Probing the EHR for Standardized Nursing Data

Baris Karacan
University of Illinois Chicago
Chicago, Illinois, USA

Andrew Boyd
University of Illinois Chicago
Chicago, Illinois, USA

Karen Dunn Lopez
University of Illinois Chicago
Chicago, Illinois, USA

Pamela Martyn-Nemeth
University of Illinois Chicago
Chicago, Illinois, USA

Daniel Fraczkowski
University of Illinois Chicago
Chicago, Illinois, USA

Haleh Vatani
University of Illinois Chicago
Chicago, Illinois, USA

Carolyn Dickens
University of Illinois Chicago
Chicago, Illinois, USA

Barbara Di Eugenio
University of Illinois Chicago
Chicago, Illinois, USA

ABSTRACT

Nursing documentation is essential for the welfare of patients and for productive communication between healthcare professionals. Currently, nursing care is documented by means of standardized and specific non-standardized nursing terminologies that various healthcare companies provide. Because of significant differences between terminologies, nursing professionals devote considerable time to map distinct terminologies by manually searching terminology databases or books. We present an automated approach that finds mappings between terminologies of two widely-used nursing care plans: it is based on UMLS as an intermediate resource, and on similarity computed via language models. According to our nursing team experts, our best-performing model found accurate mappings for approximately 54 percent of terms.

1 INTRODUCTION

Proper nursing care documentation is essential for effective communication among healthcare professionals and for monitoring patients' health status and intervention efficacy. Despite attempts to establish uniform terminologies, significant differences persist, making this a complex task. The American Nurses Association (ANA) recognizes three standardized nursing terminologies for representing nursing care [12]: North American Nursing Diagnoses Association International (NANDA-I) [6], Nursing Interventions Classification (NIC) [2], and Nursing Outcomes Classification (NOC) [8]. NANDA-I is for documenting diagnoses, while NIC and NOC are for planning interventions and describing outcomes.

In addition to the ANA, private healthcare companies such as EPIC [3] have developed their nursing plan-of-care terminologies. EPIC's documentation system uses care plan problems, goals, and interventions to represent nursing diagnoses, outcomes, and interventions, respectively. The IRB-reviewed and approved study presented here proposes a system that effectively maps the EPIC nursing plan-of-care terminologies to the corresponding standardized nursing care plan terminologies (NANDA-I, NIC, and NOC).

CHASE '23, June 21–23, 2023, Orlando, FL, USA

ACM ISBN 979-8-4007-0102-3/23/06.
https://doi.org/10.1145/3580252.3590001

2 DATASET

In this research, 130 heart failure patients' EPIC nursing care plans were analyzed, sourced from the tertiary academic medical center the researchers belong to. The care plans contained problems, goals, and interventions as subcomponents. A total of 54 unique care plan problems, 99 unique care plan goals, and 476 unique care plan interventions were extracted from 1450 care plan entries and utilized separately for the mapping task.

3 METHODOLOGY

Our two-step approach involves extracting candidate NANDA-I, NOC, and NIC terms using UMLS [11] and computing their similarities with EPIC terms utilizing different language models. UMLS, a collection of biomedical terminologies, covers most NANDA-I, NIC, and NOC terms, making it suitable for mapping tasks. After retrieving care plan terms from the dataset, we utilized the search functionality of the UMLS REST API to query these terms. To ensure accurate outputs, we only included results that belonged to the target terminology, such as NANDA-I, NIC, or NOC. We obtained results by combining matches for each query word within the care plan terms. For example, when searching for the care plan problem "Mechanical Ventilation," UMLS would return all NANDA-I results containing the words "mechanical" and "ventilation" separately.

In the first step of our approach, we recognized that some results were not found in the target terminologies' hierarchy. To address this, we retrieved the immediate parents and children for each target term, removing any terms without parents or children. We then explored the remaining terms' children and extracted the terms at the lowest level of the hierarchy, as they provide more specific information than their parent terms. For instance, in the case of the NANDA-I diagnosis "Self-care" for the care plan problem "Daily Care," we inspected its leaf children to gain specific information about the patient's required self-care type (Figure 1 depicts the taxonomy of "Self-care"). However, for care plan interventions, we discovered that matching NIC terms were granular activities that did not accurately represent NICs as a whole. To resolve this, we retrieved the immediate parents of these activities to obtain the actual NICs. Finally, we selected these actual NIC terms as potential candidates for the mapping process.

