



LREC 2026

**The Third Workshop on Patient-Oriented Language
Processing (CL4Health) @ LREC 2026**

Workshop Proceedings

Editors

**Deepak Gupta, Paul Thompson, Sophia Ananiadou and
Dina Demner-Fushman**

12 May 2026

©ELRA Language Resources Association (ELRA), 2026
These proceedings are licensed under a Creative Commons Attribution-
NonCommercial 4.0 International License (CC BY-NC 4.0)

ISBN 978-2-493814-51-7

Preface

Introduction

We are pleased to present the proceedings of the Third Workshop on Patient-Oriented Language Processing (CL4Health), held in conjunction with LREC 2026. This workshop continues to provide a dedicated forum for research on the spectrum of patient-oriented language processing and multimodal research. Following the success of the previous editions held at LREC-COLING 2024 and NAACL 2025, CL4Health has further established itself as a dedicated venue for advancing language technologies that support public health, including the summarization of health records for patient use, answering consumer health questions using reliable sources, detecting misinformation or potentially harmful content, and providing multimodal response when it more effectively supports patient understanding.

Such a venue is needed both to invigorate patient-oriented language processing research and to build a community of researchers interested in this area. The growing interest in this topic is fueled by several current trends:

- Rapid use of large language models for health counseling, diagnostic reasoning, and patient-facing question answering.
- Focus on simplifying and adapting clinical information through summarization, rewriting, and persona-based generation.
- Growth of multimodal tools for patient journeys, care pathways, and communication support.
- Increasing emphasis on safety, factual grounding, and standardized evaluation of patient-oriented systems.

The program includes contributions covering a broad spectrum of topics, including question answering, patient-oriented lay summarization, medical text simplification, and multimodal learning in the biomedical and consumer health domains. In addition to regular research papers, the workshop features four shared tasks, designed to encourage standardized evaluation and facilitate comparative analysis across systems. These shared tasks highlight current progress and persistent challenges in developing robust and trustworthy patient-oriented language technologies.

Invited Speaker

Kailai Yang, University of Manchester, UK

Towards Trustworthy and Efficient Continual Learning for Large Language Models

Current Large Language Models (LLMs) continue to face several key limitations in adapting to specific domains. First, their ability to model human values and learn rich representations remains limited. Second, general-purpose LLMs often lack sufficient domain-specific

knowledge and alignment for professional applications. Third, existing continual learning methods are still prone to catastrophic forgetting, instability in multi-objective optimization, and limited training efficiency. This research addresses these challenges from three complementary directions. First, we improve representation learning for language models through knowledge fusion and innovative paradigms, including contrastive and disentangled learning methods. Second, we strengthen domain alignment via domain knowledge injection, in-context learning, and advanced post-training approaches. Third, we improve continual learning for LLMs by proposing novel automated data-mixing strategies, selective preference optimization method for efficient training, and a plug-and-play multi-objective alignment framework that generalizes across models and objectives. Overall, this research provides practical and effective solutions to enhance the reliability and efficiency for domain adaptation for LLMs.

Brief Biography

Kailai Yang is an incoming researcher at the University of Manchester. His main research areas are artificial intelligence and large language models. He has authored over 20 papers with over 2,000 citations in total. His research as leading contributor has been published in leading conferences and journals, such as NeurIPS, ACL, EMNLP, WWW, and KDD. His work has also been featured by renowned media such as Forbes.

Shared Tasks

- **Detecting Dosing Errors from Clinical Trials (CT-DEB'26).**

The Clinical Trials Dosing Errors Benchmark 2026 (CT-DEB'26) is dedicated to automated detection of the risks of medication dosing errors within clinical trial protocols. Leveraging a curated dataset of over 29K trial records from the ClinicalTrials.gov registry, participants are challenged to predict the risk probabilities of protocols likely to result in dosing errors.

- **Automatic Case Report Form (CRF) Filling from Clinical Notes.**

Case Report Forms (CRFs) are standardized instruments in medical research that collect patient data in a consistent and reliable manner. They consist of a predefined list of items to be filled with patient information. Each item aims to collect a portion of the information relevant to a specific clinical goal (e.g., allergies, disease chronicity, test results). Automating CRF filling from clinical notes would accelerate clinical research, reduce manual burden on healthcare professionals, and create structured representations that can be directly leveraged to produce accessible, patient- and practitioner-friendly summaries.

- **ArchEHR-QA 2026: Grounded Question Answering from Electronic Health Records.**

ArchEHR-QA targets the problem of producing answers to patient questions that are supported by and explicitly linked to the underlying clinical notes. This second iteration builds on the 2025 challenge (co-located with the ACL 2025 BioNLP Workshop) by expanding the dataset and introducing four complementary subtasks covering question interpretation, clinical evidence identification, answer generation, and answer–evidence alignment.

