

Automatic Classification and Entity Relation Detection in Hungarian Spinal MRI Reports

1st András Kicsi

*Department of Software Engineering
University of Szeged
Szeged, Hungary
akicsi@inf.u-szeged.hu*

2nd Klaudia Szabó Ledenyi

*Department of Software Engineering
University of Szeged
Szeged, Hungary
ledenyik@inf.u-szeged.hu*

3rd Péter Pusztai

*Department of Software Engineering
and MTA-SZTE Research Group
on Artificial Intelligence
University of Szeged
Szeged, Hungary
pusztaip@inf.u-szeged.hu*

4th László Vidács

*Department of Software Engineering
and MTA-SZTE Research Group
on Artificial Intelligence
University of Szeged
Szeged, Hungary
lac@inf.u-szeged.hu*

Abstract—A great number of radiologic reports are created each year which incorporate the expertise of radiologists. This knowledge could be exploited via machine understanding. This could provide valuable statistics and visualization of the reports, and as training data, and it could also contribute to later automatic reporting applications. In our current work, we present our first steps toward the machine understanding of clinical reports of the spinal region, written in the Hungarian language. Our system provides an automatic classification and connection detection for various entities in the text. Our classification is achieved via bi-directional long short-term memory and conditional random fields producing 0.87-0.95 F1-score values, while the extraction of connection relies on linguistic analysis and predefined rules. The extracted information is displayed in an easily comprehensible, well-formed tree-structure.

Index Terms—nlp, radiology, machine understanding, visualization, bi-lstm-crf, clinical reports

I. MOTIVATION

Even though medical examinations usually result in graphic data, the professionals' actual medical expertise often reflects only in the reports they composed viewing these images. Machine understanding of this information

The research presented in this paper, carried out by the University of Szeged was supported by the Ministry of Innovation and the National Research, Development and Innovation Office within the framework of the Artificial Intelligence National Laboratory Programme. This research was supported by grant NKFIH-1279-2/2020 of the Ministry for Innovation and Technology, Hungary. Support was provided by the European Social Fund (EFOP-3.6.3-VEKOP-16-2017-00002), by the Hungarian Government. András Kicsi was supported by the ÚNKP-20-3 - New National Excellence Program of the Ministry for Innovation and Technology from the source of the National Research, Development and Innovation Fund.

could contribute to several aspects of the process and even enable the automatization of the final diagnosis.

The reports are customarily created in a free-form text which is easy to compose and comprehend by human experts but represents a substantial obstacle for artificial intelligence methods. Furthermore, the text is most often composed in the native language of the author, which presents additional complications.

This is true for our current scope, spinal magnetic resonance imaging (MRI) examinations. Radiologists compose reports routinely, a single specialist can create thousands of these annually. In Hungary, the reports follow only a very loose structure, and they are usually written in the Hungarian language. Radiologists are encouraged to declare not just the presence but also the absence of negative conditions. The radiologist receives the MRI images and the patient's previous history and composes both the report and a medical opinion that contains similar content but in a more concise and result-oriented manner. This process is illustrated in Fig. 1.




Fig. 1: The work of the radiologist during an examination

Precise machine understanding could contribute to this process considerably. In real-time, it could automatically check the consistency of the report and the opinion, a graphic visualization could provide a quick glance of what the radiologist might have missed, compare the current report with the results of the previous examinations and even warn the radiologist if the meaning of a sentence is unusual. Valuable statistics could be generated, even from archive data, aiding financial administration of hospitals. Further possibilities lie in the standardization of reports that could enable easier subsequent processing and bring the reports from different institutions to a common ground, contributing to patient-oriented health system improvements [1], [2]. While these are potentially beneficial uses, machine understanding of these texts could provide a pathway for automatic report generation, which is an even worthier goal that would open new doors for more precise and objective reports.

