EHRKit: A Python Natural Language Processing Toolkit for Electronic Health Record Texts

Anonymous ACL submission

Abstract

The Electronic Health Record (EHR) is an essential part of the modern medical system and impacts healthcare delivery, operations, and research. Unstructured text is attracting much attention despite structured information in the EHRs and has become an exciting research field. The success of the recent neural Natural Language Processing (NLP) method has led to a new direction for processing unstructured clinical notes. In this work, we create a python library for clinical texts, EHRKit. This library contains two main parts: MIMIC-III-specific functions and task-specific functions. The first part introduces a list of interfaces for accessing MIMIC-III NOTEEVENTS data, including basic search, information retrieval, and information extraction. The second part integrates many third-party libraries for up to 12 off-shelf NLP tasks such as named entity recognition, summarization, machine translation, etc.

1 Introduction

With the rising trend of Electronic Health Records (EHRs), massive unstructured texts (i.e., clinical and admission notes) are being created in the healthcare system. It is very important to process such data for secondary usage (Xiao et al., 2018). The main obstacle is the processing and understanding of the unstructured text. Natural Language Processing (NLP) techniques have been applied to deal with such texts (Li et al., 2021a; Shickel et al., 2018; Al-Aiad et al., 2018). Especially, deep learning-based methods achieved great success in some existing NLP tasks in the biomedical and clinical literature, such as text classification (Zhou et al., 2021; Li and Yu, 2020; Li et al., 2019; Hughes et al., 2017), named entity recognition (Song et al., 2021), text segmentation (Badjatiya et al., 2018), medical language translation and generation (Weng et al., 2019; Abacha and Demner-Fushman, 2019) and many others.

Following the success of BERT (Devlin et al., 2019), researchers developed BERT-based models trained on the clinical literature, such as BioBERT (Lee et al., 2020) and ClinicalBERT(Huang et al., 2019). With such advanced neural models, there is a need for a user-friendly programming interface that can support a variety of downstream tasks. Some existing libraries and toolkits are designed for bioinformatical and clinical needs, including the biomedical and clinical model packages of Stanza (Zhang et al., 2021), SciFive (Phan et al., 2021), UmlsBERT (Michalopoulos et al., 2021), MIMIC-Extract (Wang et al., 2020) and so on. However, we noticed a need to integrate more existing libraries with much broader coverage of clinical and biomedical NLP tasks in our toolkit. Besides, based on the analysis from Li et al. (2021a), there is limited research on generation tasks for EHR unstructured text, i.e., machine translation. Thus, we provide a pretrained model for clinical text machine translation in our toolkit, which supports three languages.

Our contributions are: 1) First, we propose EHRKit, a Python NLP toolkit for EHR unstructured texts. This toolkit contains two main components: general API functions and MIMIC-specific functions. It is user-friendly, with easy installation and quick start tutorials. 2) Second, to fill the gap in text generation for clinical texts, we release pretrained machine translation models in three languages. Besides, we evaluate existing pretrained models for summarization in a biomedical and clinical scenario. We release this toolkit and pretrained models publicly available at `placeholder.link`.

2 System Design and Architecture

We show our EHRKit architecture in Fig. 1. It consists of two modules, namely MIMIC-III Tasks and Wrapper Functions.

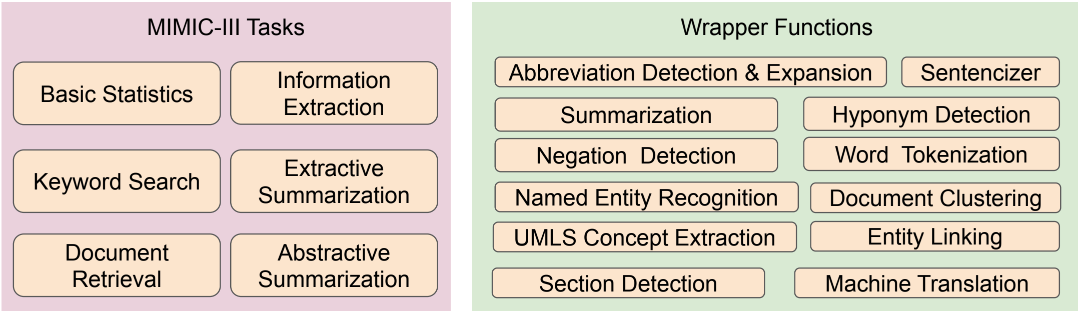

Figure 1: EHRKit Architecture.

2.1 MIMIC-III Tasks

We include some basic NLP functions for MIMIC-III NOTEEVENTS text data (Johnson et al., 2016).

