propagated to its direct ancestor through augmentation. The propagation can be interpreted either as a truth value (*binary*) or the number of descendant leaves present (*count-preserving*). The binary interpretation of the propagation results in the ancestor holding the truth value of the logical OR operation on its children. This mimics the set-based approaches described by Kosmopoulos et al. (2015). The count-preserving interpretation sets the value of the ancestor to be the sum of the values of its descendants. Through retaining the numeric information, the count-preserving interpretation allows us to track over- or under-prediction within a family of codes.

3.2 Hierarchical Evaluation

To produce hierarchical precision, recall, and F_1 up to a certain level within the hierarchy, the vectors from each layer are combined for the predictions and true labels respectively. In the case of set-based interpretation (Figure 2 B), standard metrics can be directly applied to the augmented vectors.

Since the notions of true positive (TP), false negative (FN), and false positive (FP) needed for calculation of precision, recall, and F_1 are defined for binary input, to produce the count-preserving interpretation of the hierarchical measures, we need to use a count-preserving version of TP , FP , and FN . When comparing the number of predictions and true labels mapped to an ancestor:

$$TP_{c,d} = \min(x_{c,d}, y_{c,d}) \quad (1)$$

$$FP_{c,d} = \max(x_{c,d} - y_{c,d}, 0) \quad (2)$$

$$FN_{c,d} = \max(y_{c,d} - x_{c,d}, 0) \quad (3)$$

Where c represents a particular (ancestor) label, d a specific document, $x_{c,d}$ and $y_{c,d}$ are the numbers of predicted and true prediction-level descendant labels of c in d . The functions \min and \max return the minimum and maximum of two real numbers respectively. TP (Equation 1) represents the overlap between the expected and predicted count of descendant codes of the ancestor c . FP (Equation 2) and FN (Equation 3) represent the over-prediction

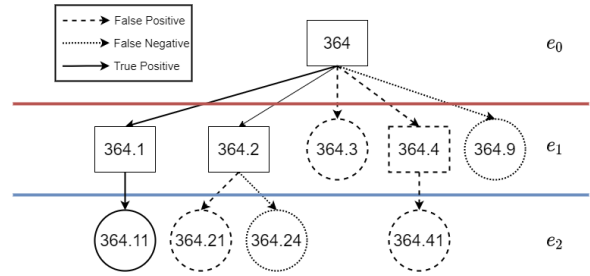


Figure 1: An example of hierarchical evaluation. Circular nodes represent leaf nodes (for non-hierarchical evaluation), borders of nodes represent set-based hierarchical evaluation, edges represent count-preserving hierarchical evaluation. Solid lines represent TP , dashed-lines represent FP , dotted lines represent FN . Levels of depth in the ontology (e_0 , e_1 , and e_2) are indicated with horizontal lines.

and under-prediction of these descendants, respectively.

Remark: Note that if $x_{c,d}$ and $y_{c,d}$ are binary, the results of these calculations are equivalent to those of standard TP , FP , FN .

	TP	FP	FN	P	R	F_1
Leaf-Only	1	3	2	25.0	33.3	28.4
Set-Based						
e_2	1	2	1	33.3	50	40
e_1	2	2	1	50	66.7	57.1
e_0	1	0	0	100	100	100
Overall	4	4	2	50	66.7	57.1
CoPHE						
e_2	1	2	1	33.3	50	40
e_1	2	2	1	50	66.7	57.1
e_0	3	1	0	75	100	86
Overall	6	5	2	54.5	75	63.1

Table 1: Evaluation of the three representations presented in Figure 2. In the case of Set-Based evaluation and CoPHE we present both the evaluation at each level of the ontology, and overall evaluation across levels.

Suppose 4 descendants of code 364 – 364.11 , 364.21 , 364.3 , and 364.41 – are predicted and 3 descendants of 364 are true for document δ – 364.11 , 364.24 , and 364.9 (as seen in Figure 1). In the count-preserving approach, according to our definitions, this results in $TP = 3$, $FP = 1$, $FN = 0$ for the highest level (364). The set-based version of hierarchical evaluation would only track if a prediction or gold standard path passes through a node, rather than their count (losing data on over- and under-prediction of each node’s descendants). Here in the highest level (code 364), however, the four prediction paths and three gold standard paths are reduced to boolean *True* for both prediction and gold standard. Hence the values for 364 become




Figure 2: A comparison between three styles of evaluation: (A) Evaluation performed only on leaf nodes (no use of hierarchical relations); (B) Set-Based Hierarchical Evaluation (Kosmopoulos et al., 2015). Four descendants of the node 364 are predicted, and three appear in the gold standard. This is reflected in the middle-level vector, but the information is lost in the top level; (C) Count-Preserving Hierarchical Evaluation. The numeric information of predictions and gold labels are preserved on higher levels.

