Year 2020 (with COVID): Observation of Scientific Literature on Clinical Natural Language Processing

Natalia Grabar^{1,2}, Cyril Grouin¹, Section Editors of the IMIA Yearbook Section on Clinical Natural Language Processing

¹ Université Paris Saclay, CNRS, Laboratoire Interdisciplinaire des Sciences du Numérique, Orsay, France

² STL, CNRS, Université de Lille, Domaine du Pont-de-bois, Villeneuve-d'Ascq cedex, France

Summary

Objectives: To analyze the content of publications within the medical NLP domain in 2020.

Methods: Automatic and manual preselection of publications to be reviewed, and selection of the best NLP papers of the year. Analysis of the important issues.

Results: Three best papers have been selected in 2020. We also propose an analysis of the content of the NLP publications in 2020, all topics included.

Conclusion: The two main issues addressed in 2020 are related to the investigation of COVID-related questions and to the further adaptation and use of transformer models. Besides, the trends from the past years continue, such as diversification of languages processed and use of information from social networks

Keywords

Natural language processing, semi-automatic selection of publication, topics, issues, 2020

Yearb Med Inform 2021;257-63

<http://dx.doi.org/10.1055/s-0041-1726528>

1 Introduction

Natural Language Processing (NLP) aims at providing methods, tools and resources designed in order to mine textual and narrative documents, and to make it possible to access the information they convey [1]. While human languages are complex (as an example, learning a human language requires many years in order to be fluent), the importance of using NLP approaches to mine documents produced by humans has been pointed out since a long time [2]. In this synopsis, we first present the selection process applied this year and then we analyze the content of some publications. More particularly, we will focus on several important issues such as robustness of the methods, reproducibility of the results, as well as the originality of the research questions addressed in 2020.

2 The Selection Process

In order to identify all papers published during the year 2020 in the field of NLP, we queried two databases: MEDLINE¹, specifically dedicated to the biomedical domain, and the Association for Computational Linguistics (ACL) anthology², a database that brings together the major NLP conferences (ACL, International Conference for Computational Linguistics (COLING), Empirical Methods in Natural Language Processing (EMNLP), International Confer-

ence on Language Resources and Evaluation (LREC), Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), etc.) and journals, since some NLP studies concerning the biomedical domain are published in conferences and journals which are not indexed by PubMed.

We applied the basic query we defined last year on MEDLINE:

```
(English[LA]
AND journal article[PT]
AND 2020[DP]
AND (medical OR clinical
OR natural)
AND "language processing")
```

all journal papers published in English in 2020, having abstract, and composed of sequences “clinical language processing” or “medical language processing” or “natural language processing”. As of 2020, January 9th, we collected 739 entries. We applied a similar query on the ACL anthology database and collected 27 entries. In order to process those 766 papers, we automatically scored the papers. Indeed, all the candidate papers are not specifically related to the NLP domain despite the use of one of the three sequences from the query. For instance, they can be related to other sections from the IMIA Yearbook (Public Health and Epidemiology Informatics, Decision Support, Knowledge Representation and Management, etc.) without providing major issues for the NLP section. Hence, we applied three sets of rules we defined in 2018 while identifying best papers in a previous edition, in order to compute global scores for each publication.

¹ <https://pubmed.ncbi.nlm.nih.gov/>

² <https://www.aclweb.org/anthology/>

The first set of rules is based upon the name of the journal (both full name and concepts found in the name):

- the positive score is assigned to the main journals in which the biomedical NLP research is usually published by the NLP community (Biomedical informatics insights, International Journal of Medical Informatics, Journal of the American Medical Informatics Association, Journal of Biomedical Informatics, BMC Bioinformatics);
- the negative score is assigned to journals not specifically related to NLP, but to other domains such as Cognitive studies and Communication disorders (Neuroscience, Human brain mapping, Operative neurosurgery, Speech therapy, ...etc.). We also dismiss survey papers and papers published in the IMIA Yearbook. We manually defined this set of journals in order to rule out those false positives.