- Basic statistical functions support counting for number of patients and number of documents, number of sentences and so on.
- Information Extraction provides helpful interfaces for investigating the data, i.e., phrases extraction, and abbreviation term extraction from a given record ID. We applied Phrase-At-Scale[1] for this function.
- Keyword Search allows users to search a record by a keyword.
- Document Retrieval allows users to search a record by ID.
- Extractive and Abstractive Summarization: Naïve Bayes (Ramanujam and Kaliappan, 2016) and DistilBART (Shleifer and Rush, 2020) for text summarization.

2.2 Wrapper Functions

This module integrates many third-party libraries and supports up to 12 functionalities for any free-text inputs.

- Abbreviation Detection and Expansion: finds abbreviation and its expansions. Function imported from ScispaCy[2].
- Sentencizer: detects sentence boundaries. We support four approaches: PyRuSH[3], Stanza, ScispaCy and Stanza Biomed.
- Hyponym Detection: finds the hyponyms of the recognized entities in the input text. Function imported from scispaCy.
- Negation Detection: detects negation in a sentence, imported from medspaCy (Eyre et al., 2021a)
- Word Tokenization: tokenizes a sentence into a list of words. Function imported from medspaCy.
- Named Entity Recognition: finds named entities, part-of-speech and universal morphological features, and dependencies of an input record. Function imported from Stanza (Zhang et al., 2021).
- Document Clustering: given the query record, selects k documents from supporting records that are most similar to the main record (K-Means clustering), measured by embedded document using pretrained BERT model (Devlin et al., 2019).
- UMLS Concept Extraction: matches the UMLS concept for the input text. Function imported from medspaCy (Eyre et al., 2021a).
- Entity Linking: finds named entities, negation entities, and linked entities in the input text. Function imported from scispaCy.
- Section Detection: rule-based method for detecting section (i.e., *allergies*, *history*). Function imported from medspaCy (Eyre et al., 2021a).
- Machine Translation: translates English texts into 17 target languages. We applied the existing MarianMT model[4], as well as our own fine-tuned models.
- Summarization: we support both extractive and abstractive summarization methods. We integrated TextRank (Mihalcea and Tarau, 2004), pretrained BART (Lewis et al., 2019),

[1] https://github.com/kavgan/phrase-at-scale
[2] https://allenai.github.io/scispacy/
[3] https://github.com/jianlins/PyRuSH.
[4] https://huggingface.co/docs/transformers/model_doc/marian

	MIMIC	Neu	MT	Summ
MIMIC-Extract	✓			
ScispaCy		✓		
medspaCy		✓		
Stanza Biomed		✓		
SciFive		✓		✓
EHRKit (ours)	✓	✓	✓	✓

Table 1: A comparison with other similar python toolkits. `MIMIC`: MIMIC Related. `Neu`: Neural Methods. `MT`: Machine Translation. `Summ`: Summarization.

T5 (Raffel et al., 2020) and SciFive summarization libraries. We also allow single and multiple documents as the input.

2.3 Other similar libraries

MIMIC-Extract[5] A pipeline for preprocessing and presenting data from MIMIC-III dataset. It provides features for data analysis, including extraction of clinical events like mortality from free text.

ScispaCy (Neumann et al., 2019) A tool that adapts SpaCy's models to process scientific, biomedical, and clinical text. It supports multiple methods for tokenization, part of speech tagging, dependency parsing, and named entity recognition.

medspaCy Based on the spaCy framework, medspaCy (Eyre et al., 2021b) is a clinical NLP python library that provides both rule-based and machine learning-based methods for processing clinical text. It supports methods for various clinical applications such as UMLS Mapping (rule-based), Section Detection, Sentence Detection, Contextual Analysis and Visualization on entities.

Stanza Biomed (Zhang et al., 2021) A set of tools for statistical, neural, and rule-based problems in computational linguistics. Its software provides a simple interface for NLP tasks. It is a widely used Python library for processing clinical texts. It provides nearly state-of-the-art performance using neural networks on tasks including tokenization, sentence segmentation, part of speech (POS) tagging, lemmatization, and dependency parsing.

SciFive (Phan et al., 2021) A pretrained neural language model for biomedical domain. Fine-tuned on PubMed Abstract [6] and PubMed Central (PMC) [7], it outperformed similar models including

[5] https://github.com/MLforHealth/MIMIC_Extract

[6] https://pubmed.ncbi.nlm.nih.gov/

[7] https://www.ncbi.nlm.nih.gov/pmc

Lang. Pair	Total	Train	Test
en → es	790,915	672,276	111,779
en→fr	2,812,305	2,390,458	407,388
en→ ro	1,165,092	990,327	161,936

Table 2: Data statistics for machine translation: we apply UFAL and select the overlapped target language pairs for our experiments.

