



Universiteit
Leiden
The Netherlands

Healthcare NLP: where are we and what is next?

Han, L.; Rayson, P.; Moore, A.; Nenadic, G.; Verberne, S.

Citation

Han, L., Rayson, P., Moore, A., Nenadic, G., & Verberne, S. (2026). Healthcare NLP: where are we and what is next? In . Spain. Retrieved from <https://hdl.handle.net/1887/4306627>

Version: Accepted Manuscript

License: [Creative Commons CC BY-SA 4.0 license](#)

Downloaded from: <https://hdl.handle.net/1887/4306627>

Note: To cite this publication please use the final published version (if applicable).

HealthcareNLP: where are we and what's next?

Lifeng Han, Paul Rayson, Suzan Verberne, Andrew Moore, and Goran Nenadic

On behalf of the 4D PICTURE consortium and beyond

Leiden University Medical Centre, NL

Lancaster University, UK

University of Manchester, UK

The Leiden Institute of Advanced Computer Science, NL

{l.han, s.verberne} @liacs.leidenuniv.nl

Accepted Tutorial by LREC 2026

Abstract

This proposed tutorial focuses on Healthcare Domain Applications of NLP, what we have achieved around HealthcareNLP and the challenges that lie ahead for the future. Existing reviews in this domain either overlook some important tasks, such as synthetic data generation for addressing privacy concerns, or explainable clinical NLP for improved integration and implementation, or fail to mention important methodologies, including retrieval augmented generation and the neural symbolic integration of LLMs and KGs. In light of this, the goal of this tutorial is to provide an overview introduction of the most important sub-areas of a patient and resource-oriented HealthcareNLP, with three layers of hierarchy: 1) data/resource layer: annotation guidelines, ethical approvals, governance, synthetic data; 2) NLP-Eval layer: NLP tasks NER, RE, sentiment, and Linking/coding with categorised methods, leading to explainable HealthAI; 3) patients layer: Patient Public Involvement and Engagement (PPIE), health literacy, translation, simplification, and summarisation (NLP tasks too), shared decision making support. A hands-on session will be included in the tutorial for audiences to use HealthcareNLP applications. The target audiences include NLP practitioners in the healthcare application domain, NLP researchers who are interested in domain applications, Healthcare researchers, and students from NLP fields. The type of tutorial is "Introductory to CL/NLP topics (HealthcareNLP)" and audiences do not need prior knowledge to attend this. Tutorial materials <https://github.com/4dpicture/HealthNLP>

1 Introduction

With the increasing capabilities of current state-of-the-art NLP models, NLP applications have been one of the hot topics in AI for social good, which includes their application in the healthcare domain. It is not a brand new topic, since there have been NLP researchers trying to deploy NLP techniques into the biomedical domain a long time ago, and the BioNLP workshop has been held since 2002 for 23 years¹. There is indeed growing interest from both researchers and practitioners, such as the recently emerged workshops including CL4Health at LREC-COLING2024 and NAACL2025 (Ananiadou et al., 2025), 6th HealthNLP at ICHI2023², Big Data Analytics in Healthcare at IEEE-BigData2023³, ClinicalNLP⁴ etc. new shared tasks such as BioLaySumm (Xiao et al., 2025), PLABA (plain language adaptation of biomedical abstract) (Ondov et al., 2025), etc.

This tutorial starts with an explanation of each sub-task into three layers (data resource, NLP, and patients), how they are defined, and what NLP approaches were used to address such tasks in history, up to the state-of-the-art models.

First, in the resource layer, we talk about 1.1) Annotation Guidelines: how the clinical documents shall be annotated and anonymised to protect privacy and address ethical concerns. 1.2) Ethical Approval and Data Governance: how data should be hosted and how access is granted for research purposes. 1.3) Synthetic

¹https://aclweb.org/aclwiki/BioNLP_Workshop

²<https://www.healthnlp.info/>

³<https://onlineacademiccommunity.uvic.ca/healthcare4bigdata4analytic/>

⁴<https://aclanthology.org/venues/clinicalnlp/>




Figure 1: An overview of HealthcareNLP tasks (up-down): 1) `data-layer` annotation guidelines, ethical approvals, governance, and synthetic data applications; 2) `NLP-Eval-layer` tasks on named entity recognition, relation extractions, sentiment analysis, entity linking, and possible NLP techniques including traditional rule/ML, LLMs with KG, retrieval-augmented generation, explainable AI (XAI); 3) `patients-layer` related event: public patient involvement and engagement; text simplification / lay-language adaptation and summarisation for public health literacy, and machine translation for multilingual setting (L2 learners in a new country), shared decision making (SDM) communication support.