$TP = 1$, $FP = 0$, $FN = 0$. Note that due to this representation the over-prediction of descendants of 364 is not captured (Figure 2 B). The binary and count-preserving evaluation can be produced for each individual code, aggregated over individual levels of depth of the label space or multiple consecutive layers of depth starting from the lowest. We refer to the latter aggregation in conjunction with count-preserving approach as Count-Preserving Hierarchical Evaluation (CoPHE; Figure 2 C). The micro-averaged precision, recall and F_1 scores (both overall and per-level) of the example in Figure 2 are presented in Table 1.

4 Metric Analysis

We produced evaluation results with CoPHE on three prior art models on MIMIC-III.³ CAML (Mullenbach et al., 2018) is the first notable LMTC model developed on MIMIC-III whose contribution to the task was the introduction of label-wise attention. Chalkidis et al. (2019) proposed an alteration on CAML called BIGRU-LWAN, (here re-

³Since BERT based models are yet to be well adapted to MIMIC-III ICD coding (Gao et al., 2021; Ji et al., 2021), we leave BERT models for a future study.

ferred to as BGLWAN) swapping CAML’s CNN encoder with a bidirectional GRU, and making use of zero-shot-motivated label description embeddings proposed by Rios and Kavuluru (2018). Dong et al. (2021) further added label-wise word- and sentence-level attention and incorporated label embeddings pre-trained from the training labelsets (instead of the label descriptions) into the prediction and attention layers producing the HLAN model.

In the context of MIMIC-III, results tend to be presented on two label sets – the 50 most frequent codes (*top50*) and the full codeset (*full*) – this stems from the big-head long-tail distribution of labels in MIMIC-III with many labels being infrequent, or not appearing within the training set at all. We have used the pre-processing and *top50* codeset dataset split of Mullenbach et al. (2018). Our results are averaged across 10 runs and serve as a comparative example for the evaluation metrics.

We have applied CoPHE up to the lowest chapter levels (one above e_0), which we simply refer here to as the chapter level (c). This choice was made due to the level above c already including the root of the ontology for some codes, leading to structural inconsistency.

We compared the flat metric results against the results of the set-based and CoPHE hierarchical measures on the three models (Table 2). The scores of CoPHE are higher than that of flat evaluation, showing that the measure is more lenient in its interpretation of errors.

The scores of CoPHE are consistently lower than that of set-based evaluation, showing that CoPHE can capture more over- and/or under-prediction of the labels, as the True Positives of the ancestor codes are better manifested in CoPHE. If all (binary) True Positives within higher levels of the hierarchy reported by the set-based measure were True Positives in CoPHE (no over- or under-predictions), this would result in a greater proportion of True Positives for CoPHE than in the set-based measure, and hence higher scores (see Table 1 as an example). In contrast, the consistent lower CoPHE scores across the three methods imply the presence of over- and/or under-prediction within families of codes that is not captured by the set-based measure.

The change of evaluation measure has not affected the ranking of the models on F_1 . This is likely due to a lack of representation of the label space structure within the explored models.

While hierarchical measures enrich the evalua-

Model	Standard Flat			Set-Based			CoPHE		
	P	R	F_1	P	R	F_1	P	R	F_1
CAML	59.3	61.4	60.1	62.6	65.5	64.2	61.1	65.4	63.1
BGLWAN	68.4	57.6	62.5	71.8	61.2	66.0	70.5	61.4	65.1
HLAN	73.9	57.4	64.2	77.0	60.3	68.5	76.1	59.1	66.5

Table 2: A comparison between flat evaluation and hierarchical evaluation – Set-Based and CoPHE – on three models from prior art using the *top50* codeset. The results are micro-averaged across labels (and ancestors for hierarchical measures). Levels of hierarchy up to and including the lowest chapter level (e_2 , e_1 , e_0 , c) are considered. Hierarchical measures report higher F_1 scores than the original flat measures. Furthermore, Set-Based evaluation F_1 scores are higher than those of CoPHE.

tion with structural information, they do not serve as a replacement of the flat measures, but rather should be viewed side-by-side. An approach may not surpass the state-of-the-art on the flat measures, whilst being better at modelling the structure of the label space based on the hierarchical score. Indication of better modelling of the structure should prompt further analysis of the model.

