

Marek Rei

Email: marek.rei@imperial.ac.uk

Homepage: www.marekrei.com

Experience

Associate Professor in Machine Learning, Imperial College London

October 2023 - current

Researching machine learning methods for planning and reasoning, representation learning, language modelling and natural language understanding. Applying the technologies in the areas of education, healthcare and sustainability. Lecturing Introduction to Machine Learning and Natural Language Processing, some of the largest courses in the department. Supervising 5 PhD students and a number of Master's students on their research projects.

Lecturer in Machine Learning, Imperial College London

October 2019 - October 2023

Researching representation learning, transfer learning and language modelling for natural language understanding.

Visiting Researcher, University of Cambridge

October 2019 - current

Researching neural machine learning algorithms for automated error detection, error correction, exercise creation and language proficiency scoring. Collaborating on the Automated Language Teaching and Assessment (ALTA) project.

Co-founder and Chief Scientific Officer, Transformative AI

October 2016 - October 2023

Researching machine learning algorithms for the early prediction of cardiac arrest in critical care patients. The system is continuously analysing ECG signals in order to detect ventricular tachyarrhythmias before they occur. We have received investor funding, as well as substantial grants from Innovate UK and the Wellcome Trust.

Affiliated Lecturer, University of Cambridge

October 2018 - October 2019

My research is on machine learning and advanced neural network architectures, with a focus on representation learning, transfer learning, language modeling and multi-objective optimization. I teach courses on machine learning and data science to both undergraduate and MPhil students, and I supervise projects for the Advanced Computer Science degree.

Senior Research Associate, University of Cambridge

July 2016 - October 2019

Working on machine learning and natural language processing, structured prediction and multi-task learning. My work contributes to the ALTA project – an industry collaboration with Cambridge Assessment, with the goal of creating innovative new technologies for language learning and teaching. Supervising MPhil and undergraduate students for their course projects, and teaching as part of the MPhil course on Advanced Computer Science.

Research Associate, University of Cambridge

November 2014 - July 2016

Performed research on machine learning algorithms for automated language analysis. Among other things, I developed a general-purpose neural structured prediction framework, and an online representation learning algorithm for improved language modeling.

Visiting Lecturer, University of Tartu

March 2015 - June 2015

As a visiting lecturer, I created and delivered a new course on machine learning and language modeling. The lectures cover topics such as n-gram language modeling, smoothing techniques, language model evaluation, neural network language models, recurrent models, backpropagation, gradient descent, representation learning, and an overview of current state-of-the-art methods for language modeling.

Language Research Engineer, SwiftKey

November 2012 - November 2014

I was a member of the Research team, working on future technologies and prototyping new solutions. One of my main projects was the neural network version of SwiftKey, which was released to very positive reviews. Since then, SwiftKey has been acquired by Microsoft, largely based on the strength of its technology.

Supervision Tutor, University of Cambridge

January 2011 - March 2012

I organized tutoring sessions for undergraduate students in the subject of Information Retrieval. This involved regular tutorials in small groups, and providing them with theoretical and practical assignments.

Public Relations Manager, Board of European Students of Technology

September 2007 - July 2008

Organizing BEST General Assembly 2008 - an international student congress with representatives from 76 universities. Handled relations with the media, companies and universities, and designed various promotional materials. Worked in a core team of 5 people together with over 40 volunteers.

Software Analyst, Webmedia

June 2007 - August 2007

Internship as a software analyst. Worked on enhancing a self-service system for a mobile service provider, designed and performed usability tests and quality control, wrote specifications and created a prototype for a new user interface.

Project manager, Board of European Students of Technology

October 2006 - May 2007

Main organizer of Key to the Future (Võti Tulevikku) 2007, a national career project for graduates in Estonia with 74 participating companies. As the project manager I led a core team of 5 people plus many volunteers, and was involved in every aspect of the project.

Software Developer, Movies.ee

June 2006 - April 2007

Developed a custom-built content management system for managing movies, actors, news, galleries, con-

tests and live cinema timetables.