- **FoodBench-QA 2026: Grounded Food & Nutrition Question Answering.**

FoodBench-QA 2026 is a shared task challenging systems to answer food and nutrition questions using evidence from nutrient databases and food ontologies. The dataset

includes realistic dietary queries, ingredient and their quantities lists, and recipe descriptions, requiring models to perform nutrient estimation, FSA traffic-light prediction, and food entity recognition/linking across three food semantic models.

Submissions

CL4Health received 36 valid submissions for the main workshop, of which 11 were accepted as oral presentations, 17 as posters, and 8 were rejected. The work covers a wide range of topics focusing on patients' well-being and proper care. The topics include question answering, medical imaging, patient-oriented lay summarization, medical text simplification, mental health issues, and communication support. We received 28 valid submissions for the shared task papers, of which 25 were accepted as posters and 3 were rejected.

We would like to thank all the authors for their submissions. We are grateful to the shared task organizers for their efforts in designing and coordinating the evaluation benchmarks. Our sincere thanks go to the invited speaker and participants for their engagement and contributions to the workshop. Finally, we acknowledge the Program Committee (listed elsewhere in this volume) for their dedicated efforts in evaluating the submissions within a short review period and identifying work ready for presentation, as well as studies that would benefit from further experimentation and analysis informed by reviewer comments. We hope that the work presented in these proceedings will contribute to ongoing research in patient-centered language technologies and foster further advances in accessible, reliable, and effective language technologies for public health.

Deepak Gupta, Paul Thompson, Sophia Ananiadou and Dina Demner-Fushman (Organizers).

Organizing Committee

Deepak Gupta, National Library of Medicine, USA
Paul Thompson, University of Manchester, UK
Sophia Ananiadou, University of Manchester, UK & Archimedes/Athena Research Center, Greece
Dina Demner-Fushman, National Library of Medicine, USA

Shared Task Organizers

Detecting Dosing Errors from Clinical Trials (CT-DEB'26)

Sohrab Ferdowsi, University of Geneva, Switzerland
Félicien Hêche, University of Geneva, Switzerland
Anthony Yazdani, University of Geneva, Switzerland
Edward Choi, KAIST, Korea
Douglas Teodoro, University of Geneva, Switzerland

Automatic Case Report Form (CRF) Filling from Clinical Notes

Pietro Ferrazzi, Fondazione Bruno Kessler, Trento, Italy
Soumitra Ghosh, Fondazione Bruno Kessler, Trento, Italy
Alberto Lavelli, Fondazione Bruno Kessler, Trento, Italy
Bernardo Magnini, Fondazione Bruno Kessler, Trento, Italy

ArchEHR-QA 2026: Grounded Question Answering from Electronic Health Records

Sarvesh Soni, National Library of Medicine, USA
Dina Demner-Fushman, National Library of Medicine, USA

FoodBench-QA 2026: Grounded Food & Nutrition Question Answering

Tome Eftimov, Jožef Stefan Institute, Ljubljana, Slovenia
Ana Gjorgjevikj, Jožef Stefan Institute, Ljubljana, Slovenia
Matej Martinc, Jožef Stefan Institute, Ljubljana, Slovenia
Gjorgjina Cenikj, Jožef Stefan Institute, Ljubljana, Slovenia
Sašo Džeroski, Jožef Stefan Institute, Ljubljana, Slovenia
Barbara Koroušić Seljak, Jožef Stefan Institute, Ljubljana, Slovenia

Program Committee

Sophia Ananiadou, University of Manchester, UK & Archimedes/Athena Research Center, Greece
Judith Jeyafreeda Andrew, Université Paris Cité, Paris, France
Mohammad Arvan, University of Illinois Chicago, USA
Zainab Awan, Queen Mary University of London, UK
Abby Blocker, University of Cape Town, South Africa
Giovanni Bonetta, Fondazione Bruno Kessler, Italy
Leonardo Campillos-Llanos, Spanish National Research Council, Spain
Gjorgjina Cenikj, Jožef Stefan Institute, Slovenia
Surabhi Datta, IMO Health, USA
Dina Demner-Fushman, National Library of Medicine, USA
Luise Duerlich, Research Institutes of Sweden, Sweden
Tome Eftimov, Jožef Stefan Institute, Slovenia