The second set of rules relies on concepts found in both the title and abstract of papers:

- the positive score is assigned to concepts typically involved in papers related to NLP. Those concepts may be related to objectives, resources, and tools (such as *natural language processing, NLP, named entity recognition, NER, part of speech, POS, tagged words, semantic, syntax, biomedical entity, meanings, electronic health record, EHR, reports, notes, clinical text, text corpus, free text, unstructured text, tweets, PubMed, Social Media, MedDRA, UMLS, annotated data, Metamap*);
- the negative score is assigned to concepts that are usually involved in studies related to disorders involving anatomical parts or language abilities (such as *word processing, language production, language comprehension, voice quality, left posterior superior temporal gyrus, pSTG, posterior superior temporal sulcus, pSTS, inferior fronto-occipital fasciculus, IFOF, dorso-*

lateral prefrontal cortex, cortex, language lateralization, chemical fragment, fragment chemistry, brain structures, verbal intelligence, cerebral, positive mismatch responses, pMMRs, prelingual, postlingual, cochlear, aphasia, SAPS, cortical, language function, infants).

The third set of rules is also applied on titles and abstracts, and covers the concepts describing the methodology used in papers:

- the positive score is assigned to papers using classical NLP methods or evaluation metrics (such as *annotation tool, text-mining, rule-based, regular expression, lexicon, CRF, recall, precision, F1-score, F-measure, accuracy, Inter-annotator agreement, Kappa, classify/classifier, detect, extract, extraction, predict, predicting, text simplification, lexical simplification*);
- the negative score is assigned to papers claiming to use the NLP methods, such as pointed out by sequences like *using*

Fig. 1 Distribution of papers according to the filter scores (violet bars indicate the total number of papers, pink bars indicate papers kept in the top-15 best papers).

natural language processing, using NLP, perform a Natural Language Processing analysis. Such papers are downgraded because the NLP claims are usually limited to the use of existing and ready-to-use NLP tools while the main contribution of papers is related to the analysis of tool results rather than to the improvements made to NLP methods by researchers who take advantage of existing tools.

For each of the 766 candidate papers, the final score ranked from 0.05 to 1 (Figure 1). On this figure, the violet bars indicate the total number of papers for each computed grade, while the pink bars indicate the papers we kept in the top-15 best papers' list. This score has been used as a meta element during the manual selection of the top-15 papers. Indeed, the section editors did not fully rely on the scores but only used them as additional information. Hence, both section editors independently browsed the abstracts, keywords and automatic scores, and assigned a *Yes / Maybe / No* score to each paper. All papers having at least one *Yes* or *Maybe* score have been kept for the next step of the selection. At this stage, 143 candidate papers remain (i.e., a subset of 18.4% of the whole dataset). We then performed an adjudication process, in order to choose the final 15 candidates to be proofread by external reviewers. We paid attention to the topics addressed by the researchers and to their geographic origin so as to provide enough diversity. As a result, out of the fifteen papers, four come from the USA, three from China, and one from each of the following countries: Finland, Germany, Italy, Peru, Spain, Turkey, and United-Kingdom. This is the first time we select a paper from South America in our top-15 best paper candidates.

In the next sections, we present the main issues and approaches addressed in the pre-selected publications.

3 Current Trends in Biomedical NLP

Since a few years, we observed the increasing use of transformer models based on word embeddings. Those models are useful to capture the context of words. When they

achieve to cover adequately the domain (e.g., clinical domain) or the properties of a given corpus (e.g., clinical texts), they allow to obtain high performances w.r.t. approaches that do not use such models. Several transformer models have been released during the last years: currently, the most famous one is BERT, a multilingual generic model provided by Google [3].

Based on this model, several other models were produced: (i) either specific to a domain such as BioBERT [4] for the biomedical domain; or (ii) specific to a language, such as BERTje [5] for Dutch, FlauBERT [6] and CamemBERT [7] for French, GottBERT [8] for German, ALBERTo [9] and UmBERTo³ for Italian, or BETO [10] for Spanish; or (iii) specific to a given task such as MRCBERT [11] for summarization or SQuADBert⁴ for question-answering, trained on the Stanford Question-Answering Dataset [12]. Conversely, almost all previous methods still disappeared, due to their relative high lifespan: Word2vec [13], GloVe [14], and ELMo [15]. New approaches are coming, with an increasing number of parameters in those new models, especially the Generative Pretrained Transformer (GPT) series provided by OpenAI: GPT-2⁵ [16] for text generation, and GPT-3 [17] for all tasks of NLP (using up to 175 billion parameters).

As a consequence, there is a harmful race for being the first to produce such resources, which implies to rapidly publish a paper in the first available conference or workshop, or on the arXiv deposit, in order to be cited. Nevertheless, this also implies that such papers would certainly never be identified as best paper candidates for the NLP chapter of the IMIA Yearbook, unless they are indexed in PubMed, or published in an NLP conference that we, co-editors of the NLP section, used to visit.