Modern deep learning methods can analyze medical images with high precision. The major setback of such an approach is that these methods tend to require immense amounts of data to work correctly. As the creation of such data requires medical expertise, its gathering presents vast time and resource costs. Traditionally, such data is gathered by radiologists manually annotating tens of thousands of MRI images, precisely pointing out the problems and their locations. Reports are similarly composed by radiologists and represent a very similar, if less graphic source of information. The expert opinion is already contained within. Thus its proper extraction can contribute valid training data even for such an image-based automatic diagnostic tool, saving resources. Our current paper deals with the extraction of such useful information from the text of the reports.

The paper presents our efforts to extract various entities and their connections from Hungarian radiologic reports of the spinal region. This process involves the classification of various anatomical locations, pathologies, and properties via machine learning trained on 487 manually annotated reports. The connections are determined based on language models, and the data is visualized in an easy to comprehend manner.

II. METHOD

The current section describes the methods used in our experiments. 487 anonymized reports have been manually annotated by a radiologist according to our classification system, this serves as our training data. Our artificial intelligence methods rely on linguistic analysis conducted via the Magyarlanc [3] tool, which provides syntactic, morphologic, constituent, and dependency analysis for general Hungarian texts with high accuracy. Our machine learning approach makes use of the various features extracted by Magyarlanc while the connections rely on its sentence parsing and constituent analysis functions.

A. Annotation

Our annotation system incorporates a few simple entities that need to be classified. These are the anatomic locations, disorders, and properties. These three classes tend to cover most of the meaningful words found in typical radiologic reports. An example of our system, converted to English for better understanding, can be seen in Fig. 2. Note that the reports are at no point translated to English in our process. An entity can consist of multiple words. A term was considered an anatomic location if it describes a specific part of the human body such as "L2" or "disc", or even as a part of a disorder itself like "disci" in "hernia disci". These entities are relatively typical and have a smaller vocabulary than the others. Disorders are the various pathologies observed by the radiologist like "hernia" or "dehydration". Positive or neutral statements also belong here such as "intact" or "status idem". The aspects under observation like liquid content or height are also considered parts of our system's disorders, as these specify the disorder. Disorders can be easily confused with properties such as in the case of "deforming" in our example. Properties are usually describing the stage or degree of a disorder, or in some cases, specify its precise location. Some examples include "3 mm", "right", and "significantly". Properties clearly have the largest vocabulary since they are much less reliant on the medical terminology than the other elements.

The annotation itself was conducted in the Brat [4] annotation tool by a radiologist and covered 487 Hungarian reports at the current phase. Brat is a tool that facilitates annotation by an easy-to-use web-based user interface, it is highly configurable, and annotation results can be downloaded in a relatively simple, tab-separated format. Note that it is a platform for the annotation, all the annotations were conducted manually. Our annotation system is the result of several meetings between the radiologist, linguists and computer scientists. A thorough set of guidelines was available for the radiologist during the annotation.

B. Classification

Our classification model is essentially a named entity recognition (NER) model. It is based on a Bi-LSTM-CRF (bi-directional long short-term memory [5], conditional random fields [6]) architecture similar to the one published by Ma et al. [7]. Since Hungarian is a morphologically rich language, character level embedding was also utilized in our solution, as suggested by Ling et al. [8]. Apart from word and character embeddings, additional predictive features such as lemmas, part-of-speech (POS) tags and part-of-sentence tags were also utilized in the model. The first layers of our NER model were embedding layers, mostly initialized with random weights. For the textual inputs, the corresponding embedding matrix is initialized with pre-trained word vectors (trained on the Hungarian Wikipedia). In a regular forward-pass, the integer encoded feature sequences were first passed through their

