

Disease Normalization with Graph Embeddings

D. Pujary^{1,2}, C. Thorne² and W. Aziz¹

¹ University of Amsterdam ² Elsevier
drv022@gmail.com c.thorne.1@elsevier.com w.aziz@uva.nl

Abstract. The detection and normalization of diseases in biomedical texts are key biomedical natural language processing tasks. Disease names need not only be identified, but also normalized or linked to clinical taxonomies describing diseases such as MeSH[®]. In this paper we describe deep learning methods that tackle both tasks. We train and test our methods on the known NCBI disease benchmark corpus. We propose to represent disease names by leveraging MeSH[®]'s graphical structure together with the lexical information available in the taxonomy using graph embeddings. We also show that combining neural named entity recognition models with our graph-based entity linking methods via multitask learning leads to improved disease recognition in the NCBI corpus.

Keywords. Disease named entity normalization, biLSTM-CRF models, graph embeddings, multi-task learning.

1 Introduction

In biomedical search applications, e.g. clinical search applications, it is of key importance to query not simply by keywords but rather by *concept*, viz., resolve names into their synonyms and return all matching documents. This is particularly true regarding diseases [4]. When we issue a query about “ovarian cancer”, we expect the system to return hits mentioning standard synonyms such as “ovarian neoplasms”. To achieve this it is necessary not only to perform *named entity recognition* (NER) to detect all entities mentioned in a document, but also to disambiguate them against databases of canonical names and synonyms, a task known as *entity normalization* or *linking* (EL).

Detecting and normalizing disease mentions are challenging tasks due to linguistic variability (e.g. abbreviations, morphological and orthographic variation, word order) [16, 24]. Interest in these tasks led to the creation of disease knowledge bases and annotated corpora. The *Medical Subject Headings* (MeSH[®]) taxonomy [18] is a repository of clinically-relevant terms covering (among others) a large range of standardized disease names, where standard names and known synonyms are organised hierarchically. The NCBI disease corpus [5] is a collection of PubMed[®] abstracts with disease mentions manually resolved to MeSH[®] or OMIM[®] concepts.

Annotated resources allow the training of supervised learning algorithms. The first such system was [16], which used a learning-to-rank approach. More recently,

Identification of APC2, a homologue of the [<i>adenomatous polyposis coli tumour</i>] ^{D011125} suppressor.	
MeSH Heading	Adenomatous Polyposis Coli
Scope Note	A polyposis syndrome due to an autosomal dominant mutation of the APC genes (GENES, APC) on CHROMOSOME 5. (...)
Tree Number(s)	C04.557.470.035.215.100 (...)
Entry Term(s)	Polyposis Syndrome, Familial (...)

Fig. 1: (Top) An NCBI corpus excerpt: title of PubMed[®] abstract 10021369. In italics, the mention, in red the MeSH[®] identifier. (Bottom) Snapshot of MeSH[®] entry for disease concept D011125. The “Scope Note” is a short (10-20 tokens) definition of the concept.

systems based on deep learning have been proposed [17, 2, 6]. The state-of-the-art for NER employs bidirectional recurrent neural networks to parameterize a conditional random field [15]. Normalization is typically based on learning a similarity metric between disease mentions and their standard names in MeSH[®], where mentions and names are represented by a composition of word embeddings [17, 2]. This typically requires first identifying potential candidate diseases, for example by exploiting handcrafted rules [17] or bidirectional recurrent neural networks [7]. While [6] leverages in addition a coherence model that enforces a form of joint disambiguation. Crucially, previous work has consistently ignored the hierarchical structure of MeSH[®] turning the taxonomy into a flat list of concepts.

In this paper we study an alternative approach to disease normalization. In particular, we acknowledge the hierarchical structure of the MeSH[®] and employ graph encoders—both supervised and unsupervised—to learn representations of disease names (i.e. MeSH[®] nodes). To make node representations easier to correlate to textual input, we propose to also exploit lexical information available in scope notes. Finally, given the similarities between NER and EL, we explore multitask learning [MTL; 1] via a shared encoder. MTL has been traditionally applied in biomedical NLP to solve NER over several, though related corpora, covering different, but related entities [25], but has seldom been used to learn related tasks over the same corpus and entities. Our findings suggest that using lexicalized node embeddings improves EL performance, while MTL allows to transfer knowledge from (graph-embedding) EL to NER, with performance topping comparable approaches for the NCBI corpus [9, 25].