Sohrab Ferdowsi, University of Geneva, Switzerland
Pietro Ferrazzi, Fondazione Bruno Kessler, Italy
Manas Gaur, University of Maryland Baltimore County, USA
Soumitra Ghosh, Fondazione Bruno Kessler, Italy
Natalia Grabar, Université de Lille, France
Tudor Groza, Bioinformatics Institute, A*STAR, Singapore
Deepak Gupta, National Library of Medicine, USA
Félicien Hêche, University of Geneva, Switzerland
Sy Hwang, University of Pennsylvania, USA
Samah Jarad-Fodeh, Yale School of Medicine, USA
Dimitrios Kokkinakis, University of Gothenberg, Sweden
Saran Krishnasamy, Gigit.ai, USA
Tiziano Labruna, University of Padova, Italy
Vojtech Lanz, Charles University, Czech Republic
Alberto Lavelli, Fondazione Bruno Kessler, Italy
Chuyuan Li, University of British Columbia, Canada
Matej Martinc, Jožef Stefan Institute, Slovenia
Vineet Mishra, HCLTech, USA
Anindita Nath, University of Texas at El Paso, USA
Aleksandr Nesterov, Artificial Intelligence Research Institute, USA
Aurélié Névél, Université de Paris-Saclay, CNRS, LISN, Orsay, France
Brian Ondov, Yale School of Medicine, USA
Tabea Pakull, University Hospital Essen, Germany
Mehdi Parviz, University of Copenhagen, Denmark
Priyanshu Priya, Indian Institute of Technology Patna, India
Roland Roller, German Research Center for Artificial Intelligence (DFKI), Germany
Roshan Santhosh, Adobe, USA
Abeed Sarker, Emory University, USA
Apoorva Singh, Fondazione Bruno Kessler, Italy
Sarvesh Soni, National Library of Medicine, USA
Suzanne Tamang, Stanford University, USA
Douglas Teodoro, Fondazione Bruno Kessler, Italy
Paul Thompson, University of Manchester, UK
Talia Tseriotou, Queen Mary University of London, UK
Grigorios Tsoumakas, Aristotle University of Thessaloniki, Greece
Shubham Vatsal, New York University, USA
Jinghua Xu, Heidelberg University, Germany
Zhicheng Yang, PAII Inc., USA
Anthony Yazdani, University of Geneva, Switzerland
Tianlin Zhang, CHN Energy, China
Pierre Zweigenbaum, LISN, CNRS, Université Paris-Saclay, France

Table of Contents

<i>FHIRPath-QA: Executable Question Answering over FHIR Electronic Health Records</i> Michael Frew, Nishit Bheda and Bryan Tripp	1
<i>COACH Meets QUORUM: A Framework and Pipeline for Aligning User, Expert, and Developer Perspectives in LLM-Generated Health Counselling</i> Yee Man Ng, Bram van Dijk, Pieter Beynen, Otto Boekesteijn, Joris Jansen, Gerard van Oortmerssen, Max J. van Duijn and Marco Spruit	15
<i>Addressing Domain Shift in Health Coaching Note Analysis through Factorized Synthetic Data Generation</i> Michael Tänzer, Iva Bojic, Ashwini Yuvraj Lawate, Andy Hau Yan Ho and Andy Khong .	26
<i>Scalable Generation of Adult-Oriented Therapeutic Reading Texts for Russian Aphasia Rehabilitation</i> Anastasia Kolmogorova, Anastasia Margolina, Alina Telnova and Igor Ilchenko	41
<i>An Open-Resource Knowledge Augmentation for Biomedical Lay Summarization</i> João Pedro Veloso and Evelin Amorim	50
<i>TabMedQA: From Structured Data to Question-Answer Datasets in Early Clinical Decision-Making</i> Gabriel Iturra Bocaz, Petra Galuščáková, Sol Gedde Vedde and Alvaro Fernandez-Quilez	61
<i>SimpliMED: Automatic Simplification of Cardiology Discharge Reports Using Large Language Models</i> Lucas Molino-Piñar, Manuel Carlos Díaz Galiano, María-Teresa Martín-Valdivia, Jose Angel Urbano-Moral and Elena Sola-Garcia	72
<i>Dutch Metaphor Extraction from Cancer Patients' Interviews and Forum Data Using LLMs and Human in the Loop</i> Lifeng Han, David Lindevelt, Sander Puts, Erik van Mulligen and Suzan Verberne	82
<i>Compressed Representations of Patient Records: A Comparative Study of Template-Based and LLM-Based Methods for Clinical Data Summarization and Visualization</i> Andreas Stöckl, Oliver Krauss and Sophie Bauernfeind	96
<i>Medical Text Rewriting for Non-Experts: A Guideline-Driven LLM Approach</i> Mana Kuramoto, Hiroyuki Nagai, Keiko Yamada, Hiroo Ide, Masayo Hayakawa, Tomohiro Nishiyama, Shoko Wakamiya and Eiji Aramaki	107
<i>Patient-Specific Care Pathway Visualisation for Medical Nursing Staff</i> Sophie Bauernfeind, Selina Adlberger, Oliver Krauss and Andreas Stöckl	117
<i>What Makes a Good Doctor Response? A Study on Text-Based Telemedicine</i> Adrian Cosma, Cosmin Dumitrache and Emilian Radoi	132
<i>Datasets for a Chatbot for Clinical Trial Search</i> Yumeng Yang, Ethan Ludmir and Kirk Roberts	139