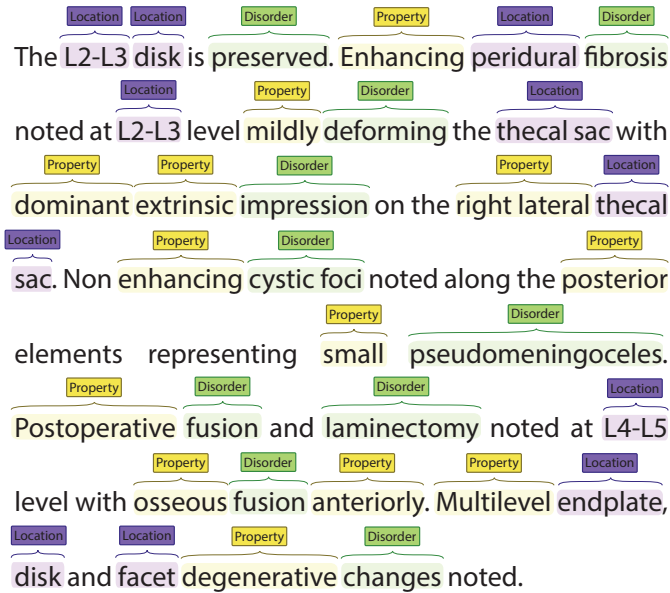


Fig. 2: An English language illustration of our annotation system

corresponding embedding layers where each feature (word, lemma, POS tag, part-of-sentence tag) was mapped to a dense vector representation. The character level vector representations of words were generated by an additional Bi-LSTM network. The main Bi-LSTM layer took the concatenated vector representations of all the features (word, lemma, POS tag, part-of-sentence tag, character) as input. The Bi-LSTM layer's output was passed through a densely-connected feed-forward layer, on top of which a CRF layer performed the final sequence tagging.

The generated embedding vectors for word, lemma, POS tag, part-of-sentence, and character had dimensions of 64, 64, 20, 20, 20, respectively. Sequence lengths to represent sentences and words (for character embedding) were fixed at 75 and 20, respectively. The main Bi-LSTM layer consisted of 100 units. A recurrent dropout of 0.3, an input dropout of 0.3, and an L2 regularization of 0.001 were used for regularization. The Bi-LSTM layer used for character encoding had 20 units. No dropouts were used in this layer. The densely-connected feed-forward layer contained 50 units with ReLU activation. RMSprop, with a learning rate of 0.001 was used as an optimizer. The batch size during training was set to 32. Our model was implemented in Tensorflow, Keras. The training ran on an NVIDIA GTX 1060 graphics card with 6 GB of memory.

The model aims to classify whole terms, not just tokens. Thus, the annotations are converted to inside-outside-beginning tags before training and prediction. Consequently, a location, for example, can consist of a single token tagged with B-Location and any number of trailing I-Location tokens.

C. Connections

The automatic extraction of connections utilizes the constituent parser of the linguistic analyzer and our classification's output. Almost all of the properties can be attributed to a disorder rather than an anatomic location. Many cases like in "compressed L5 disc", where this presumption seems faulty, are the result of flawed classification (as "compressed" should be a disorder itself here). Thus, only two kinds of connections are determined, between disorders and locations and between properties and disorders. Our system also attempts to merge the locations and the disorders that belong together (like in the case of "L2-L3 level mildly deforming the thecal sac" where the proper location would be "L2-L3 thecal sac").

Our method uses several predefined rules to cope with the task of these automatic assignments. Our method relies heavily on sentence parsing and the constituents of the sentences as the entities that belong to the same constituent are extremely likely to be connected semantically. The following rules were constructed:

- **Disorder-Location:** The system first considers only the constituent of a location. The preceding disorders are prioritized first; if more than one exists, they always receive the same treatment. If none exists, the system looks for rightmost ones. If no such disorder is found, the system broadens the search to the whole sentence, but only considers the words on the left of the location. Coordinations are also considered, the locations that are coordinated with "and", "or", or a comma always receive the same connections.
- **Property-Disorder:** If a disorder is preceded by any properties inside its constituent, they gain connections. Otherwise, the whole left side of the sentence is considered.
- **Location-Location:** Some of the locations like "disc" or "endplate" are not entirely specific, they need a vertebra or at least a region to achieve precision. A set of such locations was assembled manually, in their cases, a suitable vertebra or vertebrae are attempted to be found inside their sentence.
- **Disorder-Disorder:** Similarly to the previous problem, aspects like "height" do not convey meaningful disorders but lend specificity to others. Since they are also very typically worded with little variety, in the same manner, a list of these was gathered, and potentially suitable other disorders are found inside their sentence.