2 Datasets and Methods

Knowledge Base MeSH[®] is a bio-medical controlled vocabulary produced by the National Library of Medicine [18] covering a large number of health concepts, including 10,923 disease concepts, and is used for indexing, cataloging and

Table 1: NCBI disease corpus statistics.

Split	Abstracts	Total entities	Unique entities	Unique concept IDs	Tokens
Training	592	5,134	1,691	657	136,088
Validation	100	787	363	173	23,969
Test	100	960	424	201	24,497

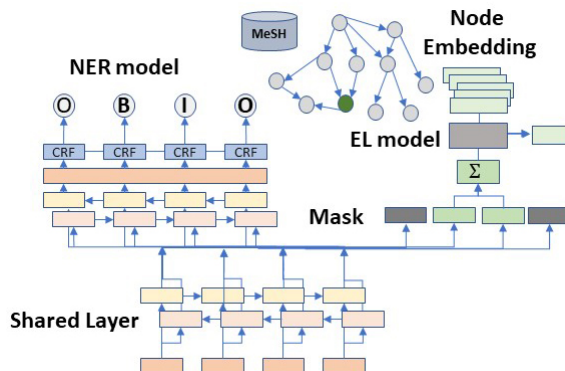


Fig. 2: Visualization of the MTL architecture.

searching bio-medical and health-related documents and information. There are MeSH[®] records of several types divided further into several categories. In this paper we consider only category *C Descriptors*, in other words on *diseases* MeSH[®] disease nodes are arrayed hierarchically (in the form of an acyclic directed graph or tree) from most generic to most specific in up to thirteen hierarchical levels, and have many attributes, such as, *MeSH Heading*, *Tree Number(s)*, *Unique ID*, *Scope Note* and *Entry Term(s)*. See Figure 1 (bottom)³.

Supervision We use the NCBI disease corpus [5] for training and validating our methods. It consists of 792⁴ PubMed[®] abstracts separated into training, validation and test subsets (see Table 1). The corpus is annotated with disease mentions mapped to either MeSH[®] or OMIM[®] [10] (whose identifiers we convert to MeSH[®] identifiers if possible using CTD [3]), Figure 1 (top) illustrates an annotated instance in the corpus.

NER For NER we use variants of [15] and [19] biLSTM-CRF models. First, we encode words in an abstract using either a fixed and pre-trained contextualized encoder, i.e. bioELMo [11], or a concatenation of a fixed and pre-trained biomedical embedding, i.e. SkipGram [22], and a trainable character-level encoder (e.g. LSTM [15] or CNN [19]). We then re-encode every token with a trainable biLSTM whose outputs are linearly transformed to the potentials of a linear-chain

³ Full MeSH[®] example: <https://meshb.nlm.nih.gov/record/ui?ui=D011125>

⁴ We drop one training abstract because it is repeated: abstract 8528200.

CRF [14]. All trainable parameters are optimized to maximize the likelihood of observations (i.e. IOB-annotated spans). At prediction time, Viterbi decoding finds the optimum IOB sequence in linear time.