data generation (connecting to NLP): how this task can be useful to address data sharing and scarcity issues. Second, in the NLP-eval layer, these tasks include 2.1) standard named entity recognition (NER) in biomedical and clinical domains such as disease, symptoms, diagnoses, treatments, and progression entity categories. 2.2) relation extraction (RE) such as drugs and side-effects, symptoms and treatments, treatments and temporality, etc. 2.3) sentiment analysis of NLP models on patient-centred text data, including patient interview transcripts, online forums, and consultations. 2.4) Clinical Coding (normalisation/linking), which automates the traditional human coders work by linking the clinical entities into clinical knowledge graphs and dictionaries; 2.5) Machine Translation (MT): on translating clinical text from one language to others to address low-resource eHealth model training obstacles as well as in a multilingual consultation environment where the patients cannot speak/understand the dominant language; 2.6) Text Simplification and Summarisation: to help

patients better understand the clinical evaluation documents and healthcare information websites. We will introduce the method applied to these tasks from historical rule-based, to statistical, and nowadays neural network models with attention mechanisms. If time allows, we will cover the state-of-the-art retrieval-augmented generation (RAG), large language models with knowledge graphs (LLM+KG), and prompt engineering, as shown in Figure 1 in the overview of HealthcareNLP picture. Third, in the patient layer, we also mention some engaging events in this field, including 3.1) public and patient involvement and engagement (PPIE) practice examples from the UK and EU projects, 3.2) patient health literacy, and 3.3) how to better support the communication between clinicians and patients for better shared decision making (SDM).

2 Tutorial organisers

- Lifeng Han, Postdoc, Leiden University Medical Center, l.han@lumc.nl, NL, researcher in ClinicalNLP from UK and EU

funded projects.

- Paul Rayson, Professor, Lancaster University, UK, p.rayson@lancaster.ac.uk, Director of UCREL research centre that combines research in corpus linguistics and NLP.
- Suzan Verberne, Professor of Natural Language Processing, Leiden University, The Netherlands, s.verberne@liacs.leidenuniv.nl, Suzan Verberne is professor of Natural Language Processing at the Leiden Institute of Advanced Computer Science (LIACS). She has supervised projects involving a large number of application domains and collaborations, many of which in health and medical contexts.

3 Target Audience

This tutorial will be developed for NLP researchers who are interested in applying their models and methods to the healthcare domain, which potentially can have a real impact on social lives and healthcare systems, making them more efficient and accurate. It is also useful for healthcare practitioners and data scientist in the hospital who go to the NLP conference to learn more knowledge on how the models they use are developed for healthcare. For UG and PG students who are interested in AI and NLP applications, this tutorial can also be of interest as an introductory level to AI for the healthcare field. Prior knowledge is not required to attend this event. By the end of the tutorial, the audiences shall acquire an overview of HealthcareNLP and CL4Health, what tasks are there, what models are being used, and what are the current challenges.

4 Outline

The tutorial aims to be 4 hours long including 30 minute break in between. The covered topics include the following:

- Data resources: annotation guidelines, anonymisation, ethical approval, governance, synthetic data usage
- HealthcareNLP Tasks and Evaluations: NER, RE, Linking, Sentiment Classification,

tion, Simplification, Translation, Summarisation, QA

- Models applied: rule, statistical, neural models, LLM+KG, prompting, RAG, Explainable and Interpretable AI
- Patients: Public patient involvement and engagement (PPIE), Patient health literacy, Shared Decision Making (SDM) ([Stiggelbout et al., 2015](#); [Frosch and Kaplan, 1999](#))
- Real-world applications of HealthcareNLP projects and models in the UK / NL, and EU project 4D PICTURE consortium <https://4dpicture.eu/>
- Hands-on experience of using some healthcareNLP platforms and applications, including our own.

5 Technical Requirements

Our technical expectations are that participants will have access to a good internet connection for online code and data experiments.