We note that the Hungarian sentence structure differs significantly from that in English. Furthermore, the sentences of radiologic reports tend to have relatively simple structures. Thus, such rules are correct in the overwhelming majority of the cases.

The detected entities are then displayed in a structured manner, their connections highlighted by grouping via frames. The result always forms a tree structure. An

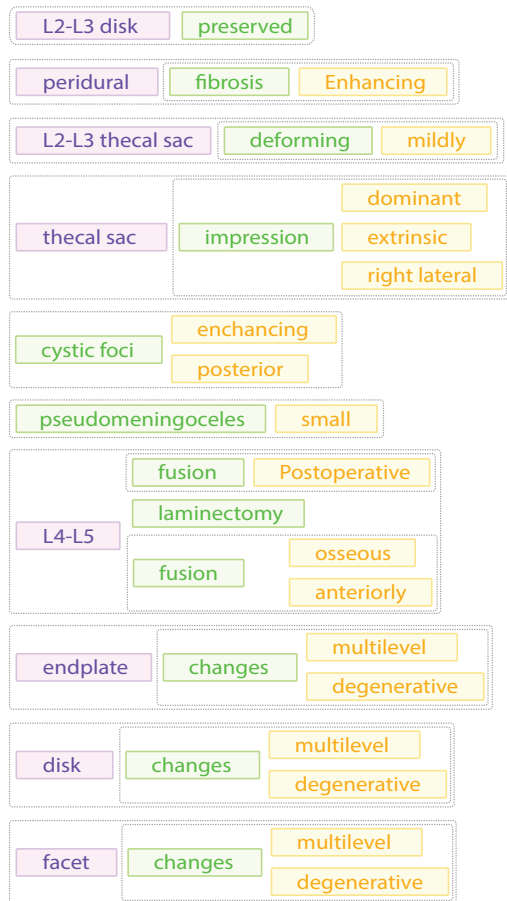


Fig. 3: An illustration of our structured visualization of the text seen in Fig. 2

illustration of our visualization system adapted to English can be seen in Fig. 3. This is a visualization of the report seen previously in Fig. 2. Note that this example was explicitly constructed as an illustration. Our system is not suitable for processing English reports which vary significantly from Hungarian reports in both terminology and sentence structure.

Our entire method is displayed in Fig. 4 with a Hungarian example. The 487 reports were manually annotated by a radiologist, and with their extracted linguistic features, these represented the training data for the Bi-LSTM-CRF model. When the system receives a new report, its text and linguistic features are given to the model for prediction. The model performs the classification of the locations, disorders and properties. The resulting entities are then submitted to our connection extraction, which use sentence and constituent parsing and predefined rules to determine the probable connections. This results in a tree visualization where the different entities are color coded and displayed in a visual format in which the connected entities are grouped together. The process is performed automatically and runs extremely quickly, making it suitable for real-time display of reports during typing. The negations are not handled in our examples, but it can be

noted that this is a relatively simple task as the negations are detected by the linguistic analysis and are almost always denying the presence of a disorder.

III. RESULTS AND DISCUSSION

The backbone of our process is formed by the Bi-LSTM-CRF model that classifies the entities into three categories. The radiologist’s manual annotation noted 7835 disorders, 6409 locations, 3490 properties. On a token level, this corresponds to 12016 disorder, 11662 location, and 5285 property tokens. Since the model works on a token level, it also assigns beginning and inside values for the tokens to mark the complete entities. Our results are displayed in Table I. During the process, 70% of the 487 reports was used for training, 10% for validation and 20% as the test set. The numbers of the table represent the results measured on this 97 reports. Since this is still a relatively small set, we also conducted three tenfold cross-validations with different random seeds, and found that the results did not show a decrease.