Nevertheless, an opposite way has also been observed, with authors searching for a green research that does not use models which are not environmentally friendly (expensive in terms of hardware, running

time, and CO₂ footprint). This way has been investigated by Poerner *et al.* [18] which proposed a GreenBioBERT model that has been produced using Word2vec to train a model on a new target domain (namely, on the COVID-19 issue) along with an alignment of vectors from the existing BioBERT model and the model trained with Word2vec.

In terms of architecture, the most commonly used for NLP tasks currently is a bidirectional LSTM, which permits to capture both left and right contexts, plus a final Conditional Random Field (CRF) layer to refine the outputs (BiLSTM-CRF). Those methods achieve excellent results but need high computation resources to train new models, currently only available in huge data centers, which implies an important CO₂ footprint. The main issue is how to use such models on new domains, new languages, or for a new task? Instead of training new models, alternative solutions exist, such as fine-tuning of existing models, domain adaptation, transfer learning, as done by Jin *et al.* [19] to predict clinical trial results. Nevertheless, such models remain expensive.

3.1 Languages Addressed

In relation with the languages processed, work with texts written in English represents the largest part of publications. Indeed, there is a significant number of existing corpora, datasets and resources available in English. Yet, we observe an increasing number of publications dedicated to other languages and a greater variety of languages: Arabic [20], Chinese [21-26], Croatian [27], Finnish [28,29], French [30, 31], German [32-34], Hebrew [35], Italian [36-38], Japanese [39, 40], Korean [41, 42], Norwegian [43], Portuguese [44], Spanish [45-48], Swedish [49], and Turkish [28]. Overall, we believe that the trend observed in previous years is continuing. We expect it will develop further in the future.

3.2 COVID-19

A lot of work has been done on the COVID-19 issue. In order to help such work, several papers present COVID-19 resources, such as the COVID-19 Open Research Dataset

³ <https://github.com/musixmatchresearch/umberto>

⁴ <https://huggingface.co/bert-large-uncased-whole-word-masking-finetuned-squad>

⁵ <https://openai.com/blog/better-language-models/>

(CORD-19) produced by Wang *et al.* [50] and available through the Kaggle⁶ platform. Besides, several works are anchored in hospitals and propose to predict conditions and events related to COVID, such as prediction of admission of patients with COVID in an intensive care unit [51], detection of COVID cases from radiological text reports [46], prognosis prediction for patients with chronic obstructive pulmonary disease [45] or for patients with hypertension [31].

In addition to clinical and hospital information, researchers investigate scientific literature looking for instance for drug repurposing recommendations [52], and for temporal evolution of research work on COVID-19 [53]. Besides, the first systematic reviews related to COVID are proposed [54], including the focus on research needs [55].

The researches also investigate social networks, focusing on analyzing public opinion and emotions on the COVID pandemics in Twitter posts [56, 57], monitoring illicit sales of COVID medication on Twitter and Instagram [58], and observing COVID symptoms and disease histories collected from a large population in Reddit, which may provide more reliable insights [59]. Another important issue is that, in the current situation, the surveillance of emerging epidemiological events becomes again very important, as around 60% of all outbreaks are detected using informal sources, which motivated online epidemiological surveillance [32].

3.3 Neurological and Psychiatric Disorders

We observed an interest for neurological and psychiatric disorders, which are mainly issued from clinical context: detection of duration of untreated psychosis [38], analysis of language in patients with aphasia [60], Alzheimer's disease [61], and autism spectrum disorder [62], generation of artificial mental health records and their evaluation [63], detection and prediction of suicide in mental illness [64–66], automatic detection of agitation and related symptoms among hospitalized patients [67], analysis

of COVID impact on people with epilepsy [36], prediction of care cost in mental health setting [68], and, more generally, the use of artificial intelligence in mental health and its biases [69].

3.4 Place of Patients

The place of patients in the healthcare context is increasing and several publications place patients at the center of their investigations. Hence, in addition to the patient-centered healthcare process, we can also mention work focusing on patient outcome [70], studying public opinion on use of free-text data in electronic medical records for research [71], studying patient feedback on the quality of care [72], analyzing the ways the patients describe their pain [73], measuring the quality of patient-doctor communication [74], analyzing the developmental crisis episodes that occur during early adulthood in social media [75], and analyzing patient experience in order to define some guidelines [76].

3.5 Social Networks

As noticed above, the social networks continue to provide important information on several issues: public opinion and emotions on the COVID pandemics in Twitter posts [56, 57], surveillance of illicit sales of COVID medication on Twitter and Instagram [58], observation of COVID symptoms and disease histories collected from a large population in Reddit [59], surveillance

of emerging epidemiological events [32], analysis of HIV-related tweets and of their relation to the HIV incidence [77], analysis of drug use on Twitter [78], analysis of developmental crisis episodes during early adulthood in social media [75].