EL with *node2vec* embeddings We want to capitalize on the structure of the disease branch of MeSH[®]. Thus, for every disease node, we learn an embedding with *node2vec* [21, 8]. The algorithm starts from randomly initialized node embeddings which are then optimized to be predictive of neighbourhood in the MeSH[®] graph (for this we ignore the directionality inherit to relations in MeSH[®]). We use the *node2vec* framework [8] to generate two types of representations. For *type-I*, we use *node2vec* as is, that is, with randomly initialized node embeddings. For *type-II*, we leverage lexical information by initializing a node’s embedding with the bioELMo encoding of that node’s *scope note* (we use average pooling). The idea behind *type-II* embeddings is to combine lexical features in the scope note and structural relations in MeSH[®]. Our EL model is a logistic regression classifier parameterized by a bilinear product between the embedding $\text{emb}(y)$ of a node y and the contextualized encoding $\text{enc}(x, m)$ of a mention m in an abstract x , i.e. $p(y|x, m, \theta) \propto \exp(\text{enc}(x, m)^\top \mathbf{W} \text{emb}(y))$, where $\theta = \{\mathbf{W}\}$ are trainable parameters. The node embedding is either a *type-I* or *type-II* embedding, and the mention embedding is the average of the bioELMo contextualized representation of the tokens in the disease mention. Unlike previous work, we normalize $p(y|x, m, \theta)$ against the set of all known entities in MeSH[®], that is, without a pre-selection of candidates. The model is optimized to maximize the likelihood of observations—triples (m, x, y) of mentions in an abstract resolved to MeSH[®] nodes.

EL with GCN encoders Graph convolution networks [GCN; 13] generalize CNNs to general graphs. A GCN is an encoder where node representations are updated in parallel by transforming messages from adjacent nodes in a graph structure. A message is nothing but a parameterized (possibly non-linear) transformation of the representation of the neighbour. Stacking L GCN layers allows information to flow from nodes as far as L hops away in graph structure. In particular we follow [20], where the encoding $\mathbf{h}_v^{(\ell)}$ of a node v in the taxonomy is computed as $\mathbf{h}_v^{(\ell)} = \sigma \left(\sum_{u \in \mathcal{N}(v)} \mathbf{W}^{(\ell)} \mathbf{h}_u^{(\ell-1)} + \mathbf{b}^{(\ell)} \right)$, where ℓ indexes the layer, $\mathcal{N}(v)$ are neighbors of v in MeSH[®] (we again discard directionality), $\mathbf{b}^{(\ell)}$ and $\mathbf{W}^{(\ell)}$ are the parameters of the ℓ th GCN layer, σ is a nonlinearity (we use ReLU), and $\mathbf{h}_v^{(0)} = \text{emb}(v)$ is the initial embedding of the disease node v . For the 0th representation of nodes, we use the bioELMo encoding of a node’s scope note. As a trainable encoder, a GCN layer is an integral part of the classifier and its parameters are updated by backpropagation from the downstream objective. Once again, our objective is to maximize the likelihood of observations under a probabilistic classifier, i.e. $p(y|x, m, \theta) \propto \exp(\text{enc}(x, m)^\top \text{gcn}(y; \theta))$, where $\text{gcn}(y; \theta) = \mathbf{h}_y^{(L)}$ and $\theta = \{(\mathbf{W}^{(\ell)}, \mathbf{b}^{(\ell)})_{\ell=1}^L\}$ are the trainable parameters. Note that unlike *node2vec* embeddings, GCN encoders are trained directly on EL supervision, which is arguably of limited availability.

Table 2: Results on test set and validation set of the NER experiments. We use [15] and [19] models as a baseline, and compare with Habibi’s et al. [9] known re-implementation of [15]. For bioELMo we use the same architecture common to both the baseline. Results reported are the mean and standard deviation of 5 runs. *Our implementation. **Result taken from original paper.

Encoder	Pre	Rec	F1	Val. F1
Lample et al*	0.824 ± 0.022	0.742 ± 0.019	0.781 ± 0.003	0.805 ± 0.007
Ma and Hovy*	0.823 ± 0.011	0.776 ± 0.023	0.799 ± 0.012	0.792 ± 0.005
bioELMo	0.878 ± 0.003	0.856 ± 0.005	0.867 ± 0.002	0.884 ± 0.001
Lample et al.** [9]	0.875**	0.836**	0.844**	-

Multitask Learning Traditionally, the problem of disease normalization is tackled by first identifying the disease names (NER) and then normalizing them (EL). We attempt to learn from both types of supervision by having a NER and an EL model share parts of their architectures. This is known as *multitask learning* [1]. In particular, we share the encoder of the NER architecture, see Figure 2, and derive mention features for the EL model from there.