TABLE I
OUR CLASSIFICATION RESULTS ON 20% OF THE MANUALLY ANNOTATED DATA

Class	Precision	Recall	F1-score	Support
B-Disorder	0.9105	0.9111	0.9108	1608
I-Disorder	0.8475	0.8589	0.8532	893
B-Location	0.9518	0.9329	0.9422	1311
I-Location	0.9404	0.9568	0.9485	1203
B-Property	0.8712	0.8628	0.8670	729
I-Property	0.8939	0.8551	0.8741	414
Micro Average	0.9101	0.9076	0.9089	6158

It is visible that locations tend to offer the best results while properties seem to lag behind. This can be partly due to the visible difference in sample size, but an even more likely cause is the size of the classes’ vocabulary. Properties can take up a wide variety of forms while locations use a fairly limited set of terms.

In the 487 reports, the model detected 7794 disorders, 6358 locations, and 3442 properties. Our system assigned 11016 connections, of which 6924 were Disorder-Location, 3382 were Property-Disorder, 425 mergings of locations, and 285 mergings of disorders. The radiologist reviewed these connections manually and found the connections precise. Some mistakes were detected where the classification itself was faulty or in cases of rare, exceptionally complex sentences.

As the refinement of the training data could further improve the results, the same reports will be annotated by at least one more radiologist, and conflicting cases will be resolved. The inclusion of typo correction is also likely to contribute to the process, this research is already underway. The next step in our machine understanding process will include the precise, ontology-based identification of the locations and disorders that enables pinpointing the tree elements of our visualization on the spine, and eventually even on the MRI images.