We assume that the investigation of social networks will go further in the future, as these networks provide independent and massive information on various events.

4 Conclusion

The NLP publications in 2020 have been heavily marked by the sanitary situation, which motivated an increasing number of works related to the COVID-19 pandemic. The authors addressed all types of content (clinical texts, clinical trials, scientific papers, social media, etc.) so as to mine information related to this major issue (adverse drug reactions, usefulness of existing treatments, psychological impact of the pandemic, ...etc.).

From a scientific perspective, we observe an increasing use of transformer models based on word embeddings. In continuation of the trend already observed in previous years, we notice that the variety of languages processed is also increasing. This observation is also related to the use of multilingual transformer models, among them BERT is the most used since it allows to process more than one hundred languages. In addition, several authors adapted those multilingual models to their data (specific domain for

Table 1 Best paper selection of articles for the IMIA Yearbook of Medical Informatics 2021 in the section 'Natural Language Processing'. The articles are listed in alphabetical order of the first author's surname.

Section

Natural Language Processing

- Iye J, Viani N, Kam J, Yin L, Verma S, Puntis S, Cardinal R, Roberts A, Stewart R, Velupillai S. Generation and evaluation of artificial mental health records for Natural Language Processing. *NPJ Digit Med* 2020;3:69.
- Jin Q, Tan C, Chen M, Liu X, Huang S. Predicting clinical trial results by implicit evidence integration. *Proc of Empirical Methods in NLP*; 2020.
- Poerner N, Waltinger U, Schütze H. Inexpensive domain adaptation of pre-trained language models: case studies on biomedical NER and Covid-19 QA. *Proc of Empirical Methods in NLP*; 2020.

⁶ <https://www.kaggle.com/allen-institute-for-ai/CORD-19-research-challenge>

a given language), which also increases the number of publications related to new languages. An opposite way also appeared with papers focusing on green research, especially to propose methods for processing new domains or new data, using existing transformer models without any new training steps for these complex models.

In the coming years, we hope that environmentally friendly solutions will be preferred to the production of new transformer models which still need more and more computing resources. In addition, we also urge the NLP community to go back to a qualitative analysis of their outputs, rather than a basic and harmful race to present numerical gains over other similar studies. The positive impact of NLP research on clinical issues should be highlighted rather than the search for better results computed by evaluation metrics.