<i>HealthTrajectory: Patient Journey Summaries and Visualizations for Patient-Clinician Communication Support</i>	
Rohmah Hidayah, Tomohiro Nishiyama, Shoko Wakamiya and Eiji Aramaki	149
<i>Medical-FLAVORS-AECC: Spanish Oncological Metaphors Dataset</i>	
Lucia Pitarch, Jordi Bernad, Sergio-Luis Ojeda-Trueba, Alec Sánchez-Montero, Maxim Ionov, Emma Anglés-Herrero, Ángel Óscar Corona Beomont and Gemma Bel-Enguix	159
<i>A Synthetic Conversational Dataset for Type 2 Diabetes Management</i>	
Stergios Ntanavaras, Maaïke de Boer and Piek T.J.M. Vossen	171
<i>Italian Medical Term Simplification: From Patient Information Leaflets to Simplified Language Resources</i>	
Maria Pia di Buono	182
<i>Evaluating Professional Acceptability of LLM-Generated Systematic Review Summaries in Healthcare: Psychiatrists' Perspectives</i>	
Paul Thompson, Artemis Boulogeorgou, Fotini Kaponi, Efstathia Soufleri and Sophia Ananiadou	190
<i>Reasoning, Contrastive, and In-Context Strategies for Opioid Use Stage Detection on Social Media</i>	
Vinu Ekanayake and Ramakanth Kavuluru	207
<i>FoodBench-QA: Overview of the Shared Task on Grounded Food and Nutrition Question Answering</i>	
Tome Eftimov, Ana Gjorgjevikj, Matej Martinc, Gjorgjina Cenikj, Sašo Džeroski and Barbara Koroušič Seljak	223
<i>Overview of the CT-DEB'26 Shared Task on Predicting Dosing Errors in Interventional Clinical Trials</i>	
Sohrab Ferdowsi, Félicien Hêche, Anthony Yazdani, Edward Choi, Sara Sansaloni-Pastor and Douglas Teodoro	234
<i>Overview of the CRF 2026 Shared Task on Clinical Case Report Forms Filling</i>	
Pietro Ferrazzi, Soumitra Ghosh, Alberto Lavelli and Bernardo Magnini	245
<i>Overview of the ArchEHR-QA 2026 Shared Task on Grounded Question Answering from Electronic Health Records</i>	
Sarvesh Soni and Dina Demner-Fushman	255
<i>Structured Radiology Intelligence: Extracting Structured Data from MRI Reports Using LLMs</i>	
Sushvin Marimuthu, Parameswari Krishnamurthy, Dipti Misra Sharma, Goldwin H, Anu Eapen, Betty Simon and Anuradha Chandramohan	268
<i>Useful to Whom? A Persona-Driven Evaluation of Knowledge-Adapted Health Question Reformulation via LLM Simulation</i>	
Jooyeon Lee, Luan Huy Pham and Özlem Uzuner	281
<i>MedGore: An Approach and a Dataset for Identification of Sensitive Medical Images</i>	
Soumya Gayen, Rory Mulcahey, Russell Loane, Dina Demner-Fushman and Deepak Gupta	296