3 Experiments and Results

Hyperparameters We train our models using ADAM [12] and initial learning rate of 10^{-3} . We split our abstracts into sentences using NLTK⁵ and process sentences independently. We use 200-dimensional *word2vec* embeddings and 1,024-dimensional bioELMo encodings. For NER models, we learn 60-dimensional character embeddings. We train MeSH[®] node embeddings using *node2vec* for 100 epochs, whereas GCN-based EL models are trained for 500 epochs. All models employ 1,024-dimensional node embeddings. We stack $L = 2$ GCN layers with 2,048 hidden and 1,024 output units. We use 0.5 dropout regularization [23]⁶.

Metrics and Results We report precision (Pre), recall (Rec) and (micro averaged) F1 scores (F1) for NER. For EL it is customary to test whether the target canonical entity (a MeSH[®] concept identifier) occurs among the top k predictions [6]. Hence we report precision at confidence rank k (Pre@ k) and mean reciprocal rank (MRR) defined as $MRR = 1/|E| \cdot \sum\{1/\text{rank}_i \mid 1 \leq i \leq |E|\}$. We use early stopping for model selection based on performance on the validation set, and report the average result of 5 independent runs on the test and validation sets.

Table 2 shows the results for NER. By replacing the word and character embeddings altogether with bioELMo embeddings, the results improve by a large

⁵ <https://www.nltk.org/>

⁶ The code of our experiments is available at: <https://github.com/druv022/Disease-Normalization-with-Graph-Embeddings>.

Table 3: Results on test set and validation set of our EL models for different type of MeSH encoding. Results reported are the mean and standard deviation of 5 runs. We consider unlexicalized *node2vec* as our baseline. *Result taken from original paper. **We cannot compare fully due to different methodology. “S.N” refers to Scope Note.

Disease emb.	MRR	F1	Pre	Pre@30	Val. MRR
bioELMo (S.N.)	0.748 ± 0.002	0.715 ± 0.004	0.715 ± 0.002	0.844 ± 0.004	0.791 ± 0.001
<i>node2vec</i> I	0.749 ± 0.002	0.718 ± 0.004	0.720 ± 0.004	0.819 ± 0.006	0.800 ± 0.003
<i>node2vec</i> II	0.757 ± 0.001	0.721 ± 0.004	0.724 ± 0.001	0.842 ± 0.004	0.804 ± 0.006
GCN	0.744 ± 0.006	0.710 ± 0.008	0.710 ± 0.007	0.831 ± 0.005	0.803 ± 0.007
DNorm* [16]	-	0.782**	-	-	-
NormCo* [6]	-	0.840**	0.878**	-	-

margin. Increasing the number of parameters in the models did not affect model performance. Furthermore, we improve on [9]’s re-implementation of Lample et al.’s biLSTM-CRF model on the NCBI corpus (+0.028 F1 score points).

We tested our EL baseline—*node2vec* type-I, encoding only the MeSH[®] structural information—against disease embeddings generated by averaging the bioELMo embedding of Scope Notes (see Table 3), and found it to perform better except on Pre@30. Node lexicalization, *node2vec* type-II, yields results on a par or better than the bioELMo baseline. Using GCN for training, we do not achieve any improvement. This suggests that both structural and lexical information in MeSH[®] are important for normalization, and that taxonomy encoding is best achieved when computed independently from the EL task. This strategy seems on the other hand to underperform other normalization methods reported in the literature [6, 16]. Our systems however do not incorporate many of their optimizations such as: normalization to disease synonym rather than MeSH[®] identifier, abbreviation resolution, re-ranking and filtering, or coherence models.

Table 4 shows, finally, that using MTL results in an improved NER score (+0.009 F1 score points w.r.t. our best NER model in Table 2). This indicates a certain degree of transfer from the EL to the NER task. It also yields a improvement w.r.t. similar NER and MTL for the NCBI corpus [9, 25] (+0.028 points for NER and +0.015 for MTL).