5 Conclusion

We have proposed the use of hierarchical evaluation measures in the LMTC task involving hierarchical label spaces and provided an example in the task of automated ICD coding. Unlike the approaches in prior art of ICD coding, which penalise all mispredictions equally, the proposed hierarchical evaluation measures adjust the penalty based on the performance on the ancestral levels. We have described a means to represent the hierarchy according to depth within the ontology. Finally we have proposed the use of count-preserving evaluation which captures data on both over and under-prediction in ancestral levels, as opposed to an existing set-based (binary) hierarchical evaluation approach.

6 Future Work

We intend to use the proposed hierarchical evaluation metrics alongside the flat metrics from prior art in our future experiments, particularly in the ones incorporating ontological structure within the model (similar to Rios and Kavuluru (2018)). The code for the proposed metric along with the level-based representation of the ontology have been made public in order to aid future work. The hierarchical evaluation can be applied to other LMTC tasks using different hierarchical label spaces. Finally, not all structured label spaces follow a tree structure – it is important to explore the possibil-

ity of similar measures for ontologies with more generic graphs, *e.g.*, SNOMED CT.⁴

Acknowledgements

This work was supported by the United Kingdom Research and Innovation (grant EP/S02431X/1), UKRI Centre for Doctoral Training in Biomedical AI at the University of Edinburgh, School of Informatics, Health Data Research UK (HDR UK) National Text Analytics and Phenomics Projects, and the Wellcome Institutional Translation Partnership Award (P111009). The authors would like to thank Francisco Vargas and Andreas Grivas for their feedback on the visualisation of CoPHE.

References

- Ilias Chalkidis, Manos Fergadiotis, Sotiris Kotitsas, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. An empirical study on large-scale multi-label text classification including few and zero-shot labels. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7503–7515.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, and Ion Androutsopoulos. 2019. Large-scale multi-label text classification on eu legislation. *arXiv preprint arXiv:1906.02192*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Hang Dong, Víctor Suárez-Paniagua, William Whiteley, and Honghan Wu. 2021. Explainable automated coding of clinical notes using hierarchical label-wise attention networks and label embedding initialisation. *Journal of Biomedical Informatics*, 116:103728.

⁴<https://www.nlm.nih.gov/healthit/snomedct/index.html>

- Matúš Falis, Maciej Pajak, Aneta Lisowska, Patrick Schrenpf, Lucas Deckers, Shadia Mikhael, Sotirios A Tsaftaris, and Alison Q O’Neil. 2019. Ontological attention ensembles for capturing semantic concepts in icd code prediction from clinical text. volume 47, page 168.
- Shang Gao, Mohammed Alawad, Michael Todd Young, John Gounley, Noah Schaefferkoetter, Hong-Jun Yoon, Xiao-Cheng Wu, Eric B Durbin, Jennifer Doherty, Antoinette Stroup, et al. 2021. Limitations of transformers on clinical text classification. *IEEE journal of biomedical and health informatics*.
- Shaoxiong Ji, Matti Hölttä, and Pekka Marttinen. 2021. Does the magic of bert apply to medical code assignment? a quantitative study. *arXiv preprint arXiv:2103.06511*.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, H Lehman Li-Wei, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data*, 3(1):1–9.
- Aris Kosmopoulos, Ioannis Partalas, Eric Gaussier, Georgios Paliouras, and Ion Androutsopoulos. 2015. Evaluation measures for hierarchical classification: a unified view and novel approaches. *Data Mining and Knowledge Discovery*, 29(3):820–865.
- Nikolaos Manginas, Ilias Chalkidis, and Prodrimos Malakasiotis. 2020. Layer-wise guided training for bert: Learning incrementally refined document representations. *arXiv:2010.05763 [cs]*.
- Diana Maynard, Wim Peters, and Yaoyong Li. 2008. Evaluating evaluation metrics for ontology-based applications: Infinite reflection. In *LREC*.
- Julian McAuley and Jure Leskovec. 2013. Hidden factors and hidden topics: understanding rating dimensions with review text. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 165–172.
- James Mullenbach, Sarah Wiegrefe, Jon Duke, Jimeng Sun, and Jacob Eisenstein. 2018. Explainable prediction of medical codes from clinical text. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1101–1111.
- Anthony Rios and Ramakanth Kavuluru. 2018. Few-shot and zero-shot multi-label learning for structured label spaces. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3132–3142.