6 Diversity Considerations

The tutorial includes an MT sub-topic on translating healthcare domain data into low-resource languages so that researchers can help to build Healthcare AI/NLP systems for low-resource communities.

7 Reading List

Healthcare and Clinical NLP ([Jerfy et al., 2024](#); [Nazi and Peng, 2024](#); [Elvas et al., 2025](#)).

8 Presenters

Dr. Lifeng Han is currently a researcher on the 4D Picture project <https://4dpicture.eu/> researching the cancer patient journey support. He had his first postdoctoral position with the University of Manchester from 2022 to 2024 on HealthcareNLP funded by the UKRI. Prior to these, he finished his PhD thesis "An investigation into multi-word expressions in machine translation" ([Han, 2022](#)) from ADAPT centre, Ireland. He initiated the joint track between two workshops MWE and ClinicalNLP at MWE2023 with EACL ([Bhatia et al., 2023](#))

when he was co-chairing MWE2023. He gave a main conference tutorial to LREC2022 in person / on-site on “Meta-Evaluation of Translation Evaluation Methods: a systematic up-to-date overview” at Marseille, France, 2022. His recent publications mainly focused on HealthcareNLP in the past three years. He has 50+ peer-reviewed publications, including journal articles and NLP main conference papers.

Professor Paul Rayson is based in the School of Computing and Communications at Lancaster University, UK and Director of the UCREL interdisciplinary research centre which carries out research in corpus linguistics and natural language processing (NLP). A long term focus of his work is semantic multilingual NLP in extreme circumstances where language is noisy e.g. in historical, learner, speech, email, txt and other CMC varieties. Along with domain experts, he has applied his research in the areas of metaphor in cancer narratives, dementia detection, mental health, online child protection, cyber security, learner dictionaries, and text mining of biomedical literature, historical corpora, and financial narratives. He was a co-investigator of the five-year ESRC Centre for Corpus Approaches to Social Science (CASS) which was designed to bring the corpus approach to bear on a range of social sciences.

Dr. Andrew Moore is a research software engineer at Lancaster University, also currently working on semantic tagging and emotion analysis within the 4D picture project, with a PhD on empirical evaluation methodology for target dependent sentiment analysis (Moore, 2021) as well as conference papers on this topic. He has worked on semantic tagging while working at Lancaster University creating, maintaining, and extending the Python Multilingual Ucrel Semantic Analysis System (PyMUSAS)⁵ software with Professor Paul Rayson. He has also previously worked within industry at the World Intellectual Property Organization (WIPO) working on topics such as Machine Translation Quality Estimation and semantic search.

Professor Goran Nenadic is based in Department of Computer Science, University of Manchester, with expertise in HealthNLP, as co-/PI of multiple UK funded projects. Current

research projects focus on large-scale extraction and curation of biomedical information and clinical/epidemiological findings, by combining rule-based and data-driven approaches. He is also interested in processing healthcare social media. He leads the UK healthcare text analytics network (Healtex) and was the founding Editor-in-Chief of Journal of Biomedical Semantics.

9 Other Information

The number of estimated audience is 50-100. This estimation is partially from Dr Han's last tutorial at LREC2022.

10 Ethics Statement

All the examples and corpora used in the tutorial are either public data from shared task challenges approved by the shared task organisers or anonymised data samples from our research projects.

11 Topic-aware Literature

- De-identification (Shaji et al., 2025)
- Healthcare Relation Extraction (Cui et al., 2023; Tu et al., 2023)
- Clinical NER and Linking (Belkadi et al., 2023; Romero et al., 2025a)
- Clinical Coding with explainability (Dong et al., 2021)
- Biomedical Text Simplification - PLABA task (Li et al., 2024; Attal et al., 2023) and Summarisation (PerAnSum) (Romero et al., 2025b)
- Synthetic Data for HealthcareNLP (Ramesh et al., 2024; Belkadi et al., 2025; Ren et al., 2025)
- Annotation Guidelines (Schulz et al., 2023)
- Machine Translation in biomedical domain (Mehandru et al., 2022; Han et al., 2024)
- Sentiment Modelling (Dumbach et al., 2024)
- AI for patient-clinician communication (Han et al., 2025)

⁵<https://ucrel.github.io/pymusas/>

- Embedded ethics (Bak et al., 2025)
- LLM prompting (in-context learning) (Dong et al., 2024)
- Future Perspectives in Figure 2

12 Acknowledgement

We thank Ida Korfage, Sheila Payne, and Judith Spek for their feedback on the abstract of this work. The 4D PICTURE: Funded by the European Union under Horizon Europe Work Programme 101057332. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Health and Digital Executive Agency (HaDEA). Neither the European Union nor the granting authority can be held responsible for them. The UK team are funded under the Innovate UK Horizon Europe Guarantee Programme, UKRI Reference Number: 10041120.