Error analysis Regarding NER, most errors arise when dealing with multi-token entities: e.g., for *sporadic T-PLL* only the head *T-PLL* is detected. On the other hand, the bioELMo model is able to identify abbreviated disease names much better than the two baselines: e.g., *T-PLL* occurring alone is always correctly detected. This might be because bioELMo starts at character-level and gradually captures token- and phrase-level information.

Our best EL model makes errors by confusing diseases with their MeSH[®] ancestors and neighbors: e.g., it confuses *D016399- Lymphoma*, *T-Cell* with *D015458- Leukemia*, *T-Cell*, with which it shares an ancestor: *D008232- Lym-*

Table 4: Results on test set of our MTL experiment where we report precision, recall and F1-score for NER task, and for EL we report MRR and precision. Results reported are the mean and standard deviation of 3 runs. *Result taken from original paper.

Model	NER			EL	
	Pre	Rec	F1	MRR	Pre@30
NER & GCN	0.880 ± 0.003	0.872 ± 0.008	0.876 ± 0.003	0.747 ± 0.003	0.816 ± 0.006
NER	0.875 ± 0.006	0.869 ± 0.001	0.872 ± 0.003	-	-
GCN	-	-	-	0.745 ± 0.001	0.816 ± 0.001
Wang* [25]	0.857 ± 0.009*	0.864 ± 0.004*	0.861 ± 0.003*	-	-

phoproliferative Disorders. Often, it resolves correctly the first instance of a concept but returns instead a neighbour afterwards: e.g., in “*Occasional missense mutations in ATM were also found in tumour DNA from patients with B-cell non-Hodgkins lymphomas (B-NHL) and a B-NHL cell line*” *B-cell non-Hodgkins lymphomas* is correctly resolved to *D016393* but *B-NHL* (its abbreviation) is mapped to *D008228* (its child) and the second occurrence to *D018239* (another form of cancer). This might be due to the following facts: Both NCBI and MeSH[®] are not particularly large resources, hence the mention and disease encodings learnt by the model are not sufficiently discriminative. In addition, GCN and *node2vec* training tends to ignore the direction of MeSH[®] taxonomical edges. Finally, DNorm and NormCo normalize diseases to concepts indirectly via their synonyms, allowing them to exploit better lexical features⁷.

4 Conclusions

In this paper, we address the problem of normalizing disease names to MeSH[®] concepts (canonical identifiers) by adapting state-of-the-art neural graph embeddings (GCN and *node2vec*) that exploit both MeSH[®]’s hierarchical structure and the description of diseases. Our graph-based disease node encoding is the first of its kind, to best of our knowledge. We also apply multi-tasking to transfer information between disease detection (NER) and resolution (EL) components of the task, leveraging their common signals to improve on the single models. We observe that bioELMo embeddings lead to substantial improvement in NER performance. We demonstrate that node lexicalization does improve over either pure structural or lexical embeddings, and that MTL gives rise to state-of-the-art performance for NER in the corpus used (NCBI corpus).

On the other hand, we do not currently outperform other disease normalization approaches [6, 16], and often confuse neighbour MeSH[®] concepts with their true targets. In the future we intend to experiment with the optimizations reported (especially: resolution to synonym rather than concept identifier and

⁷ They also include a wide range of optimizations such as re-ranking, coherence models or abbreviation resolution.

re-ranking) by [6, 16]. MeSH[®] is also a comparatively small graph with short disease descriptions. As further work we plan to enrich it by linking it to larger scale resources (e.g., Wikipedia).