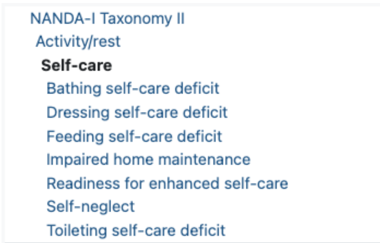

Figure 1: Taxonomy of NANDA-I term "Self-care"

NANDA-I Diagnosis	Score
Impaired Tissue Integrity	0.94
Risk for Impaired Tissue Integrity	0.80
Impaired Skin Integrity	0.60
Risk for Impaired Skin Integrity	0.48
Ineffective Peripheral Tissue Perfusion	0.46

Table 1: Top 5 NANDA-I diagnoses for EPIC "Tissue Integrity"

The second step in our approach involves using text similarity measures to compute the similarity between the source EPIC care plan term and its respective candidate terms. As each term is represented by a single phrase, we vectorized them to capture their semantic information by creating individual phrase embeddings. These embeddings can then be compared using an appropriate similarity measure to find terms with similar meanings. To achieve this, we utilized several existing language models, including BERT [4] and five of its variants: Sentence-BERT [10], Bio-BERT [7], Bio-Clinical-BERT [1], Blue-BERT [9], and PubMed-BERT [5]. BERT and Sentence-BERT are general-purpose language models, while the remaining models are domain-specific, pre-trained on various biomedical corpora to better understand language representations in our dataset. We employed all mentioned models via HuggingFace Transformers and calculated pairwise similarity scores between source and candidate term embeddings using cosine similarity. Our team's nursing experts assessed the top ten candidate terms for each model. Table 1 displays an example of the top five NANDA-I diagnoses generated by the Sentence-BERT model, along with their cosine similarity scores.

4 EVALUATION

Two nursing professionals in our team manually mapped 31 of 54 care plan problems to NANDA-I terms. However, 15 problems were too broad or vague for mapping, while eight required further discussion. At least one language model captured 29 of the experts' suggested mappings. The remaining two suggestions were present in UMLS but could not be captured due to a lack of overlapping words with the care plan problems used in the UMLS REST API query. For the 29 accurately mapped care plan problems, experts made 39 different NANDA-I suggestions. Sentence-BERT and Bio-Clinical-BERT captured 36 suggestions, outperforming BERT-base (35), Bio-BERT (34), Blue-BERT (34), and PubMed-BERT (33). Domain-specific models did not outperform general-purpose models, as they were predominantly trained on physician terms.

5 DISCUSSION AND CONCLUSIONS

Our approach showed promise for care plan problems and their mappings to NANDA-I terms. However, goals and outcomes are more challenging due to their length and complexity. For instance, the longest care plan problem in our dataset, "Chronic Conditions and Co-morbidities," only contains four words, while average goals and interventions have over ten words. This generates more candidates in UMLS queries and leads to performance degradation.

In the EPIC system, the "care plan templates" determine the domain of a care plan problem. However, we did not receive access to them from the EPIC data experts. Given our focus on heart failure patients, we need to restrict our analysis to heart failure-related care. For example, the "Respiratory Status" care plan problem in EPIC is only linked to domains unrelated to heart failure. Therefore, we excluded goals and outcomes associated with this problem from our study.

Consequently, evaluating mappings for goals and interventions without care plan templates is not trivial. We are currently working on retrieving care plan templates for our dataset, as it is critical to eliminate irrelevant terms. Note that we do not have direct access to patient data, but we need to obtain it through the appropriate procedures within our health system. Finally, the overall performance of our approach is bound by the capabilities of the UMLS search algorithm. Therefore, the resulting target term needs to share at least one common word with the source term. Otherwise, the target term would not be among the candidate terms generated by UMLS.

REFERENCES

[1] Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew McDermott. 2019. Publicly available clinical BERT embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*. Association for Computational Linguistics, Minneapolis, Minnesota, USA, (June 2019), 72–78. DOI: 10.18653/v1/W19-1909.

[2] Howard K Butcher, Gloria M Bulechek, Joanne McCloskey Dochterman, and Cheryl M Wagner. 2018. *Nursing interventions classification (NIC)-E-Book*. Elsevier Health Sciences.

[3] Epic Systems Corporation. 2022. Epic electronic health record. Available at: https://www.epic.com/. (2022).

[4] 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)*. Association for Computational Linguistics, (June 2019), 4171–4186. DOI: 10.18653/v1/N19-1423.

[5] Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2020. Domain-specific language model pretraining for biomedical natural language processing. (2020). eprint: arXiv:2007.15779.

[6] T Heather Herdman. 2011. *Nursing diagnoses 2012-14: definitions and classification*. John Wiley & Sons.

[7] 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. *Bioinformatics*, 36, 4, 1234–1240.

[8] Sue Moorhead, Marion Johnson, Meridean L Maas, and Elizabeth Swanson. 2018. *Nursing Outcomes Classification (NOC)-e-book: Measurement of health outcomes*. Elsevier Health Sciences.

[9] Yifan Peng, Shankai Yan, and Zhiyong Lu. 2019. Transfer learning in biomedical natural language processing: an evaluation of bert and elmo on ten benchmarking datasets. In *Proceedings of the 2019 Workshop on Biomedical Natural Language Processing (BioNLP 2019)*, 58–65.

[10] Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, (Nov. 2019). https://arxiv.org/abs/1908.10084.

[11] Peri L Schuyler, William T Hole, Mark S Tuttle, and David D Sherertz. 1993. The umls metathesaurus: representing different views of biomedical concepts. *Bulletin of the Medical Library Association*, 81, 2, 217.

[12] Bonnie L Westra, Connie White Delaney, Debra Konicek, and Gail Keenan. 2008. Nursing standards to support the electronic health record. *Nursing outlook*, 56, 5, 258–266.