References

- Nadkarni PM, Ohno-Machado L, Chapman WW. Natural Language Processing: an introduction. *J Am Med Inform Assoc* 2011;18:544–51.
- Friedman C, Hripesak G. Natural language processing and its future in medicine. *Acad Med* 1999 Aug;74(8):890–5.
- Devlin J, Chang M-W, Lee K, Toutanova K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding [Internet]. 2019. Available from: <https://arxiv.org/abs/1810.04805v2>
- Lee J, Yoon W, Kim S, Kim D, Kim S, So CH, et al. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* 2020;36(4): 1234–40.
- de Vries W, Van Cranenburgh A, Bisazza A, Caselli T, van Noord G, Nissim M. BERTje: A Dutch BERT Model [Internet]. 2019. Available from: <https://arxiv.org/abs/1912.09582v1>
- Le H, Vial L, Frej J, Segonne V, Coavoux M, Lecouteux B, et al. FlauBERT: Unsupervised Language Model Pre-training for French. In: *Proc of LREC*. Marseille, France: European Language Resources Association; 2020. p. 2479–90.
- Martin L, Muller B, Ortiz Suárez PJ, Dupont Y, Romary L, De La Cergerie É, et al. CamemBERT: a tasty French language model. In: *Proc. of ACL. Association for Computational Linguistics*; 2020. p. 7203–19.
- Scheible R, Thomczyk F, Tippmann P, Jaravine V, Boeker M (2020). GottBERT: a pure German Language Model [Internet]. 2020. Available from: [arXiv preprint arXiv:2012.02110](https://arxiv.org/abs/2012.02110)
- Polignano M, Basile P, de Gemmis M, Semeraro G & Basile V. AIBERTO: Italian BERT Language Understanding Model for NLP Challenging Tasks Based on Tweets. In: *Proc of CLiC-it*, volume 2481. CEUR; 2019.
- Canete J, Chaperon G, Fuentes R, Ho J-H, Kang H, Pérez J. Spanish Pre-Trained BERT Model and Evaluation Data. In *PML4DC at ICLR*; 2020.
- Jain S, Tang G, Chi LS. MRCBert: A Machine Reading Comprehension Approach for Unsupervised Summarization [Internet]. 2021. Available from: <https://arxiv.org/abs/2105.00239v1>
- Rajpurkar P, Zhang J, Lopyrev K, Liang P. SQuAD: 100,000+ questions for machine comprehension of text. In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas: Association for Computational Linguistics; 2016. p. 2383–92.
- Mikolov T, Chen K, Corrado G & Dean J. Efficient Estimation of Word Representations in Vector Space [Internet]. 2013. Available from: <https://arxiv.org/abs/1301.3781v3>
- Pennington J, Socher R, Manning CD. GloVe: Global Vectors for Word Representation. In: *Empirical Methods in Natural Language Processing (EMNLP)*; 2014. p. 1532–43.
- Peters ME, Neumann M, Iyyer M, Gardner M, Clark C, Lee K, Zettlemoyer L. Deep contextualized word representations. In: *Proc. of NAACL*; 2018.
- Radford A, Wu J, Child R, Luan D, Amodei D, Sutskever I. Language Models are Unsupervised Multitask Learners. 2019. Available from: <https://d4mucfpksyww.cloudfront.net/better-language-models/language-models.pdf>.
- Brown TB, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P, et al. D. Language Models are Few-Shot Learners [Internet]. 2020. Available from: <https://arxiv.org/abs/2005.14165v4>
- Poerner N, Waltinger U, Shütze H. Inexpensive Domain Adaptation of Pretrained Language Models: Case Studies on Biomedical NER and Covid-19 QA. In: *Proc. of EMNLP*; 2020. p. 1482–90.
- Jin Q, Tan C, Chen M, Liu X & Huang S. Predicting Clinical Trial Results by Implicit Evidence Integration. In: *Proc. of EMNLP*; 2020. p. 1461–77.
- Faris H, Habib M, Faris M, Alomari M & Alomari A. Medical speciality classification system based on binary particle swarms and ensemble of one vs. rest support vector machines. *J Biomed Inform* 2020;109:103525.
- Chen C-H, Hsieh J-G, Cheng S-L, Lin Y-L, Lin P-H, Jeng J-H. Early short-term prediction of emergency department length of stay using natural language processing for low-acuity outpatients. *Am J Emerg Med* 2020;38(11):2368–73.
- Li X, Lin X, Ren H, Guo J. Ontological organization and bioinformatic analysis of adverse drug reactions from package inserts: Development and usability study. *J Med Internet Res* 2020;22(7):e20443.
- Wang Z, Huang H, Cui L, Chen J, An J, Duan H, et al. Using natural language processing techniques to provide personalized educational materials for chronic disease patients in China: Development and assessment of a knowledge-based health recommender system. *JMIR Med Inform* 2020;8(4):e17642.
- Wu C, Luo G, Guo C, Ren Y, Zheng A, Yang C. An attention-based multi-task model for named entity recognition and intent analysis of Chinese online medical questions. *J Biomed Inform* 2020;108:103511.
- Xia H, An W, Li J, Zhang ZJ. Outlier knowledge management for extreme public health events: Understanding public opinions about COVID-19 based on microblog data. *Socioecon Plann Sci* 2020;10(09):41–1.
- Zhang Z, Zhu L, Yu P. Multi-level representation learning for Chinese medical entity recognition: Model development and validation. *JMIR Med Inform* 2020;8(5):e17637.
- Krsnik I, Glavas G, Krsnik M, Miletic D, Štajduhar I. Automatic annotation of narrative radiology reports. *Diagnostics (Basel)* 2020;10(4):196.
- Güngör O, Güngör T, Uskudarli S. Exseqreg: Explaining sequence-based nlp tasks with regions with a case study using morphological features for named entity recognition. *PLoS One* 2020;15(12):e0244179.
- Moen H, Hakala K, Peltonen L-M, Matinlinna H-M, Suhonen H, Terho K, et al. Assisting nurses in care documentation: from automated sentence classification to coherent document structures with subject headings. *J Biomed Semantics* 2020;11(1):10.
- Grabar N, Dalloux C, Claveau V. CAS: corpus of clinical cases in French. *J Biomed Semantics* 2020;11(1):7.
- Neuraz A, Lerner I, Digan W, Paris N, Tsopra R, Rogier A, et al. Natural language processing for rapid response to emergent diseases: Case study of calcium channel blockers and hypertension in the COVID-19 pandemic. *J Med Internet Res* 2020;22(8):e20773.
- Abbood A, Ullrich A, Busche R, Ghazzi S. EventEpi-A natural language processing framework for event-based surveillance. *PLoS Comput Biol* 2020;16(11):e1008277.
- Ferrario A, Demiray B, Yordanova K, Luo M, Martin M. Social reminiscence in older adults' everyday conversations: Automated detection using natural language processing and machine learning. *J Med Internet Res* 2020;22(9):e19133.
- Wulff A, Mast M, Hassler M, Montag S, Marschollek M, Jack T. Designing an openEHR-based pipeline for extracting and standardizing unstructured clinical data using natural language processing. *Methods Inf Med* 2020;59(S02):e64–e78.
- Barash Y, Guralnik G, Tau N, Soffer S, Levy T, Shimon O, et al. Comparison of deep learning models for natural language processing-based classification of non-english head ct reports. *Neuroradiology* 2020;62(10):1247–56.
- Lanzone J, Cenci C, Tombini M, Ricci L, Tufo T, Piccoli M, et al. Glimpsing the impact of COVID19 lock-down on people with epilepsy: A text mining approach. *Front Neurol* 2020 ;11 :870.
- Mensa E, Colla D, Dalmaso M, Giustini M, Mamo C, Pitidis A, et al. Violence detection explanation via semantic roles embeddings. *BMC Med Inform Decis Mak* 2020;20(1):263.
- Viani N, Kam J, Yin L, Bittar A, Dutta R, Patel R, et al. Temporal information extraction from mental health records to identify duration of untreated psychosis. *J Biomed Semantics* 2020;11(1):2.
- Nakatani H, Nakao M, Uchiyama H, Toyoshiba H, Ochiai C. Predicting inpatient falls using natural language processing of nursing records obtained from Japanese electronic medical