<i>Diagnostic Reasoning with Large Language Models for a Rare Disease: Case Study of Primary Ciliary Dyskinesia</i>	
Swati Rajwal, Mary Ellen M. Fain, Lokesh Guglani and Abeed Sarker	307
<i>Beyond One-Size-Fits-All: Multi-Agent Refinement Framework for Persona-Based Biomedical Summarization</i>	
Rohan Charudatt Salvi, Chirag Chawla, Md. Shad Akhtar and Shweta Yadav	320
<i>Automated Detection of Dosing Errors in Clinical Trial Narratives: A Multi-Modal Feature Engineering Approach with LightGBM</i>	
Mohammad AL-Smadi	333
<i>CaresAI at CT-DEB'26: Detecting Dosing Errors In Clinical Trials Using Domain-Specific Transformer Embeddings and Classification Models</i>	
Leon Hamnett, Favour Igwezeke, Joseph Itopa Abubakar and Mary Adetutu Adewunmi	344
<i>CGU-ILALab at FoodBench-QA 2026: Comparing Traditional and LLM-based Approaches for Recipe Nutrient Estimation</i>	
Wei-Chun Chen, Yu-Xuan Chen, I-Fang Chung and Ying-Jia Lin	362
<i>FCProfiler: Structured and Deterministic Pipeline for Recipe-Level Nutrient Estimation</i>	
Siyoon Kim	369
<i>DocUA at CRF Filling 2026: LLM StructCore — Schema-Guided Reasoning Condensation and Deterministic Compilation</i>	
Serhii Zabolotnii	375
<i>GREYC at CRF Filling 2026: Rewrite Before You Extract - Rewriting Clinical Notes for Automated CRF</i>	
Jesus Lovon-Melgarejo, J��r��mie Pantin and Ga��l Dias	384
<i>Innov8rs at CRF Filling 2026: An Iterative Multi-LLM Ensemble Pipeline with Dynamic Few-Shot Retrieval and Data-Driven Precision Filtering</i>	
Samminga Sainath Rao, Sumit Mishra and Chanchal Suman	390
<i>Aurum at CRF Filling 2026: Modular DSPy Extractors with Qwen3-Max for Multilingual CRF Filling</i>	
Vinay Babu Ulli, Jyoti Kumari and Anindita Mondal	395
<i>sebis at CRF Filling 2026: A Two-Stage Local LLM Pipeline for Medical CRF Filling</i>	
Katharina Sommer, Tristan Till and Florian Matthes	402
<i>Polimi at CRF Filling 2026: Prompt-Based Information Extraction from Italian Clinical Notes</i>	
Vittorio Torri and Francesca Ieva	412
<i>Cohere Labs Community at FoodBench-QA 2026: The Cake Makes the Ingredients</i>	
Ravi Ranjan, Roshan Santhosh and Lucien Carroll	428
<i>HiTZ-IXA at ArchEHR-QA 2026: Evidence Alignment Through Self-Consistency and Prompt Curation in Memory-Constrained Environments</i>	
Xabier Irastortza-Urbieta, Maite Oronoz and Alicia P��rez	434

<i>BIT.UA-AAUBS at ArchEHR-QA 2026: Evaluating Open-Source and Proprietary LLMs via Prompting in Low-Resource QA</i>	
Richard A. A. Jonker, Alexander Christiansen, Alexandros Maniatis, Rúben Garrido, Rogério Braunschweiger de Freitas Lima, Roman Jurowetzki and Sérgio Matos	441
<i>WisPerMed at ArchEHR-QA 2026: Retrieval-Augmented Prompting for Grounded EHR Question Answering</i>	
Jan-Henning Büns, Tabea Margareta Grace Pakull, Hendrik Damm, Bohao Chu, Christoph M. Friedrich, Felix Nensa, Elisabeth Livingstone, Peter A. Horn and Norbert Fuhr	455
<i>sebis at ArchEHR-QA 2026: How Much Can You Do Locally? Evaluating Grounded EHR QA on a Single Notebook</i>	
Ibrahim Ebrar Yurt, Fabian Tobias Karl, Tejaswi Chopra and Florian Matthes	469
<i>MedEvi-NS at ArchEHR-QA 2026: Using Clinical Reasoning Principles to Improve Zero-shot Capabilities of Large Language Models in Evidence Alignment</i>	
Mengxuan Sun and Nicolay Rusnachenko	482
<i>UIC-AIHealth4All at ArchEHR-QA 2026: Answer-First Evidence Grounding for Clinical Question Answering</i>	
Mohammad Arvan, Hossein Haeri, Natalie Parde and Rebecca Feinstein	487
<i>tt501 at ArchEHR-QA 2026: Few-Shot Prompting with Retrieval-Augmented Generation for Grounded Clinical EHR Question Answering</i>	
Tai Tan Tran	497
<i>HealthNLP_Retrievers at ArchEHR-QA 2026: Cascaded LLM Pipeline for Grounded Clinical Question Answering</i>	
Md Biplob Hosen, Md Alomgeer Hussein, Md Akmol Masud, Omar Faruque, Tera L Reynolds and Lujie Karen Chen	506
<i>Yale-DM-Lab at ArchEHR-QA 2026: Deterministic Grounding and Multi-Pass Evidence Alignment for EHR Question Answering</i>	
Elyas Irankhah and Samah Fodeh	515
<i>Razreshili at ArchEHR-QA 2026: Evidence Alignment via LLM Prompting and Cross-Encoder Fine-tuning</i>	
Arina Zemchyk	524
<i>Neural at ArchEHR-QA 2026: One Method Fits All: Unified Prompt Optimization for Clinical QA over EHRs</i>	
Abrar Majeedi, Viswanatha Reddy Gajjala, Sai Prasanna Teja Reddy Bogireddy and Siddhant Rai	530
<i>OptiMed at ArchEHR-QA 2026: GEPA Prompt Optimization and Multi-Agent Majority Voting for EHR-Grounded Question Answering</i>	
Feras AlMannaa, Talia Tseriotou and Maria Liakata	539
<i>TAMU-NLP at ArchEHR-QA 2026: Grounded Clinical QA with Evidence Identification and Intent-Aware Answer Generation</i>	
Xinqi Su, Rongrong Wang, Sunyang Fu, Hongfang Liu and Ruihong Huang	551