GN is supported by the grant “Assembling the Data Jigsaw: Powering Robust Research on the Causes, Determinants, and Outcomes of MSK Disease”, and the grant “Integrating hospital outpatient letters into the healthcare data space” (EP/V047949/1; funder: UKRI/EP SRC).

References

- Sophia Ananiadou, Dina Demner-Fushman, Deepak Gupta, and Paul Thompson, editors. 2025. *Proceedings of the Second Workshop on Patient-Oriented Language Processing (CL4Health)*. Association for Computational Linguistics, Albuquerque, New Mexico.
- Kush Attal, Brian Ondov, and Dina Demner-Fushman. 2023. A dataset for plain language adaptation of biomedical abstracts. *Scientific Data*, 10(1):8.
- Marieke Bak, Laura Hartman, Charlotte Graafland, Ida J Korfage, Alena Buyx, Maartje Schermer, 4D PICTURE Consortium, et al. 2025. Ethical design of data-driven decision support tools for improving cancer care: embedded ethics review of the 4d picture project. *JMIR cancer*, 11(1):e65566.
- Samuel Belkadi, Lifeng Han, Yuping Wu, and Goran Nenadic. 2023. *Exploring the value of pre-trained language models for clinical named entity recognition*. In *2023 IEEE International Conference on Big Data (BigData)*, pages 3660–3669.
- Samuel Belkadi, Libo Ren, Nicolo Micheletti, Lifeng Han, and Goran Nenadic. 2025. Mlm4synmed: Masked language modelling for synthetic free-text medical records generation. In *2025 IEEE 13th International Conference on Healthcare Informatics (ICHI)*, pages 531–542. IEEE.
- Archana Bhatia, Kilian Evang, Marcos Garcia, Voula Giouli, Lifeng Han, and Shiva Taslimipour, editors. 2023. *Proceedings of the 19th Workshop on Multiword Expressions (MWE 2023)*. Association for Computational Linguistics, Dubrovnik, Croatia.
- Yang Cui, Lifeng Han, and Goran Nenadic. 2023. *MedTem2.0: Prompt-based temporal classification of treatment events from discharge summaries*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, pages 160–183, Toronto, Canada. Association for Computational Linguistics.
- Hang Dong, Víctor Suárez-Paniagua, William Whiteley, and Honghan Wu. 2021. Explainable automated coding of clinical notes using hierarchical label-wise attention networks and label embedding initialisation. *Journal of biomedical informatics*, 116:103728.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, Xu Sun, Lei Li, and Zhifang Sui. 2024. *A survey on in-context learning*. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1107–1128, Miami, Florida, USA. Association for Computational Linguistics.
- Philipp Dumbach, Leo Schwinn, Tim Löhr, Phi Long Do, and Bjoern M Eskofier. 2024. Artificial intelligence trend analysis on healthcare podcasts using topic modeling and sentiment analysis: a data-driven approach. *Evolutionary Intelligence*, 17(4):2145–2166.
- Luis B Elvas, Ana Almeida, and João C Ferreira. 2025. Natural language processing in medical text processing: A scoping literature review. *International Journal of Medical Informatics*, page 106049.
- Dominick L Frosch and Robert M Kaplan. 1999. Shared decision making in clinical medicine: past research and future directions. *American journal of preventive medicine*, 17(4):285–294.
- Lifeng Han. 2022. *An investigation into multi-word expressions in machine translation*. Ph.D. thesis, Dublin City University.




Figure 2: NLP-related tasks and their mappings to applications (Future Perspectives)

Lifeng Han, Serge Gladkoff, Gleb Erofeev, Irina Sorokina, Betty Galiano, and Goran Nenadic. 2024. Neural machine translation of clinical text: an empirical investigation into multilingual pre-trained language models and transfer-learning. *Frontiers in Digital Health*, 6:1211564.