Software Developer, Geomedia

February 2006 - June 2006

Developed software for conducting large-scale population surveys.

Education

PhD, University of Cambridge

2009 - 2013

My thesis is about minimally supervised dependency-based methods for natural language processing, under the supervision of Prof Ted Briscoe. I worked on topics such as distributional and semantic similarity, speculation detection, entailment detection, and unsupervised parse reranking. I applied various machine learning techniques and developed systems that achieved state-of-the-art performance.

MPhil, University of Cambridge

2008 - 2009

MPhil in Computer Speech, Text and Internet Technologies. Covered topics such as speech recognition, speech synthesis, parsing, discourse, dialogue systems, machine translation and question answering.

BSc, Tallinn University of Technology

2005 - 2008

Bachelor's Degree in Computer Science from the Tallinn University of Technology. Specialised in Network Applications and Logic, graduated *cum laude*.

Membership

I am a member of the following organizations:

Member of the Association for the Advancement of Artificial Intelligence (AAAI)

Member of the Association for Computational Linguistics (ACL)

Member of the Scientific Board at the Estonian Centre of Excellence in Artificial Intelligence (EXAI)

Member of the Special Interest Group on Representation Learning (SigREP)

Member of the Special Interest Group on NLP and Education (SigEDU)

Member of the AI4Health Training Board at Imperial College London

Member of the AI Network at Imperial College London

Member of the Leonardo Centre for Sustainable Business

Organization

I helped organize the following events:

Area Chair for the 19th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2026)

Area Chair for the International Joint Conference on Natural Language Processing & Asia-Pacific Chapter of the Association for Computational Linguistics (AAACL 2025)

Area Chair for the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP 2025)

Senior Area Chair for the 63rd Annual Meeting of the Association for Computational Linguistics (ACL 2025)

Area Chair for the 31st International Conference on Computational Linguistics (COLING 2025)

Area Chair for the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024)

Area Chair for the 18th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2024)

Member of the Estonian Science Agency (ETAG) steering committee for ICT target evaluation (2023)

Area Chair for the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023)

Area Chair for the 61st Annual Meeting of the Association for Computational Linguistics (ACL 2023)

Senior Area Chair for the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022)

Organizer of the 10th Conference on Automated Knowledge Base Construction (AKBC 2022)

Area Chair for the 2021 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2021)

Organizer of the 2020 Workshop on Representation Learning for NLP (RepL4NLP 2020)

Area Chair for the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)

Area Chair for the 57th Annual Meeting of the Association for Computational Linguistics (ACL 2019)

Area Chair for the Eighth Joint Conference on Lexical and Computational Semantics (*SEM 2019)

Organizer of the 2019 Workshop on Representation Learning for NLP (RepL4NLP 2019)

Organizer of the Workshop on Machine Translation and Natural Language Processing at Makerere University 2018

Awards, Grants and Scholarships

UKRI AI Hub in Generative Models. Co-Investigator. EPSRC. 2024-2029.

UKRI AI Centre for Doctoral Training in Digital Healthcare (AI4Health CDT). Co-investigator. EPSRC. 2024-2033.

Gemma Academic Program GCP Credit Award, 2024

EU-Horizon-RIA grant for Combatting Diet Related Non-Communicable Disease through Enhanced Surveillance (CoDiet). Co-investigator, 2023-2027.

Young IT Scientist of the Year, Presidential Award. Estonia, 2021

Nominated for the President's Award for Excellence in Education, by the Imperial College Department of Computing, 2021

King's College Junior Research Fellowship, 2017-2021

Cambridge-Africa Alborada Research award, 2017-2018

Best Paper Award for Research Inspired by Human Language Learning, CoNLL 2018

Language Sciences Research Incubator Fund, 2017

EPSRC Doctoral Training award, 2009-2012

Cambridge European Trust award, 2009-2012

Computer Laboratory Local Industry Fund, 2009-2012

Churchill College Studentship, 2009-2012

EPSRC Collaborative Training award, 2008

Cambridge European Trust award, 2008

Estonian Academy of Sciences research award, 2008

AS Elion Scholarship, 2007

Lecturing

The courses I have created and lectured:

Machine Learning, Imperial College London, 2026 Spring

Natural Language Processing, Imperial College London, 2026 Spring

Introduction to Machine Learning, Imperial College London, 2025 Autumn

Introduction to Machine Learning, Imperial College London, 2025 Spring

Natural Language Processing, Imperial College London, 2025 Spring

Introduction to Machine Learning, Imperial College London, 2024 Autumn

Introduction to Machine Learning, Imperial College London, 2024 Spring

Natural Language Processing, Imperial College London, 2024 Spring

Introduction to Machine Learning, Imperial College London, 2023 Autumn

Introduction to Machine Learning, Imperial College London, 2023 Spring

Natural Language Processing, Imperial College London, 2023 Spring

Introduction to Machine Learning, Imperial College London, 2022 Autumn

Introduction to Machine Learning, Imperial College London, 2022 Spring

Introduction to Machine Learning, Imperial College London, 2021 Autumn

Introduction to Machine Learning, Imperial College London, 2021 Spring
Introduction to Machine Learning, Imperial College London, 2020 Autumn
Introduction to Machine Learning, Imperial College London, 2020 Spring
Interpreting the Black Box: Explainable Neural Network Models, University of Cambridge, 2019
Data Science: Principles and Practice, University of Cambridge, 2018
Constructing and Evaluating Word Embeddings, University of Cambridge, 2017
Constructing and Evaluating Word Embeddings, University of Cambridge, 2016
Machine Learning for Language Modelling, University of Tartu, 2015

PhD Students

Nightingale AI: A Foundation Model from Millions of Electronic Health Records

Joshua Placidi

PhD, Imperial College London. 2025 - current.

Co-supervised with Prof Aldo Faisal.

Unsupervised Text Classification with Retrieval Augmented Models

Tyler Bonnet

PhD, Imperial College London. 2024 - current.

External Tool Use for Complex Problem Solving with Language Models

Lisa Alazraki

PhD, Imperial College London. 2023 - current.

Planning and Reasoning with Large Language Models

Nikolai Rozanov

PhD, Imperial College London. 2023 - current.

Robust Neural Models for Generalising to Unseen NLI Tasks

Joe Stacey

PhD, Imperial College London. 2020 - current.

Noise Reduction for Robust Language Processing Models

Zhenhao Li

PhD, Imperial College London, 2019 - 2025.

Co-supervised with Prof Lucia Specia.

Generative Neural-led Visual Dialog

Nihir Vedda

PhD, Imperial College London, 2019 - 2024.

Co-supervised with Prof Lucia Specia.

Masters Students

Detect-then-Correct: Advancing the Error Correction Abilities of Language Models

Stefan Creasta

MSc in Computing (AI & ML), Imperial College London, 2025.

Fine-tuning with RAG for Improving LLM Learning of New Skills

Humaid Ibrahim

*MSc in Artificial Intelligence, Imperial College London, 2025.***Multimodal State Space Models for Modelling Electronic Health Records**

Reynard Adha Ryanda

*MSc in Advanced Computing, Imperial College London, 2025.***Optimal Data Mixing Strategies for Language Model Pre-Training**

Kevin Zhou

*MSc in Advanced Computing, Imperial College London, 2025.***Improving Robustness for Black-box Large Language Models**

Aran Ubhi

*MSc in Computing (AI & ML), Imperial College London, 2024.***Temporal Multi-modal Modelling of Document Sequences**

Cindy Huang

*MSc in Computing (AI & ML), Imperial College London, 2024.***Improving the Robustness of Large Language Models to Prompt Perturbations**

Aryan Agrawal

*MSc in Advanced Computing, Imperial College London, 2024.***Predictive Temporal Modelling of Document Sequences**

Mireia Hernandez Caralt

*MSc in Artificial Intelligence, Imperial College London, 2023.***Conditioning Language Generation on Future Downstream Impact**