References

- [1] Rich Caruana. Multitask learning. *Machine Learning*, 28(1):41–75, 1997.
- [2] Hyejin Cho, Wonjun Choi, and Hyunju Lee. A method for named entity normalization in biomedical articles: Application to diseases and plants. *BMC Bioinformatics*, 18(1):451, 2017.
- [3] Allan Peter Davis, Cynthia J. Grondin, Robin J. Johnson, Daniela Sciaky, Roy McMorran, Jolene Wieggers, Thomas C. Wieggers, and Carolyn J. Mattingly. The Comparative Toxicogenomics Database: update 2019. *Nucleic Acids Research*, 47(D1):D948–D954, 2018.
- [4] Rezarta Islamaj Dogan and Zhiyong Lu. An inference method for disease name normalization. In *AAAI Fall Symposium: Information Retrieval and Knowledge Discovery in Biomedical Text*, 2012.
- [5] Rezarta Islamaj Doğan, Robert Leaman, and Zhiyong Lu. NCBI disease corpus. *Journal of Biomedical Informatics*, 47(C):1–10, 2014.
- [6] Raghav Mehta Chun-Nan Hsu Dustin Wright, Yannis Katsis. NormCo: Deep disease normalization for biomedical knowledge base construction. In *Proceedings of AKBC 2019*, 2019.
- [7] Nathan Greenberg, Trapit Bansal, Patrick Verga, and Andrew McCallum. Marginal likelihood training of BiLSTM-CRF for biomedical named entity recognition from disjoint label sets. In *Proceedings of EMNLP 2018*, 2018.
- [8] Aditya Grover and Jure Leskovec. Node2Vec: Scalable feature learning for networks. In *Proceedings of KDD 2016*, 2016.
- [9] Maryam Habibi, Leon Weber, Mariana Neves, David Luis Wiegandt, and Ulf Leser. Deep learning with word embeddings improves biomedical named entity recognition. *Bioinformatics*, 33(14):i37–i48, 2017.
- [10] Ada Hamosh, Alan F. Scott, Joanna S. Amberger, Carol A. Bocchini, and Victor A. McKusick. Online Mendelian Inheritance in Man (OMIM), a knowledge base of human genes and genetic disorders. *Nucleic Acids Research*, 33(suppl.1):D514–D517, 2005.
- [11] Qiao Jin, Bhuwan Dhingra, William W Cohen, and Xinghua Lu. Probing biomedical embeddings from language models. *CoRR*, abs/1904.02181, 2019.
- [12] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. 2014. Proceedings of ICLR 2014.
- [13] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *CoRR*, abs/1609.02907, 2016.
- [14] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of ICML 2001*, 2001.

- [15] Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. Neural architectures for named entity recognition. *CoRR*, abs/1603.01360, 2016.
- [16] Robert Leaman, Rezarta Islamaj Doğan, and Zhiyong Lu. DNorm: disease name normalization with pairwise learning to rank. *Bioinformatics*, 29(22):2909–2917, 2013.
- [17] Haodi Li, Qingcai Chen, Buzhou Tang, Xiaolong Wang, Hua Xu, Baohua Wang, and Dong Huang. CNN-based ranking for biomedical entity normalization. *BMC Bioinformatics*, 18(11):385, 2017.
- [18] C.E. Lipscomb. Medical subject headings (MeSH). *Bull Medical Library Association*, 88(3):265–266, 2000.
- [19] Xuezhe Ma and Eduard H. Hovy. End-to-end sequence labeling via Bidirectional LSTM-CNNs-CRF. *CoRR*, abs/1603.01354, 2016.
- [20] Diego Marcheggiani and Ivan Titov. Encoding sentences with graph convolutional networks for semantic role labeling. *CoRR*, abs/1703.04826, 2017.
- [21] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. DeepWalk: Online learning of social representations. In *Proceedings of KDD 2014*, 2014.
- [22] S. Pyysalo, F. Ginter, H. Moen, T. Salakoski, and S. Ananiadou. Distributional semantics resources for biomedical text processing. In *Proceedings of LBM 2013*, 2013.
- [23] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [24] Camilo Thorne and Roman Klinger. On the semantic similarity of disease mentions in PubMed[®] and Twitter. In *Proceedings of NLDB 2018*, 2018.
- [25] Xuan Wang, Yu Zhang, Xiang Ren, Yuhao Zhang, Marinka Zitnik, Jingbo Shang, Curtis Langlotz, and Jiawei Han. Cross-type biomedical named entity recognition with deep multi-task learning. *Bioinformatics*, 35(10):1745–1752, 2018.