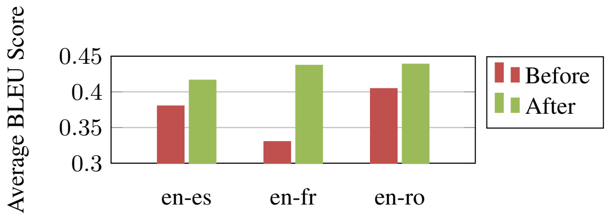

Figure 2: BLEU score: before and after pretraining.

BioBERT and T5 on multiple NLP tasks: named entity relation, relation extraction, natural language inference, and question answering.

Tab. 1 lists our EHRKit and other similar toolkits. We compare from different perspectives by focusing on the functionalities. **MIMIC Related**: if supports MIMIC-related functions. We consider MIMIC an essential data source that plays an important role in research. We can find that only MIMIC-Extract and EHRKit support these related functions, and users can apply them directly to the MIMIC data. **Neural Methods**: if this toolkit supports neural methods and embedding methods. As we can see, the majority contain such features. **Machine Translation** and **Summarization**: if this toolkit supports (neural) generation tasks like machine translation and summarization. In this case, only SciFive and EHRKit support such features. Our toolkit provides diverse functionalities and is easy to use based on these perspectives.

3 Performance Evaluation

3.1 Machine Translation

We report the performance of the Machine Translation function from our EHRKit and compare it with the baseline model, MarianMT. Our training sets and test sets are obtained from the UFAL Medical Corpus[8]. These data are from various medical text sources, such as titles of medical Wikipedia articles, medical term-pairs, patents, and documents from the European Medicines Agency (Braune et al.,

[8] https://ufal.mff.cuni.cz/ufal_medical_corpus

	PubMed			MIMIC-CXR		
	R-1	R-2	R-L	R-1	R-2	R-L
Pegasus (Zhang et al., 2020)	45.97	20.15	28.25	65.11	**52.90**	**61.88**
BigBird (Zaheer et al., 2020)	46.32	**20.65**	**42.33**	63.85	51.09	60.55
BART (Lewis et al., 2019)	44.16	20.28	36.80	62.09	49.02	58.65
SciFive (Phan et al., 2021)	**48.83**	15.81	37.06	**65.17**	52.45	61.80

Table 3: Summarization evaluation: we evaluate selected models and report ROUGE-1, ROUGE-2 and ROUGE-L.

Dataset	Train	Valid	Test
PubMed	112K	6.6K	6.7K
MIMIC-CXR	91,544	2000	600

Table 4: Data statistics for PubMed and MIMIC-CXR summarization datasets. Words are counted before tokenization.

2018). We evaluate language pairs, including English (en) to Spanish (es), English to French (fr), and English to Romanian (ro), as those are the three language pairs that EHRKit and UFAL mutually support.

For data pre-processing, we first exclude general domain data from UFAL, such as parliament proceedings. Next, we randomly shuffle the medical-domain corpora and split it into two parts by 85% and 15%, as our training set and test set, respectively. For each language pair, we use all of the available parallel data. Tab. 2 summarizes the number of sentences that we extract from UFAL.

Subsequently, we evaluate model performance using the BLEU score (Papineni et al., 2002). We finetune our model with the training sets (After) and compare it with the baseline model, MarianMT (Before). As shown in Fig. 2, our model improves significantly after finetuning, with an average 16.80% increase in BLEU score. Among the three selected language pairs, we can observe that English to French has the best improvement - it achieves 32.40% performance gain. We conjecture that this occurs because we have significantly more training data in English to French.

3.2 Summarization

Since there are many existing pretrained models for summarization in the general NLP field, we investigate how they perform in the biomedical literature and clinical data. Summarization corpora from the clinical scenario are very challenging to be obtained, so we chose the existing PubMed (Cohan et al., 2018) and MIMIC-CXR (Johnson et al., 2019) as our principal datasets.

PubMed dataset consists of 133k biomedical scientific publications from the PubMed database. Each input document is a scientific article, and the reference summary is the associated abstract. MIMIC-CXR is a de-identified, Protected Health Information removed dataset of chest radiographs, with a DICOM format and free-text radiology reports. We use a subset from the MIMIC-CXR for the MEDIQA 2021 Radiology report summarization shared task (Delbrouck et al., 2021). Each example contains three fields: (a) findings field is the original human-written radiology findings text, (b) impression field is the human-written radiology impression text, and (c) background field is the background information of the study in text format. We show the statistics in Tab. 4.