Jorge Gallego Feliciano

*MSc in Advanced Computing, Imperial College London, 2023.***Editing Pre-trained Language Models for Domain Adaptation**

Joon Jeon

*MSc in Computing (AI & ML), Imperial College London, 2023.***Removing model biases by directly editing model parameters**

Joy Liu

*MSc in Computing (AI & ML), Imperial College London, 2023.***Pre-trained Generative Language Models for Unsupervised Grammatical Error Correction**

Yulong Lin

*Computer Science Tripos (Part III), University of Cambridge, 2023.***Being Right for the Right Reasons: Using Auxiliary Datasets During Training to Create More Robust NLP Models**

Marine De Backer

*MSc in Advanced Computing, Imperial College London, 2022.***Capturing the Gap in Pre-trained Language Models**

Hangyu Li

*MSc in Advanced Computing, Imperial College London, 2022.***Entity Relation Extraction within the Telecommunications Domain**

Alba Espinosa Rastoll

MSc in Artificial Intelligence, Imperial College London, 2022.

ICD Coding with Transformers

Clarence Ng

*MSc in Computing (AI & ML), Imperial College London, 2022.***Finding the Needle in a Haystack: Zero-shot Rationale Extraction for Long Text Classifiers**

Kamil Bujel

*MEng in Mathematics and Computer Science, Imperial College London, 2022.***Hierarchical Models for Grammatical Error Correction**

Stuart Mesham

*MPhil in Advanced Computer Science, University of Cambridge, 2022.**Co-supervision with Christopher Bryant and Zheng Yuan.***Multi-level Optimization in Transformer Models**

Tom Taylor

*MSc in Computing (AI & ML), Imperial College London, 2021.***Sustainability Initiative Detection in Company Reports**

Mircea-Dan Hirlea

*MSc in Artificial Intelligence, Imperial College London, 2021.***Pre-training Contextual Language Decoders**

Jordan Clive

*MSc in Computing (AI & ML), Imperial College London, 2021.***Temporal Modelling in Electronic Health Records for Diagnosis Prediction**

Ziyue Xu

*MSc in Computing (AI & ML), Imperial College London, 2021.***Capturing Outliers with Neural Episodic Memory**

Michael Tanzer

*MSc in Computing (AI & ML), Imperial College London, 2020.***Twice-cooked Contextual Embeddings**

Sun Whang

*MSc in Artificial Intelligence, Imperial College London, 2020.***Unsupervised Error Detection**

Simone Primarosa

*MPhil in Advanced Computer Science, University of Cambridge, 2020.**Co-supervising together with Dr Chris Bryant.***Supervised Interpretability for Text Classification**

Kamil Bujel

*MEng in Mathematics and Computer Science, Imperial College London, 2020.***Joint modelling of compositionality, metaphor and emotion**

Verna Dankers

*Master's dissertation, University of Amsterdam, 2019.**Supervising in collaboration with Dr Ekaterina Shutova and Dr Martha Lewis.***Joint Text Classification on Multiple Levels with Multiple Labels**

Miruna Pislari

MPhil in Advanced Computer Science, University of Cambridge, 2019.

Few-Shot Learning in Distributional Semantics

Jeroen Van Hautte

*MPhil in Advanced Computer Science, University of Cambridge, 2019.**Supervising in collaboration with Dr Guy Emerson.***Language Modelling for Neural Error Detection**

Samuel Bell

*MPhil in Advanced Computer Science, University of Cambridge, 2018.***Neural Architectures for Visual Question Answering**

Joshua Wong

*Computer Science Tripos, University of Cambridge, 2018.***Supervised Attention for Neural Error Correction**

Gladys Tien

*MPhil in Advanced Computer Science, University of Cambridge, 2018.**Supervising in collaboration with Dr Zheng Yuan.***Domain Adaptation for Neural Named Entity Recognition**

Stefan Hosein

*MPhil in Advanced Computer Science, University of Cambridge, 2017.***Hybrid Language Modeling for Text Prediction**

Devan Kuleindiren

*Computer Science Tripos, University of Cambridge, 2017.***Convolutional Neural Networks for Automated Essay Assessment**

Youmna Farag

*MPhil in Advanced Computer Science, University of Cambridge, 2016.***Predicting Word Difficulty for Lexical Acquisition**

Shushan Arakelyan

MPhil in Advanced Computer Science, University of Cambridge, 2016.

