



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/196486/>

Version: Published Version

---

**Proceedings Paper:**

Taneja, S., Boyce, R., Reynolds, W. et al. (2021) Introducing information retrieval for biomedical informatics students. In: Jurgens, D., Kolhatkar, V., Li, L., Mieskes, M. and Pedersen, T., (eds.) Proceedings of the Fifth Workshop on Teaching NLP. Proceedings of the Fifth Workshop on Teaching NLP, 10-11 Jun 2021, Online. Association for Computational Linguistics, pp. 96-98. ISBN: 9781954085367.

<https://doi.org/10.18653/v1/2021.teachingnlp-1.16>

---

© 2021 Association for Computational Linguistics. Licensed on a Creative Commons Attribution 4.0 International License (<https://creativecommons.org/licenses/by/4.0/>).

**Reuse**

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.

# Introducing Information Retrieval for Biomedical Informatics Students

Sanya B. Taneja    Richard D. Boyce    William T. Reynolds    Denis Newman-Griffis

University of Pittsburgh

Pittsburgh, PA, USA

{sbt12, rdb20, wtr8, dnewmangriffis}@pitt.edu

## Abstract

Introducing biomedical informatics (BMI) students to natural language processing (NLP) requires balancing technical depth with practical know-how to address application-focused needs. We developed a set of three activities introducing introductory BMI students to information retrieval with NLP, covering document representation strategies and language models from TF-IDF to BERT. These activities provide students with hands-on experience targeted towards common use cases, and introduce fundamental components of NLP workflows for a wide variety of applications.

## 1 Introduction

Natural language processing (NLP) technologies have become a fundamental tool for biomedical informatics (BMI) research, with uses ranging from mining new protein-protein interactions from scientific literature to recommending medications in clinical care. In most cases, however, the research question at hand is a biological or clinical one, and NLP tools are used off-the-shelf or lightly adapted rather than being the focus of research and development. When introducing BMI students to NLP, instructors must therefore navigate a balance between the computational and linguistic insights behind why NLP technologies work the way they do, and practical know-how for the application-focused settings where most students will go on to use NLP.

We developed a set of three activities designed to expose introductory BMI students to the fundamentals of NLP and provide hands-on experience with NLP techniques. These activities were designed for use in the Foundations of Biomedical Informatics I course at the University of Pittsburgh, a survey course which introduces students to core methods and topics in biomedical informatics. The course is required for all students in the Biomedical Informatics Training Program; students have a range

of experience in computer science, and no background in artificial intelligence or NLP is required. The sequence of activities, implemented as Jupyter notebooks, comprise a single assignment focused on information retrieval, a common use case for NLP in all areas of BMI research. The assignment followed lectures focused on information retrieval, word embeddings, and language models (presented online in Fall 2020). §2 gives an overview of the three activities, and we note directions for further refinement of the activities in §3.

## 2 Assignment Details

Our set of three activities was designed with two primary learning goals in mind:

- Expose introductory BMI students to fundamental strategies for text representation and language models, geared towards information retrieval in biomedical contexts; and
- Provide students with hands-on experience creating NLP workflows using pre-built tools.

Our Jupyter notebooks provide a sequence of code samples to analyze and execute, combined with background questions to assess understanding of the computational and linguistic insights behind the NLP technologies students are using.

**Notebook 1: Fundamentals of document analysis.** In this notebook, students were first introduced to basic preprocessing tasks in NLP workflows such as tokenization, stemming, casing, and stop-word removal. Using a corpus from the Natural Language Toolkit (NLTK) (Bird et al., 2009), the notebook demonstrated two indexing techniques - inverted indexing and creation of a weighted document-term matrix using term frequency-inverse document frequency (TF-IDF). Students then implemented a synthetic information retrieval task involving a collection of 12 documents mapped to 20 queries as a reference set. The students evaluated the information retrieval system

with synthetic results for two queries comprising numeric values for each document in the query. Students measured system performance using TREC evaluation measures including recall, precision, interpolated precision-recall average, and mean average precision. The evaluation measures were implemented using the `pytrec_eval` library (Van Gysel and de Rijke, 2018).

**Notebook 2: Introduction to word embeddings.** In this notebook, students were introduced to word embeddings as a text representation tool for NLP. The students first created embeddings using singular value decomposition (SVD) of a co-occurrence matrix using the corpus from Notebook 1. This established the idea of capturing semantic similarity in texts as opposed to a tradition bag-of-words model. The notebook then used pretrained word2vec embeddings (Mikolov et al., 2013) to demonstrate a more refined approach of SVD. The students were able to visualize both the embedding approaches in the notebook. This was particularly important as most students did not have prior experience with embeddings and plotting the word embeddings can lead to greater insight into the variable semantic similarity that embedding representations provide over lexical features.