We evaluate selected pretrained abstractive methods using ROUGE (Lin, 2004) in Tab. 3. Among the four models, we can observe that SciFive has a high R-1 score, but BigBird (Zaheer et al., 2020) and Pegasus (Zhang et al., 2020) achieve a better score on R-2 and R-L, respectively, on the two datasets. This evaluation shows that it is challenging to determine which model is the best in our specific scenario, though SciFive was pretrained for this purpose. In the future, more work can be done to improve automatic summarization for biomedical and clinical texts.

4 Conclusion

In this work, we propose a python library for clinical texts, EHRKit. This toolkit contains two main components: general API functions and MIMIC-specific functions. In the future, we will investigate more EHR-NLP tasks including machine translation for more languages, multi-document summarization and question answering (Li et al., 2021a). Besides, we plan to investigate better-performed NLP models for these tasks, for example, BERT-based models (Lee et al., 2020; Li et al., 2021b).

References

Asma Ben Abacha and Dina Demner-Fushman. 2019. On the summarization of consumer health questions. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 2228–2234. Association for Computational Linguistics.

Ahmad Al-Aiad, Rehab Duwairi, and Manar Fraihat. 2018. Survey: Deep learning concepts and techniques for electronic health record. In *15th IEEE/ACS International Conference on Computer Systems and Applications, AICCSA 2018, Aqaba, Jordan, October 28 - Nov. 1, 2018*, pages 1–5. IEEE Computer Society.

Pinkesh Badjatiya, Litton J. Kurisinkel, Manish Gupta, and Vasudeva Varma. 2018. Attention-based neural text segmentation. In *Advances in Information Retrieval - 40th European Conference on IR Research, ECIR 2018, Grenoble, France, March 26-29, 2018, Proceedings*, volume 10772 of *Lecture Notes in Computer Science*, pages 180–193. Springer.

Fabienne Braune, Alex. Fraser, and Barry Haddow. 2018. D1.2: Report on improving translation with monolingual data.

Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. A discourse-aware attention model for abstractive summarization of long documents. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 615–621, New Orleans, Louisiana. Association for Computational Linguistics.

Jean-Benoit Delbrouck, Cassie Zhang, and Daniel Rubin. 2021. QIAI at MEDIQA 2021: Multimodal radiology report summarization. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 285–290, 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.

Hannah Eyre, Alec B Chapman, Kelly S Peterson, Jianlin Shi, Patrick R Alba, Makoto M Jones, Tamara L Box, Scott L DuVall, and Olga V Patterson. 2021a. Launching into clinical space with medspacy: a new clinical text processing toolkit in python. In *AMIA Annual Symposium Proceedings 2021*.

Hannah Eyre, Alec B Chapman, Kelly S Peterson, Jianlin Shi, Patrick R Alba, Makoto M Jones, Tamara L Box, Scott L DuVall, and Olga V Patterson. 2021b. Launching into clinical space with medspacy: a new clinical text processing toolkit in python. *arXiv preprint arXiv:2106.07799*.

Kexin Huang, Jaan Altosaar, and Rajesh Ranganath. 2019. Clinicalbert: Modeling clinical notes and predicting hospital readmission. *arXiv:1904.05342*.

Mark Hughes, Irene Li, Spyros Kotoulas, and Toyotaro Suzumura. 2017. Medical text classification using convolutional neural networks. *CoRR*, abs/1704.06841.

AEWP Johnson, Tom Pollard, Roger Mark, Seth Berkowitz, and Steven Horng. 2019. Mimic-cxr database. *PhysioNet10*, 13026:C2JT1Q.

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.

Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinform.*, 36(4):1234–1240.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.

Fei Li and Hong Yu. 2020. ICD coding from clinical text using multi-filter residual convolutional neural network. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8180–8187. AAAI Press.

Irene Li, Jessica Pan, Jeremy Goldwasser, Neha Verma, Wai Pan Wong, Muhammed Yavuz Nuzumlali, Benjamin Rosand, Yixin Li, Matthew Zhang, David Chang, Richard Andrew Taylor, Harlan M. Krumholz, and Dragomir R. Radev. 2021a. Neural natural language processing for unstructured data in electronic health records: a review. *CoRR*, abs/2107.02975.

Irene Li, Prithviraj Sen, Huaiyu Zhu, Yunyao Li, and Dragomir Radev. 2021b. Improving cross-lingual text classification with zero-shot instance-weighting. In *Proceedings of the 6th Workshop on Representation Learning for NLP (RepL4NLP-2021)*, pages 1–7.