PhD Examinations

From Graphs to Truth: Towards efficient knowledge graph fusion for factual verification

Weihang Zhang

*Imperial College London, 2025.***On Robustness for Natural Language Processing**

Emanuele La Malfa

*University of Oxford, 2023.***A Personalised Approach to Lexical Complexity**

Sian Gooding

University of Cambridge, 2023.

Publications

Reverse Engineering Human Preferences with Reinforcement Learning

Lisa Alazraki, Tan Yi-Chern, Jon Ander Campos, Maximilian Mozes, Marek Rei, Max Bartolo

The Thirty-Ninth Annual Conference on Neural Information Processing Systems (NeurIPS 2025). San Diego,

United States, 2025.

****Spotlight paper****

No Need for Explanations: LLMs can implicitly learn from mistakes in-context

Lisa Alazraki, Maximilian Mozes, Jon Ander Campos, Yi Chern Tan, Marek Rei, Max Bartolo
In Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP 2025). Suzhou, China, 2025.

Fine-tuning with RAG for Improving LLM Learning of New Skills

Humaid Ibrahim, Nikolai Rozanov, Marek Rei
under review. 2025.

The “Question Neighbourhood” Approach for Systematic Evaluation of Code-Generating LLMs

Shahin Honarvar, Marek Rei, Alastair Donaldson
IEEE Transactions on Software Engineering. 2025.

Predictive Multimodal Modeling of Diagnoses and Treatments in EHR

Cindy Shih-Ting Huang, Clarence Boon Liang Ng, Marek Rei
Proceedings of the 1st Workshop on Multimodal Large Language Models (MLLMs) in Clinical Practice (MICCAI MLLMs workshop 2025). Daejeon, Republic of Korea, 2025.

AgentCoMa: A Compositional Benchmark Mixing Commonsense and Mathematical Reasoning in Real-World Scenarios

Lisa Alazraki, Lihu Chen, Ana Brassard, Joe Stacey, Hossein A. Rahmani, Marek Rei
under review. ArXiv, 2025.

How to Improve the Robustness of Closed-Source Models on NLI

J Stacey, L Alazraki, A Ubhi, B Ermis, A Mueller, M Rei
under review. ArXiv, 2025.

Enhancing LLM Robustness to Perturbed Instructions: An Empirical Study

Aryan Agrawal, Lisa Alazraki, Shahin Honarvar, Marek Rei
ICLR 2025 Workshop - Building Trust in LLMs and LLM Applications: From Guardrails to Explainability to Regulation. Singapore, 2025.

Assessing the effectiveness of interdependent corporate sustainability choices

Simone Cenci, Matteo Burato, Marek Rei, Maurizio Zollo
npj Climate Action (Nature). 2025.

KU AIGEN ICL EDI@ BC8 Track 3: Advancing Phenotype Named Entity Recognition and Normalization for Dysmorphology Physical Examination Reports

Hajung Kim, Chanhwi Kim, Jiwoong Sohn, Tim Beck, Marek Rei, Sunkyu Kim, T Ian Simpson, Joram M Posma, Antoine Lain, Mujeen Sung, Jaewoo Kang
ArXiv, 2025.

Automatic genetic phenotype normalization from dysmorphology physical examinations: an overview of the BioCreative VIII—Task 3 competition

Davy Weissenbacher, Xinwei Zhao, Jessica RC Priestley, Katherine M Szigety, Sarah F Schmidt, Karen O'Connor, Ekin Soysal, Kirk Roberts, Jiewei Qi, Ling Luo, Zhihao Yang, Hongfei Lin, Hajung Kim, Chanhwi Kim, Jiwoong Sohn, Tim Beck, Marek Rei, Sunkyu Kim, T Ian Simpson, Joram M Posma, Antoine Lain, Mujeen Sung, Jaewoo Kang, Areej Alhassan, Viktor Schlegel, Monira Aloud, Riza Batista-Navarro, Goran Nenadic, Ying-Jia Lin, Zhi-Quan Feng, Hung-Yu Kao, Ian M Campbell, Graciela Gonzalez-Hernandez
Database. 2025.