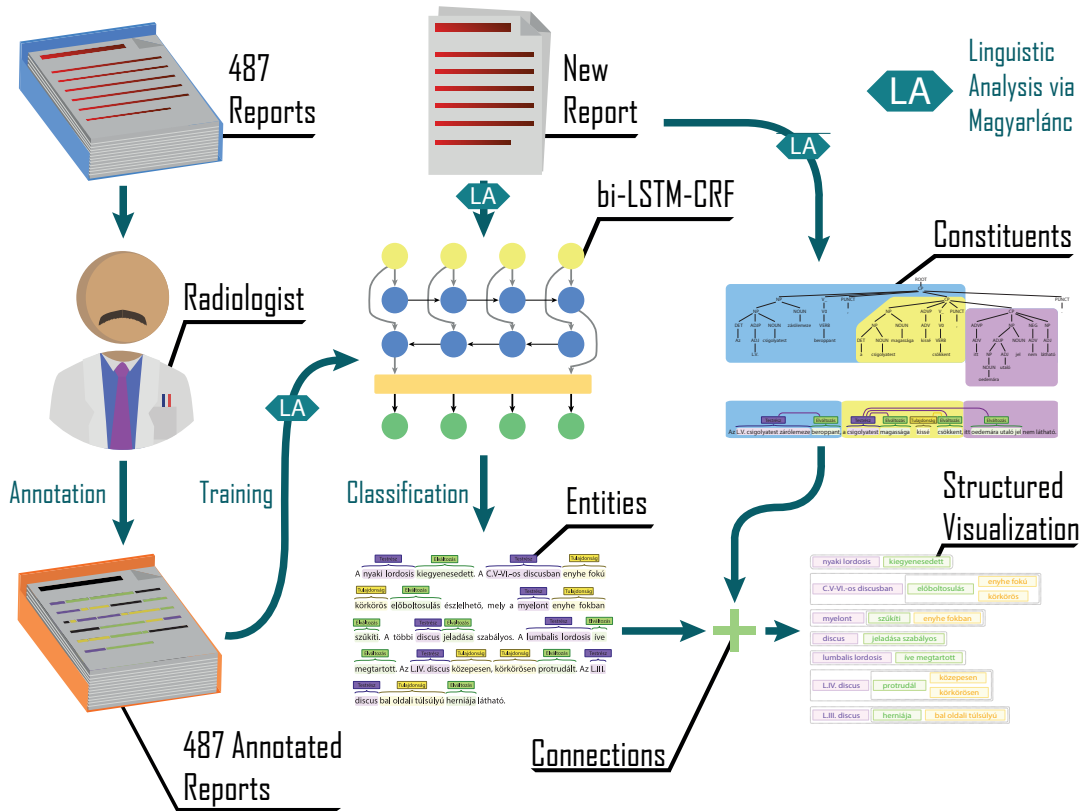


Fig. 4: Our proposed method for automatic visualization

The extension of our process to a new language would require new training data for our detection model, and a new set of possible names would have to be defined for our ontology. If such exists, a similar ontology from the target language could be extended to accommodate this. New linguistic analysis tools with at least morphological and constituent analysis capabilities should be enlisted to handle the new language. New rules would be required for the connections based on manual experiments. This would be a relatively high-effort transition, but besides the obvious benefits, it would also enable a standardized translation of the reports between different languages.

IV. RELATED WORK

Words in sentences follow distinct patterns. Sequence tagging is a task where the class type of an individual element can depend on the class type of neighboring elements. To unravel and exploit these patterns and dependencies recurrent neural networks have proven to be very effective. Bi-LSTM-CRF is a widely used architecture in these types of tasks [9]. In the medical domain, Bi-LSTM-CRF and its derivative architectures are very popular in drug name recognition (DNR), clinical concept extraction (CCE) and adverse drug event recognition (ADER) tasks [10]–[13]. In recent years, these networks have been widely used and further improved for named entity recognition tasks from Chinese medical reports [14], [15]. Yin et al. used features extracted by convolutional

neural networks to enrich the semantic information of the characters and applied a self-attention mechanism to capture the dependencies between characters [16]. Li et al. implemented attention mechanism into their Bi-LSTM-CRF architecture, which enabled their model to capture more useful context information and alleviated the problem of missing information caused by long distances between related elements in the sequence [17]. Cai et al. suggested that named entity recognition on Chinese medical reports can be improved by making entity boundary detection more accurate [18]. They proposed the utilization of POS tags using a BiLSTM-CRF architecture with a self-matching attention layer. Zhao et al. used a lattice LSTM-CRF system with adversarial training [19]. The advantage of lattice LSTM was that it could integrate word- and character-level information, while adversarial training improved the model’s robustness by adding perturbations to the training data. Zhang et al. pointed out that recent NER models use only one layer to encode information in texts [20]. They argued that these shallow text representations could not capture in-depth features and pose a limitation on model performance. They proposed a stacked architecture combining an LSTM and a gated recurrent unit (GRU) layer with a final CRF layer as the classifier. Bi-LSTM-CRF systems have also been successfully applied in information extraction tasks from Portuguese medical texts. Lopes et al. used Bi-LSTM-CRF

neural networks with in-domain and out-of-domain word embeddings [21]. They found that applying in-domain embeddings results in better model performance than using out-of-domain embeddings (even if the in-domain embeddings are trained on a much smaller dataset than the out-of-domain embeddings). In their recent work, they compared the performance of a baseline CRF and different Bi-LSTM-CRF derivatives in medical named entity recognition task [22]. They showed that models trained on public clinical texts could effectively extract information from previously unseen hospital clinical texts.

An accurate and automatic NER model on clinical texts opens the way to many possibilities. Such exciting applications could be automatic opinion generation, smart statistics generation or visual summarization of the medical records. Our goal is to develop a system that can automatically understand free-text radiologic reports and visualize them in real-time. In this study, a hybrid system was used based on Bi-LSTM-CRF that could extract medical named entities effectively, while other components of our system grouped the corresponding entities. A framework for a visualization system was also constructed to visualize entities and their relations in a tree-like structure. In the next iteration of this framework, the tree elements will be connected to the corresponding parts of a schematic spine image, this task needs proper identification of locations which is already well underway. This kind of machine learning based, real-time visualization systems are relatively scarce in the literature [23] and to the best of our knowledge they are entirely non-existent for the Hungarian language. An earlier version of our training solution was published previously in Hungarian [24] as well as various supplementing solutions [25]–[27] that could contribute to the future progress of the process but are not in the scope of the current paper.