- records: Case-control study. *JMIR Med Inform* 2020;8(4):e16970.
40. Ujice S, Yada S, Wakamiya S, Aaramaki E. Identification of adverse drug event-related Japanese articles: Natural language processing analysis. *JMIR Med Inform* 2020;8(11):e22661.
 41. Cho I, Lee M & Kim Y. What are the main patient safety concerns of healthcare stakeholders: a mixed-method study of web-based text. *Int J Med Inform* 2020;140:104162.
 42. Lee KH, Kim HJ, Kim YJ, Kim JH, Song EY. Extracting structured genotype information from free-text hla reports using a rule-based approach. *J Korean Med Sci* 2020;35(12):e78.
 43. Eskildsen NK, Eriksson R, Christensen SB, Aghasipour TS, Bygso MJ, Brunak S, et al. Implementation and comparison of two text mining methods with a standard pharmacovigilance method for signal detection of medication errors. *BMC Med Inform Decis Mak* 2020;20(1):94.
 44. Lopes F, Teixeira C & Oliveira HG. Comparing different methods for named entity recognition in Portuguese neurology text. *J Med Syst* 2020;44(4):77.
 45. Graziani D, Soriano JB, Rio-Bermudez CD, Morena D, Diaz T, Castillo M, et al. Characteristics and prognosis of COVID-19 in patients with COPD. *J Clin Med* 2020;9(10):3259.
 46. Lopez-Úbeda P, Diaz-Galiano MC, Martín-Noguerol T, Luna A, Urena-Lopez LA, Martín-Valdivia MT. COVID-19 detection in radiological text reports integrating entity recognition. *Comput Biol Med* 2020;127:104066.
 47. Najafabadipour M, Zanin M, Rodriguez-González A, Torrente M, Garcia BN, Bermudez JLC, et al. Reconstructing the patient's natural history from electronic health records. *Artif Intell Med* 2020;105:101860.
 48. Santiso S, Pérez A, Casillas A, Oronoz M. Neural negated entity recognition in Spanish electronic health records. *J Biomed Inform* 2020;105:103419.
 49. Caccamisi A., Jorgensen L, Dalianis H, Rosenlund M. Natural language processing and machine learning to enable automatic extraction and classification of patients' smoking status from electronic medical records. *Ups J Med Sci* 2020;125(4):316–24.
 50. Wang LL, Lo K, Chaddrasekhar Y, Reas R, Yang J, Burdick D, et al. CORD-19: The COVID-19 Open Research Dataset [Internet]. 2020. Available from: <https://arxiv.org/abs/2004.10706v4>.
 51. Izquierdo JL, Ancochea J, Savana COVID-19 Research Group; Soriano JB. Clinical characteristics and prognostic factors for intensive care unit admission of patients with COVID-19: Retrospective study using machine learning and natural language processing. *J Med Internet Res* 2020;22(10):e21801.
 52. Gates LE, Hamed AA. The anatomy of the SARS-CoV-2 biomedical literature: Introducing the CovidX network algorithm for drug repurposing recommendation. *J Med Internet Res* 2020;22(8):e21169.
 53. Ebadi A, Xi P, Tremblay S, Spencer B, Pall R, Wong A. Understanding the temporal evolution of COVID-19 research through machine learning and natural language processing. *Scientometrics* 2020;11:1–15.
 54. Wang LL, Lo K. Text mining approaches for dealing with the rapidly expanding literature on COVID-19. *Brief Bioinform* 2021;22(2):781–99.
 55. Doanvo A, Qian X, Ramjee D, Piontkivska H, Desai A, Majumder M. Machine learning maps research needs in COVID-19 literature. *Patterns (NY)* 2020;1(9):100123.
 56. Boon-Itt S, Skunkan Y. Public perception of the COVID-19 pandemic on Twitter: Sentiment analysis and topic modeling study. *JMIR Public Health Surveill* 2020;6(4):e21978.
 57. Dyer J, Kolic B. Public risk perception and emotion on Twitter during the Covid-19 pandemic. *Appl Netw Sci* 2020;5(1):99.
 58. Mackey TK, Li J, Purushothaman V, Nali M, Shah N, Bardier C, et al. Big data, natural language processing, and deep learning to detect and characterize illicit COVID-19 product sales: Inveillance study on Twitter and Instagram. *JMIR Public Health Surveill* 2020;6(3):e20794.
 59. Picone M, Inoue S, Defelice C, Naujokas MF, Sinrod J, Cruz VA, et al. Social listening as a rapid approach to collecting and analyzing COVID-19 symptoms and disease natural histories reported by large numbers of individuals. *Popul Health Manag* 2020;23(5):350–60.