High-precision Biomedical Text Corpora for Multi-Entity Recognition: A CoDiet study

Antoine D Lain, Stephanie Go, Alisha Mahmud, Shruti Rajendra, Ainara Cano San Jose, Katerina Loupasaki, Georgios Theodoridis, Maider Bizkarguenaga Uribearte, Yajie Gu, Olga Deda, Ricardo Diogo Alves Conde, Nieves Embade, Angela de Diego Rodriguez, Nerea Burguera, Danai Rossiou, Ruben Gil Redondo, Domniki Gallou, Itziar Tueros, Rakesh Velmurugan, Vasiliki Gkanali, Mercedes Caro Burgos, Petros Pousinis, George Alektoridis, Sara Arranz, Nasos Nikolopoulos, Xingchen Yan, Rebeca Fernandez Carrion, Thomas Rowlands, Donghee Choi, Marek Rei, Chris Cave-Ayland, Adrian D Alessandro, CoDiet consortium, Tim Beck, Joram M Posma
bioRxiv, 2025.

DiffuseDef: Improved Robustness to Adversarial Attacks via Iterative Denoising

Zhenhao Li, Huichi Zhou, Marek Rei, Lucia Specia
In Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL 2025). Vienna, Austria, 2025.

Meta-Reasoning Improves Tool Use in Large Language Models

Lisa Alazraki, Marek Rei
Findings of the Association for Computational Linguistics: NAACL 2025. Albuquerque, New Mexico, 2025.

SoK: Membership Inference Attacks on LLMs are Rushing Nowhere (and How to Fix It)

Matthieu Meeus, Igor Shilov, Shubham Jain, Manuel Faysse, Marek Rei, Yves-Alexandre de Montjoye
IEEE Conference on Secure and Trustworthy Machine Learning (SaTML 2025). Copenhagen, Denmark, 2025.
Best paper award

StateAct: Enhancing LLM Base Agents via Self-prompting and State-tracking

Nikolai Rozanov, Marek Rei
In Proceedings of the Workshop for Research on Agent Language Models (REALM 2025). Vienna, Austria, 2025.

Atomic Inference for NLI with Generated Facts as Atoms

Joe Stacey, Pasquale Minervini, Haim Dubossarsky, Oana-Maria Camburu, Marek Rei
In Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024). Miami, USA, 2024.

Distilling Robustness into Natural Language Inference Models with Domain-Targeted Augmentation

Joe Stacey, Marek Rei
In Findings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024). Bangkok, Thailand, 2024.

Continuous Predictive Modeling of Clinical Notes and ICD Codes in Patient Health Records

Mireia Hernandez Caralt, Clarence Boon Liang Ng, Marek Rei
In Proceedings of the Biomedical Natural Language Processing Workshop (BioNLP 2024). Bangkok, Thailand, 2024.

Predicting cell type-specific epigenomic profiles accounting for distal genetic effects

Alan E Murphy, William Beardall, Marek Rei, Mike Phuycharoen, Nathan G Skene
Nature Communications. 2024.

Prompting open-source and commercial language models for grammatical error correction of English learner text

Christopher Davis, Andrew Caines, Øistein Andersen, Shiva Taslimipoor, Helen Yannakoudakis, Zheng Yuan, Christopher Bryant, Marek Rei, Paula Buttery
In Findings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024). Bangkok, Thailand, 2024.

Did the Neurons Read your Book? Document-level Membership Inference for Large Language Models

Matthieu Meeus, Shubham Jain, Marek Rei, Yves-Alexandre de Montjoye
The 33rd USENIX Security Symposium (2024). Philadelphia, PA, USA, 2024.