V. CONCLUSIONS

The paper described our efforts in creating an automatic framework for the machine understanding of Hungarian spinal region reports. Our current goals were the classification of anatomic locations, disorders and their properties in the free-form text of the reports. This was achieved via a Bi-LSTM-CRF solution trained on 487 manually annotated reports. The classification produced F1-score values between 0.8670 and 0.9485, with a micro-average of 0.9089. The detected entities are connected to each other based on sentence and constituent parsing and pre-defined rules specific to the Hungarian reports' sentence structure. The gathered information is displayed in a tree-structure with the connected components grouped together. The process is suitable for real-time visualization, even while the report is being written.

Our future plans concern the refinement of our current data, proper identification of locations and disorders, and eventually their connection to the original MRI images themselves.

REFERENCES

- [1] J. Rojo, J. Hernandez, and J. M. Murillo, "A personal health trajectory API: Addressing problems in health institution-oriented systems," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 12128 LNCS. Springer, jun 2020, pp. 519–524.
- [2] A. Roehrs, C. A. da Costa, and R. da Rosa Righi, "OmniPHR: A distributed architecture model to integrate personal health records," *Journal of Biomedical Informatics*, vol. 71, pp. 70–81, jul 2017.
- [3] J. Zsibrita, V. Vincze, and R. Farkas, "magyarlanc: A Toolkit for Morphological and Dependency Parsing of Hungarian," p. 763–771, 2013.
- [4] P. Stenetorp, S. Pyysalo, G. Topić, T. Ohta, S. Ananiadou, and J. Tsujii, "brat: A Web-based Tool for NLP-Assisted Text Annotation," in *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*. Avignon, France: Association for Computational Linguistics, April 2012, pp. 102–107.
- [5] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [6] J. D. Lafferty, A. McCallum, and F. C. N. Pereira, "Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data," in *Proceedings of the Eighteenth International Conference on Machine Learning*, ser. ICML '01. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2001, pp. 282–289.
- [7] X. Ma and E. Hovy, "End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF," in *54th Annual Meeting of the Association for Computational Linguistics, ACL 2016 - Long Papers*, vol. 2, 2016, pp. 1064–1074.
- [8] W. Ling, T. Luís, L. Marujo, R. F. Astudillo, S. Amir, C. Dyer, A. W. Black, and I. Trancoso, "Finding function in form: Compositional character models for open vocabulary word representation," in *Conference Proceedings - EMNLP 2015: Conference on Empirical Methods in Natural Language Processing*, 2015, pp. 1520–1530.
- [9] X. Ma and E. Hovy, "End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF," in *54th Annual Meeting of the Association for Computational Linguistics, ACL 2016 - Long Papers*, vol. 2, 2016, pp. 1064–1074.
- [10] Z. Liu, M. Yang, X. Wang, Q. Chen, B. Tang, Z. Wang, and H. Xu, "Entity recognition from clinical texts via recurrent neural network," *BMC Medical Informatics and Decision Making*, vol. 17, no. S2, p. 67, jul 2017.
- [11] I. Jauregi Unanue, E. Zare Borzeshi, and M. Piccardi, "Recurrent neural networks with specialized word embeddings for health-domain named-entity recognition," *Journal of Biomedical Informatics*, vol. 76, pp. 102–109, dec 2017.
- [12] Y. Chen, C. Zhou, T. Li, H. Wu, X. Zhao, K. Ye, and J. Liao, "Named entity recognition from Chinese adverse drug event reports with lexical feature based BiLSTM-CRF and tri-training," *Journal of Biomedical Informatics*, vol. 96, aug 2019.
- [13] Y. Zhang, L. Guo, D. Huang, K. Huang, J. Li, and Z. Pan, "English Drug Name Entity Recognition Method Based on Attention Mechanism BiLSTM-CRF," in *Proceedings of IEEE 14th International Conference on Intelligent Systems and Knowledge Engineering, ISKE 2019*. Institute of Electrical and Electronics Engineers Inc., nov 2019, pp. 831–836.
- [14] B. Ji, R. Liu, W. S. Xu, S. S. Li, J. T. Tang, J. Yu, and Q. Li, "A BiLSTM-CRF method to Chinese electronic medical record named entity recognition," in *ACM International Conference Proceeding Series*. Association for Computing Machinery, dec 2018.
- [15] B. Ji, R. Liu, S. Li, J. Yu, Q. Wu, Y. Tan, and J. Wu, "A hybrid approach for named entity recognition in Chinese electronic medical record," *BMC Medical Informatics and Decision Making*, vol. 19, no. S2, p. 64, apr 2019.
- [16] M. Yin, C. Mou, K. Xiong, and J. Ren, "Chinese clinical named entity recognition with radical-level feature and self-attention mechanism," *Journal of Biomedical Informatics*, vol. 98, oct 2019.