 60. Themistocleous C, Webster K, Afthinos A, Tsapikini K. Part of speech production in patients with primary progressive aphasia: An analysis based on natural language processing. *Am J Speech Lang Pathol* 2021;30(15):466–80.
 61. Reeves S, Williams V, Costela FM, Palumbo R, Umoren O, Christopher MM, et al. Narrative video scene description task discriminates between levels of cognitive impairment in Alzheimer's disease. *Neuropsychology* 2020;34(4):437–46.
 62. Chojnicka I, Wawer A. Social language in autism spectrum disorder: A computational analysis of sentiment and linguistic abstraction. *PLoS One* 2020;15(3):e229985.
 63. Ive J, Viani N, Kam J, Yin L, Verma S, Puntis S, et al. Generation and evaluation of artificial mental health records for natural language processing. *NPJ Digit Med* 2020;3:69.
 64. Senior M, Burghart M, Yu R, Kormilitzin A, Liu Q, Vaci N, et al. Identifying predictors of suicide in severe mental illness: A feasibility study of a clinical prediction rule (oxford mental illness and suicide tool or OxMIS). *Front Psychiatry* 2020;11:268.
 65. Levis M, Westgate CL, Gui J, Watts BV, Shiner B. Natural language processing of clinical mental health notes may add predictive value to existing suicide risk models. *Psychol Med* 2020;1–10.
 66. Jayasinghe L, Bittar A, Dutta R & Stewart R. Clinician-recalled quoted speech in electronic health records and risk of suicide attempt: a case-cross-over study. *BMJ Open* 2020;10(4):e36186.
 67. Hart KL, Pellegrini AM, Forester BP, Berretta S, Murphy SN, Perlis RH, et al. Distribution of agitation and related symptoms among hospitalized patients using a scalable natural language processing method. *Gen Hosp Psychiatry* 2021;68:46–51.
 68. Colling C, Khondoker M, Patel R, Fok M, Harland R, Broadbent M, et al. Predicting high-cost care in a mental health setting. *BJPsych Open* 2020;6(1):e10.
 69. Straw I, Callison-Burch C. Artificial intelligence in mental health and the biases of language based models. *PLoS One* 2020;15(12):e0240376.
 70. Hernandez-Boussard I, Blayney DW, Brooks JD. Leveraging digital data to inform and improve quality cancer care. *Cancer Epidemiol Biomarkers Prev* 2020;29(4):816–22.
 71. Ford E, Oswald M, Hassan L, Bozentko K, Nenadic G, Cassell J. Should free-text data in electronic medical records be shared for research? A citizens' jury study in the UK. *J Med Ethics* 2020;46(6):367–77.
 72. Nawab K, Ramsey G, Schreiber R. Natural language processing to extract meaningful information from patient experience feedback. *Appl Clin Inform* 2020;11(2):242–52.
 73. Tighe PJ, Sannapaneni B, Fillingim RB, Doyle C, Kent M, Shickel B, et al. Forty-two million ways to describe pain: Topic modeling of 200,000 PubMed pain-related abstracts using natural language processing and deep learning-based text generation. *Pain Med* 2020;21(11):3133–60.
 74. Cuffy C, Hagiwara N, Vrana S, McInnes BT. Measuring the quality of patient-physician communication. *J Biomed Inform* 2020;112:103589.
 75. Agarwal S, Guntuku SC, Robinson OC, Dunn A, Ungar LH. Examining the phenomenon of quarter-life crisis through artificial intelligence and the language of twitter. *Front Psychol* 2020;11:341.
 76. Cammel SA, De Vos MS, van Soest D, Hettne KM, Boer F, Steyerberg EW, et al. How to automatically turn patient experience free-text responses into actionable insights: a natural language programming (NLP) approach. *BMC Med Inform Decis Mak* 2020;20(1):97.
 77. Stevens R, Bonett S, Bannon J, Chittamuru D, Slaff B, Browne SK, et al. (2020). Association between HIV-related tweets and HIV incidence in the United States: Infodemiology study. *J Med Internet Res* 2020;22(6):e17196.
 78. Tassone J, Yan P, Simpson M, Mendhe C, Mago V, Choudhury S. Utilizing deep learning and graph mining to identify drug use on Twitter data. *BMC Med Inform Decis Mak* 2020;20(Suppl 11):304.