Irene Li, Michihiro Yasunaga, Muhammed Yavuz Nuzumlali, Cesar Caraballo, Shiwani Mahajan, Harlan M. Krumholz, and Dragomir R. Radev. 2019. A neural topic-attention model for medical term abbreviation disambiguation. *CoRR*, abs/1910.14076.

Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.

George Michalopoulos, Yuanxin Wang, Hussam Kaka, Helen H. Chen, and Alexander Wong. 2021. Umlsbert: Clinical domain knowledge augmentation of contextual embeddings using the unified medical language system metathesaurus. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 1744–1753. Association for Computational Linguistics.

Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, pages 404–411.

Mark Neumann, Daniel King, Iz Beltagy, and Waleed Ammar. 2019. ScispaCy: Fast and Robust Models for Biomedical Natural Language Processing. In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 319–327, Florence, Italy. Association for Computational Linguistics.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 311–318, Philadelphia. Association for Computational Linguistics.

Long N Phan, James T Anibal, Hieu Tran, Shaurya Chanana, Erol Bahadroglu, Alec Peltekian, and Grégoire Altan-Bonnet. 2021. Scifive: a text-to-text transformer model for biomedical literature. *arXiv preprint arXiv:2106.03598*.

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. *J. Mach. Learn. Res.*, 21:140:1–140:67.

Nedunchelian Ramanujam and Manivannan Kaliappan. 2016. An automatic multidocument text summarization approach based on naive bayesian classifier using timestamp strategy. *The Scientific World Journal*, 2016.

Benjamin Shickel, Patrick J. Tighe, Azra Bihorac, and Parisa Rashidi. 2018. Deep EHR: A survey of recent advances in deep learning techniques for electronic health record (EHR) analysis. *IEEE J. Biomed. Health Informatics*, 22(5):1589–1604.

Sam Shleifer and Alexander M. Rush. 2020. Pre-trained summarization distillation. *CoRR*, abs/2010.13002.

Bosheng Song, Fen Li, Yuansheng Liu, and Xiangxiang Zeng. 2021. Deep learning methods for biomedical named entity recognition: a survey and qualitative comparison. *Briefings Bioinform.*, 22(6).

Shirly Wang, Matthew BA McDermott, Geeticka Chauhan, Marzyeh Ghassemi, Michael C Hughes, and Tristan Naumann. 2020. Mimic-extract: A data extraction, preprocessing, and representation pipeline for mimic-iii. In *Proceedings of the ACM conference on health, inference, and learning*, pages 222–235.

Wei-Hung Weng, Yu-An Chung, and Peter Szolovits. 2019. Unsupervised clinical language translation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019*, pages 3121–3131. ACM.

Cao Xiao, Edward Choi, and Jimeng Sun. 2018. Opportunities and challenges in developing deep learning models using electronic health records data: a systematic review. *J. Am. Medical Informatics Assoc.*, 25(10):1419–1428.

Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. 2020. Big bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems*, 33:17283–17297.

Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR.

Yuhao Zhang, Yuhui Zhang, Peng Qi, Christopher D Manning, and Curtis P Langlotz. 2021. Biomedical and clinical english model packages for the stanza python nlp library. *Journal of the American Medical Informatics Association*, 28(9):1892–1899.

Tong Zhou, Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao, Kun Niu, Weifeng Chong, and Shengping Liu. 2021. Automatic ICD coding via interactive shared representation networks with self-distillation mechanism. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 5948–5957. Association for Computational Linguistics.

A Limitations

As our work is a tool for processing clinical texts, we do not propose model-based novelty as one of our main contributions. Users may find that we conducted evaluations and built high-level user interfaces instead of proposing new models.

As EHRKit relies on many other existing libraries, we suggest that users install compatible and correct versions for robust usage.

B Potential Risk

This work is an open-source tool for clinical text processing. We did not use any user-sensitive data for training or testing, and this tool does not contain any related functionalities. Users should avoid using such data as inputs.

C Experiments

Our models were trained on a 4 Nvidia 3090 GPUs with a batch size of 8. We train all of our models using Adagrad with 0.15 learning rate and have an accumulator of 0.1.

C.1 Machine Translation

The training time varies on language pairs. The total trial, training and evaluation time is about 60 ours.

C.2 Summarization

During training we are regularly measuring the loss and the ROUGE-1 F-score on the validation set of the dataset in order to monitor the learning of our model. We end the training when the validation loss stops improving. The overall trial, training and evaluation time is about 20 hours.