- [17] L. Li, J. Zhao, L. Hou, Y. Zhai, J. Shi, and F. Cui, "An attention-based deep learning model for clinical named entity recognition of Chinese electronic medical records," *BMC Medical Informatics and Decision Making*, vol. 19, no. 5, pp. 1–11, dec 2019.
- [18] X. Cai, S. Dong, and J. Hu, "A deep learning model incorporating part of speech and self-matching attention for named entity recognition of Chinese electronic medical records," *BMC Medical Informatics and Decision Making*, vol. 19, no. S2, apr 2019.
- [19] S. Zhao, Z. Cai, H. Chen, Y. Wang, F. Liu, and A. Liu, "Adversarial training based lattice LSTM for Chinese clinical named entity recognition," *Journal of Biomedical Informatics*, vol. 99, nov 2019.
- [20] R. Zhang, W. Lu, S. Wang, X. Peng, R. Yu, and Y. Gao, "Chinese clinical named entity recognition based on stacked neural network," *Concurrency Computation*, apr 2020.
- [21] F. Lopes, C. Teixeira, and H. Gonalo Oliveira, "Contributions to Clinical Named Entity Recognition in Portuguese," in *18th BioNLP Workshop and Shared Task*, 2019, pp. 223–233.
- [22] F. Lopes, C. Teixeira, and H. Gonalo Oliveira, "Comparing Different Methods for Named Entity Recognition in Portuguese Neurology Text," *Journal of Medical Systems*, vol. 44, no. 4, pp. 1–20, apr 2020.
- [23] W. Ruan, N. Appasani, K. Kim, J. Vincelli, H. Kim, and W. S. Lee, "Pictorial visualization of EMR summary interface and medical information extraction of clinical notes," in *CIVEMSA 2018 - 2018 IEEE International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications, Proceedings*. Institute of Electrical and Electronics Engineers Inc., aug 2018.
- [24] A. Kicsi, P. Pusztai, K. Szabó Ledenyi, E. Szabó, G. Berend, V. Vincze, and L. Vidács, "Információkinyerés magyar nyelvű gerinc mr leletekből," in *XV. Magyar Számítógépes Nyelvészeti Konferencia (MSZNY 2019)*, Szeged, 2019, p. 177–186.
- [25] A. Kicsi, P. Pusztai, E. Szabó, and L. Vidács, "Szaknyelvi annotációk javításának statisztikai alapú támogatása," in *XVI. Magyar Számítógépes Nyelvészeti Konferencia (MSZNY 2020)*, Szeged, 2020, p. 115–128.
- [26] A. Kicsi, K. Szabó Ledenyi, P. Németh, P. Pusztai, L. Vidács, and T. Gyimóthy, "Elírások automatikus detektálása és javítása radiológiai leletek szövegében," in *XVI. Magyar Számítógépes Nyelvészeti Konferencia (MSZNY 2020)*, Szeged, 2020, p. 191–204.
- [27] A. Kicsi, K. Szabó Ledenyi, P. Pusztai, P. Németh, and L. Vidács, "Entitások azonosítása és szemantikai kapcsolatok feltárása radiológiai leletekben," in *XVI. Magyar Számítógépes Nyelvészeti Konferencia (MSZNY 2020)*, Szeged, 2020, p. 15–28.