Correspondence to:

Natalia Grabar

E-Mail: natalia.grabar@univ-lille.fr

Cyril Grouin

E-Mail: cyril.grouin@limsi.fr

Content Summaries of Best Papers for the Natural Language Processing Section of the 2021 IMIA Yearbook

Jin Q, Tan C, Chen M, Liu X, Huang S

Predicting clinical trial results by implicit evidence integration

Proc of Empirical Methods in NLP; 2020

The clinical trial result prediction (CTRP) task is based on medical literature containing PICO (how the Intervention group compares with the Comparison group in terms of the measured Outcomes in the studied Population). The authors proposed an EBM-Net model which is a transformer model that uses unstructured sentences as implicit evidences and a fine-tuning approach. They compared their fine-tuned model w.r.t. the BioBERT model and other approaches (MeSH ontology, bag-of-words, ...etc.) and achieved better results on the COVID-19 clinical trials' dataset (22 clinical trials from the CORD-19 dataset).

Poerner N, Waltinger U, Schütze H

Inexpensive domain adaptation of pre-trained language models: case studies on biomedical NER and Covid-19 QA

Proc of Empirical Methods in NLP; 2020

The authors highlight the expensive cost of domain adaptation while training a model on target-domain text. This cost is expressed in terms of hardware requirement, high running time and negative impact on the CO2 footprint. The authors investigated a solution they called GreenBioBERT that relies first on a Word2vec training stage on target-domain texts (PubMed, PMC, CORD-19), and second on the alignment of word vectors with the vectors from BioBERT. They applied their model on two issues: 8 biomedical NER tasks in English, and question-answering (QA) on COVID-19 issue. The authors achieved competitive results using BioBERT and a better precision on a few tasks; on the COVID-19 QA task, their model achieved better results than the SQuADBERT model (designed for QA). In this paper, the authors proposed a useful method to use existing pretrained language models in order to adapt them to new datasets, new tasks, new languages, etc.

Ive J, Viani N, Kam J, Yin L, Verma S, Puntis S, Cardinal R, Roberts A, Stewart R, Velupillai S

Generation and evaluation of artificial mental health records for Natural Language Processing

NPJ Digit Med 2020;3:69

The main problem for biomedical NLP is the difficult access to clinical documents and the inherent complexity to completely de-identify documents. The solution proposed by the authors consists in generating artificial discharge summaries. In this paper, the authors produced artificial summaries in mental health, based on the MIMIC-III data. Then, they used their artificial texts in order to train models using the Keras toolkit for classification tasks. The authors observed that models trained on their synthetic data perform as well as models trained on real data. Since the synthetic discharge summaries have been produced taking as input the MIMIC-III data, the authors cannot share their resources. Nevertheless, the method is reproducible.