GigitAI at ArchEHR-QA 2026: Prompting Strategies and Constitutional AI for Clinical Question Answering

Saran Krishnasamy and Inez Wihardjo	565
---	-----

Workshop Program

Tuesday, May 12, 2026

09:00–09:10 Opening remarks

09:10–10:10 Session 1: Oral Presentations

09:10–09:30 *FHIRPath-QA: Executable Question Answering over FHIR Electronic Health Records*

Michael Frew, Nishit Bheda and Bryan Tripp

09:30–09:50 *COACH Meets QUORUM: A Framework and Pipeline for Aligning User, Expert, and Developer Perspectives in LLM-Generated Health Counselling*

Yee Man Ng, Bram van Dijk, Pieter Beynen, Otto Boekesteijn, Joris Jansen, Gerard van Oortmerssen, Max J. van Duijn and Marco Spruit

09:50–10:10 *Addressing Domain Shift in Health Coaching Note Analysis through Factorized Synthetic Data Generation*

Michael Tänzer, Iva Bojic, Ashwini Yuvraj Lawate, Andy Hau Yan Ho and Andy Khong

10:10–10:31 Session 2: Virtual Poster Presentations

10:10–10:17 *Scalable Generation of Adult-Oriented Therapeutic Reading Texts for Russian Aphasia Rehabilitation*

Anastasia Kolmogorova, Anastasia Margolina, Alina Telnova and Igor Ilchenko

10:17–10:24 *An Open-Resource Knowledge Augmentation for Biomedical Lay Summarization*

João Pedro Veloso and Evelin Amorim

10:24–10:31 *TabMedQA: From Structured Data to Question-Answer Datasets in Early Clinical Decision-Making*

Gabriel Iturra Bocaz, Petra Galuščáková, Sol Gedde Vedde and Alvaro Fernandez-Quilez

10:31–11:00 Coffee Break and Poster Session

SimpliMED: Automatic Simplification of Cardiology Discharge Reports Using Large Language Models

Lucas Molino-Piñar, Manuel Carlos Díaz Galiano, María-Teresa Martín-Valdivia, Jose Angel Urbano-Moral and Elena Sola-Garcia

Tuesday, May 12, 2026 (continued)

Dutch Metaphor Extraction from Cancer Patients' Interviews and Forum Data Using LLMs and Human in the Loop

Lifeng Han, David Lindevelt, Sander Puts, Erik van Mulligen and Suzan Verberne

Compressed Representations of Patient Records: A Comparative Study of Template-Based and LLM-Based Methods for Clinical Data Summarization and Visualization

Andreas Stöckl, Oliver Krauss and Sophie Bauernfeind

Medical Text Rewriting for Non-Experts: A Guideline-Driven LLM Approach

Mana Kuramoto, Hiroyuki Nagai, Keiko Yamada, Hiroo Ide, Masayo Hayakawa, Tomohiro Nishiyama, Shoko Wakamiya and Eiji Aramaki

Patient-Specific Care Pathway Visualisation for Medical Nursing Staff

Sophie Bauernfeind, Selina Adlberger, Oliver Krauss and Andreas Stöckl

What Makes a Good Doctor Response? A Study on Text-Based Telemedicine

Adrian Cosma, Cosmin Dumitrache and Emilian Radoi

Datasets for a Chatbot for Clinical Trial Search

Yumeng Yang, Ethan Ludmir and Kirk Roberts

HealthTrajectory: Patient Journey Summaries and Visualizations for Patient-Clinician Communication Support

Rohmah Hidayah, Tomohiro Nishiyama, Shoko Wakamiya and Eiji Aramaki

Medical-FLAVORS-AECC: Spanish Oncological Metaphors Dataset

Lucia Pitarch, Jordi Bernad, Sergio-Luis Ojeda-Trueba, Alec Sánchez-Montero, Maxim Ionov, Emma Anglés-Herrero, Ángel Óscar Corona Beomont and Gemma Bel-Enguix

A Synthetic Conversational Dataset for Type 2 Diabetes Management

Stergios Ntanavaras, Maaïke de Boer and Piek T.J.M. Vossen

Italian Medical Term Simplification: From Patient Information Leaflets to Simplified Language Resources

Maria Pia di Buono

Tuesday, May 12, 2026 (continued)