**Notebook 3: Introduction to BERT and clinicalBERT.** As all the activities are designed for introductory students, YouTube tutorials were used to provide background on neural networks and design decisions in language models for students with minimal background in NLP. Students were introduced to NLP workflows and language models using the Transformers library. Transformers is an open-source library developed by Hugging Face that provides a collection of pretrained models and general-purpose architectures for natural language processing (Wolf et al., 2020), based on the Transformer architecture (Vaswani et al., 2017). This includes BERT (Devlin et al., 2019) and clinicalBERT (Alsentzer et al., 2019), which are used in this notebook to implement Named Entity Recognition and Medical Language Inference tasks.

The notebook guided the students through a Medical Language Inference task to infer knowledge from clinical texts. Language inference in medicine is an important application of NLP as clinical notes such as those containing past medical history of a patient contain vital information that is utilized by clinicians to draw useful inferences (Romanov and Shivade, 2018). The task also

introduced BMI students to challenges unique to applying NLP on clinical texts such as domain-specific vocabulary, diversity in abbreviations, content structure, and named entity recognition with clinical jargon. The students used the MedNLI (Romanov and Shivade, 2018) dataset created from MIMIC III (Johnson et al., 2016) clinical notes for this task. Building on knowledge from the previous notebooks, they implemented workflows to compare the performance of BERT and clinicalBERT models for prediction of labels in clinical texts. The students were encouraged to understand the importance of domain representation in pretraining data and the process of fine-tuning NLP models for domain-specific language.

### 3 Discussion

We designed three activities to demonstrate fundamental concepts and workflows for application of NLP to introductory BMI students. While the scope of NLP in the biomedical field is much larger than one assignment, we developed the activities to provide students with a modular workflow of components that are applicable to other NLP applications besides information retrieval; i.e., text preprocessing, indexing, execution, and evaluation.

One practical challenge for clinically-focused exercises is the limited availability of benchmark datasets. Most clinical datasets require Data Use Agreements and individual training requirements that are cumbersome for a classroom setting. Thus, the first two notebooks use the NLTK corpus which is a popular dataset for introducing NLP concepts without the challenges present in clinical datasets. While MIMIC (Johnson et al., 2016) is a valuable, relatively accessible source for clinical text, available annotations for it are limited. Students can thus be introduced to fine-tuning of language models for the specifics of medical language, but instructors must anticipate challenges in providing train/test splits for supervised machine learning.

We further take advantage of popular pre-built libraries (like Transformers) so that students can focus on the application rather than constructing neural networks. In an application-focused setting, technical knowledge of neural NLP systems is less necessary, but users of those systems still need to understand what kinds of regularities they rely on and when they may be unreliable. The activities are thus designed to reflect the perspective of the practical challenges that students will face when

working in biomedical NLP.

Finally, one important area we are investigating as we refine and extend these teaching materials is the ethical considerations of biomedical AI technologies. The intersection of medical and AI ethics poses several challenging questions for designing, training, and applying AI technologies in the biomedical setting (Char et al., 2018; Keskinbora, 2019), and sensitive information often described in medical text presents further ethical questions for NLP systems (Lehman et al., 2021). Thus, determining what BMI students have a responsibility to understand about the NLP tools they use, and how we most effectively teach that information in limited course time, is key to broadening the responsible use of NLP in BMI research.

All materials presented in this paper are available from [https://github.com/dbmi-pitt/bioinf\\_teachingNLP](https://github.com/dbmi-pitt/bioinf_teachingNLP).

## Acknowledgments

This work was supported in part by the National Library of Medicine of the National Institutes of Health under award number T15 LM007059.

## References

- Emily Alsentzer, John Murphy, William Boag, Weihung Weng, Di Jin, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. "O'Reilly Media, Inc."
- Danton S Char, Nigam H Shah, and David Magnus. 2018. [Implementing Machine Learning in Health Care - Addressing Ethical Challenges](#). *The New England journal of medicine*, 378(11):981–983.
- 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 {A}merican 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.
- 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.
- Kadircan H Keskinbora. 2019. [Medical ethics considerations on artificial intelligence](#). *Journal of Clinical Neuroscience*, 64:277–282.
- Eric Lehman, Sarthak Jain, Karl Pichotta, Yoav Goldberg, and Byron C. Wallace. 2021. [Does bert pre-trained on clinical notes reveal sensitive data?](#) *arXiv preprint arXiv:2104.07762*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. *arXiv preprint arXiv:1310.4546*.
- Alexey Romanov and Chaitanya Shivade. 2018. [Lessons from natural language inference in the clinical domain](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1586–1596, Brussels, Belgium. Association for Computational Linguistics.
- Christophe Van Gysel and Maarten de Rijke. 2018. Pytrec\_eval: An extremely fast python interface to trec\_eval. In *SIGIR*. ACM.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is All you Need](#). In I Guyon, U V Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan, and R Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.