Lifeng Han, David Lindevelt, Sander Puts, Erik van Mulligen, and Suzan Verberne. 2025. Dutch metaphor extraction from cancer patients' interviews and forum data using llms and human in the loop. *arXiv preprint arXiv:2511.06427*.

Aadit Jerfy, Owen Selden, and Rajesh Balkrishnan. 2024. The growing impact of natural language processing in healthcare and public health. *INQUIRY: The Journal of Health Care Organization, Provision, and Financing*, 61:00469580241290095.

Zihao Li, Samuel Belkadi, Nicolo Micheletti, Lifeng Han, Matthew Shardlow, and Goran Nenadic. 2024. Investigating large language models and control mechanisms to improve text readability of biomedical abstracts. In *2024 IEEE 12th International Conference on Healthcare Informatics (ICHI)*, pages 265–274. IEEE Computer Society.

Nikita Mehandru, Samantha Robertson, and Nilofar Salehi. 2022. Reliable and safe use of machine translation in medical settings. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 2016–2025.

Andrew Moore. 2021. *Empirical evaluation methodology for target dependent sentiment analysis*. Phd thesis, Lancaster University (United Kingdom).

Zabir Al Nazi and Wei Peng. 2024. Large language models in healthcare and medical domain: A review. In *Informatics*, volume 11, page 57. MDPI.

Brian Ondov, William Xia, Kush Attal, Ishita Unde, Jerry He, and Dina Demner-Fushman. 2025. Lessons from the trec plain language adaptation of biomedical abstracts (plaba) track. *ArXiv*, pages arXiv–2507.

Krithika Ramesh, Nupoor Gandhi, Pulkit Madaan, Lisa Bauer, Charith Peris, and Anjalie Field. 2024. [Evaluating differentially private synthetic data generation in high-stakes domains](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15254–15269, Miami, Florida, USA. Association for Computational Linguistics.

Libo Ren, Samuel Belkadi, Lifeng Han, Warren Del-Pinto, and Goran Nenadic. 2025. Synthetic4health: generating annotated synthetic clinical letters. *Frontiers in Digital Health*, 7:1497130.

Pablo Romero, Lifeng Han, and Goran Nenadic. 2025a. Medication extraction and entity linking using stacked and voted ensembles on llms. In *Proceedings of the Second Workshop on Patient-Oriented Language Processing (CL4Health)*, pages 303–315.

Pablo Romero, Libo Ren, Lifeng Han, and Goran Nenadic. 2025b. The manchester bees at peransumm 2025: Iterative self-prompting with claude and o1 for perspective-aware healthcare answer summarisation. In *Proceedings of the Second Workshop on Patient-Oriented Language Processing (CL4Health)*, pages 340–348.

Stefan Schulz, Warren Del-Pinto, Lifeng Han, Markus Kreuzthaler, Sareh Aghaei Dinani, and Goran Nenadic. 2023. [Towards principles of ontology-based annotation of clinical narratives](#). In *International Conference on Biomedical Ontology*.

Dhivin Shaji, Angel Paul, Lifeng Han, Warren Del-Pinto, Goran Nenadic, and Suzan Verberne. 2025. De-identifying clinical texts using biomedical bert and comprehensive risk assessment. In *2025 IEEE 13th International Conference on Healthcare Informatics (ICHI)*, pages 683–684. IEEE.

Anne M Stiggelbout, Arwen H Pieterse, and Johanna CJM De Haes. 2015. Shared decision making: concepts, evidence, and practice. *Patient education and counseling*, 98(10):1172–1179.

Hangyu Tu, Lifeng Han, and Goran Nenadic. 2023. Extraction of medication and temporal relation from clinical text using neural language models. In *2023 IEEE International Conference on Big Data (BigData)*, pages 2735–2744. IEEE.

Chenghao Xiao, Kun Zhao, Xiao Wang, Siwei Wu, Sixing Yan, Tomas Goldsack, Sophia Ananiadou, Noura Al Moubayed, Liang Zhan, William K Cheung, et al. 2025. Overview of the biolay-summ 2025 shared task on lay summarization of biomedical research articles and radiology reports. In *Proceedings of the 24th Workshop on Biomedical Language Processing*, pages 365–377.