11:00–11:40 Session 3: Oral Presentations

11:00–11:20 *Evaluating Professional Acceptability of LLM-Generated Systematic Review Summaries in Healthcare: Psychiatrists' Perspectives*
Paul Thompson, Artemis Boulogeorgou, Fotini Kaponi, Efstathia Soufleri and Sophia Ananiadou

11:20–11:40 *Reasoning, Contrastive, and In-Context Strategies for Opioid Use Stage Detection on Social Media*
Vinu Ekanayake and Ramakanth Kavuluru

11:40–12:40 Session 4: Shared Task Overviews

11:40–11:55 *FoodBench-QA: Overview of the Shared Task on Grounded Food and Nutrition Question Answering*
Tome Eftimov, Ana Gjorgjevikj, Matej Martinc, Gjorgjina Cenikj, Sašo Džeroski and Barbara Koroušič Seljak

11:55–12:10 *Overview of the CT-DEB'26 Shared Task on Predicting Dosing Errors in Interventional Clinical Trials*
Sohrab Ferdowsi, Félicien Hêche, Anthony Yazdani, Edward Choi, Sara Sansaloni-Pastor and Douglas Teodoro

12:10–12:25 *Overview of the CRF 2026 Shared Task on Clinical Case Report Forms Filling*
Pietro Ferrazzi, Soumitra Ghosh, Alberto Lavelli and Bernardo Magnini

12:25–12:40 *Overview of the ArchEHR-QA 2026 Shared Task on Grounded Question Answering from Electronic Health Records*
Sarvesh Soni and Dina Demner-Fushman

Tuesday, May 12, 2026 (continued)

12:40–13:01 Session 5: Virtual Poster Presentations

12:40–12:47 *Structured Radiology Intelligence: Extracting Structured Data from MRI Reports Using LLMs*
Sushvin Marimuthu, Parameswari Krishnamurthy, Dipti Misra Sharma, Goldwin H, Anu Eapen, Betty Simon and Anuradha Chandramohan

12:47–12:54 *Useful to Whom? A Persona-Driven Evaluation of Knowledge-Adapted Health Question Reformulation via LLM Simulation*
Jooyeon Lee, Luan Huy Pham and Özlem Uzuner

12:54–13:01 *MedGore: An Approach and a Dataset for Identification of Sensitive Medical Images*
Soumya Gayen, Rory Mulcahey, Russell Loane, Dina Demner-Fushman and Deepak Gupta

13:01–14:00 Lunch Break

14:00–14:35 Session 6: Oral Presentations

14:00–14:20 *Diagnostic Reasoning with Large Language Models for a Rare Disease: Case Study of Primary Ciliary Dyskinesia*
Swati Rajwal, Mary Ellen M. Fain, Lokesh Guglani and Abeed Sarker

14:20–14:35 *Beyond One-Size-Fits-All: Multi-Agent Refinement Framework for Persona-Based Biomedical Summarization*
Rohan Charudatt Salvi, Chirag Chawla, Md. Shad Akhtar and Shweta Yadav

14:35–15:05 *Invited Talk: Towards Trustworthy and Efficient Continual Learning for Large Language Models*
Kailai Yang

Tuesday, May 12, 2026 (continued)

15:05–16:01 Session 7: Shared Task Virtual Poster Presentations

15:05–15:12 *Automated Detection of Dosing Errors in Clinical Trial Narratives: A Multi-Modal Feature Engineering Approach with LightGBM*
Mohammad AL-Smadi

15:12–15:19 *CaresAI at CT-DEB'26: Detecting Dosing Errors In Clinical Trials Using Domain-Specific Transformer Embeddings and Classification Models*
Leon Hamnett, Favour Igwezeke, Joseph Itopa Abubakar and Mary Adetutu Adewunmi

15:19–15:26 *CGU-ILALab at FoodBench-QA 2026: Comparing Traditional and LLM-based Approaches for Recipe Nutrient Estimation*
Wei-Chun Chen, Yu-Xuan Chen, I-Fang Chung and Ying-Jia Lin

15:26–15:33 *FCProfiler: Structured and Deterministic Pipeline for Recipe-Level Nutrient Estimation*
Siyoon Kim

15:33–15:40 *DocUA at CRF Filling 2026: LLM StructCore — Schema-Guided Reasoning Condensation and Deterministic Compilation*
Serhii Zabolotnii

15:40–15:47 *GREYC at CRF Filling 2026: Rewrite Before You Extract - Rewriting Clinical Notes for Automated CRF*
Jesus Lovon-Melgarejo, Jérémie Pantin and Gaël Dias

15:47–15:54 *Innov8rs at CRF Filling 2026: An Iterative Multi-LLM Ensemble Pipeline with Dynamic Few-Shot Retrieval and Data-Driven Precision Filtering*
Samminga Sainath Rao, Sumit Mishra and Chanchal Suman

15:54–16:01 *Aurum at CRF Filling 2026: Modular DSPy Extractors with Qwen3-Max for Multilingual CRF Filling*
Vinay Babu Ulli, Jyoti Kumari and Anindita Mondal

16:01–17:00 Coffee Break and Poster session

sebis at CRF Filling 2026: A Two-Stage Local LLM Pipeline for Medical CRF Filling
Katharina Sommer, Tristan Till and Florian Matthes

Polimi at CRF Filling 2026: Prompt-Based Information Extraction from Italian Clinical Notes
Vittorio Torri and Francesca Ieva

Tuesday, May 12, 2026 (continued)

Cohere Labs Community at FoodBench-QA 2026: The Cake Makes the Ingredients

Ravi Ranjan, Roshan Santhosh and Lucien Carroll

HiTZ-IXA at ArchEHR-QA 2026: Evidence Alignment Through Self-Consistency and Prompt Curation in Memory-Constrained Environments

Xabier Irastortza-Urbieta, Maite Oronoz and Alicia Pérez

BIT.UA-AAUBS at ArchEHR-QA 2026: Evaluating Open-Source and Proprietary LLMs via Prompting in Low-Resource QA

Richard A. A. Jonker, Alexander Christiansen, Alexandros Maniatis, Rúben Garrido, Rogério Braunschweiger de Freitas Lima, Roman Jurawetzki and Sérgio Matos

WisPerMed at ArchEHR-QA 2026: Retrieval-Augmented Prompting for Grounded EHR Question Answering

Jan-Henning Büns, Tabea Margareta Grace Pakull, Hendrik Damm, Bohao Chu, Christoph M. Friedrich, Felix Nensa, Elisabeth Livingstone, Peter A. Horn and Norbert Fuhr

sebis at ArchEHR-QA 2026: How Much Can You Do Locally? Evaluating Grounded EHR QA on a Single Notebook

Ibrahim Ebrar Yurt, Fabian Tobias Karl, Tejaswi Chopra and Florian Matthes

MedEvi-NS at ArchEHR-QA 2026: Using Clinical Reasoning Principles to Improve Zero-shot Capabilities of Large Language Models in Evidence Alignment

Mengxuan Sun and Nicolay Rusnachenko

UIC-AIHealth4All at ArchEHR-QA 2026: Answer-First Evidence Grounding for Clinical Question Answering

Mohammad Arvan, Hossein Haeri, Natalie Parde and Rebecca Feinstein

17:00–17:56 Session 8: Shared Task Virtual Poster Presentations

17:00–17:07 *tt501 at ArchEHR-QA 2026: Few-Shot Prompting with Retrieval-Augmented Generation for Grounded Clinical EHR Question Answering*
Tai Tan Tran

17:07–17:14 *HealthNLP_Retrievers at ArchEHR-QA 2026: Cascaded LLM Pipeline for Grounded Clinical Question Answering*
Md Biplob Hosen, Md Alomgeer Hussein, Md Akmol Masud, Omar Faruque, Tera L Reynolds and Lujie Karen Chen

Tuesday, May 12, 2026 (continued)

- 17:14–17:21 *Yale-DM-Lab at ArchEHR-QA 2026: Deterministic Grounding and Multi-Pass Evidence Alignment for EHR Question Answering*
Elyas Irankhah and Samah Fodeh
- 17:21–17:28 *Razreshili at ArchEHR-QA 2026: Evidence Alignment via LLM Prompting and Cross-Encoder Fine-tuning*
Arina Zemchyk
- 17:28–17:35 *Neural at ArchEHR-QA 2026: One Method Fits All: Unified Prompt Optimization for Clinical QA over EHRs*
Abrar Majeedi, Viswanatha Reddy Gajjala, Sai Prasanna Teja Reddy Bogireddy and Siddhant Rai
- 17:35–17:42 *OptiMed at ArchEHR-QA 2026: GEPA Prompt Optimization and Multi-Agent Majority Voting for EHR-Grounded Question Answering*
Feras AlMannaa, Talia Tseriotou and Maria Liakata
- 17:42–17:49 *TAMU-NLP at ArchEHR-QA 2026: Grounded Clinical QA with Evidence Identification and Intent-Aware Answer Generation*
Xinqi Su, Rongrong Wang, Sunyang Fu, Hongfang Liu and Ruihong Huang
- 17:49–17:56 *GigitAI at ArchEHR-QA 2026: Prompting Strategies and Constitutional AI for Clinical Question Answering*
Saran Krishnasamy and Inez Wihardjo
- 17:56–18:00 Closing Remarks**