The alignment of companies' sustainability behavior and emissions with global climate targets

Simone Cenci, Matteo Burato, Marek Rei, Maurizio Zollo
Nature Communications, 2023.

When and Why Does Bias Mitigation Work?

Abhilasha Ravichander, Joe Stacey, Marek Rei
In Findings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023). Singapore, 2023.

Competitive Pressure and Emission Reduction: Unravelling the Link

Simone Cenci, Hossein Asgharian, Lu Liu, Marek Rei, Maurizio Zollo
SSRN, 2023.

Climbing empirical fitness landscapes requires companies to manage interconnected sustainability choices

Simone Cenci, Matteo Burato, Marek Rei, Maurizio Zollo
SSRN. 2023.

Finding the Needle in a Haystack: Unsupervised Rationale Extraction from Long Text Classifiers

Kamil Bujel, Andrew Caines, Helen Yannakoudakis, Marek Rei
ArXiv, 2023.

Does competitive pressure drive effective corporate environmental actions?

Simone Cenci, Hossein Asgharian, Lu Liu, Marek Rei, Maurizio Zollo
SSRN. 2023.

On the application of Large Language Models for language teaching and assessment technology

Andrew Caines, Luca Benedetto, Shiva Taslimipoor, Christopher Davis, Yuan Gao, Oeistein Andersen, Zheng Yuan, Mark Elliott, Russell Moore, Christopher Bryant, Marek Rei, Helen Yannakoudakis, Andrew Mullooly, Diane Nicholls, Paula Buttery
In Proceedings of the AIED 2023 Workshop on Empowering Education with LLMs (AIED LLM 2023). Tokyo, Japan, 2023.

Modelling Temporal Document Sequences for Clinical ICD Coding

Clarence Ng, Diogo Santos, Marek Rei
In Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2023). Dubrovnik, Croatia, 2023.

An Extended Sequence Tagging Vocabulary for Grammatical Error Correction

Stuart Mesham, Christopher Bryant, Marek Rei, Zheng Yuan
In Findings of the 17th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2023). Dubrovnik, Croatia, 2023.

Multimodal Conversation Modelling for Topic Derailment Detection

Zhenhao Li, Marek Rei, Lucia Specia
In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022). Abu Dhabi, United Arab Emirates, 2022.

Logical Reasoning with Span-Level Predictions for Interpretable and Robust NLI Models

Joe Stacey, Pasquale Minervini, Haim Dubossarsky, Marek Rei

In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022). Abu Dhabi, United Arab Emirates, 2022.

Control Prefixes for Parameter-Efficient Text Generation

Jordan Clive, Kris Cao, Marek Rei

In Proceedings of the Second workshop on Generation, Evaluation & Metrics (GEM 2022). Abu Dhabi, United Arab Emirates, 2022.

An Analysis of Corporate Sustainability Behaviour Through the Lens of Empirical Fitness Landscapes

Simone Cenci, Marek Rei, Maurizio Zollo

SSRN pre-print under review, 2022.

Supervising Model Attention with Human Explanations for Robust Natural Language Inference

Joe Stacey, Yonatan Belinkov, Marek Rei

In Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2022). Virtual Conference, 2022.

Business sustainability behaviour and alignment with climate targets

Simone Cenci, Matteo Burato, Marek Rei, Maurizio Zollo

Research Square pre-print under review, 2022.

Memorisation versus Generalisation in Pre-trained Language Models

Michael Tänzler, Sebastian Ruder, Marek Rei

In Proceedings of the 60th annual meeting of the Association for Computational Linguistics (ACL 2022). Dublin, Ireland, 2022.

Probing for targeted syntactic knowledge through grammatical error detection

Christopher Davis, Christopher Bryant, Andrew Caines, Marek Rei, Paula Buttery

In Proceedings of the 26th Conference on Computational Natural Language Learning (CoNLL 2022). Abu Dhabi, United Arab Emirates, 2022.

Guiding Visual Question Generation

Nihir Vedd, Zixu Wang, Marek Rei, Yishu Miao, Lucia Specia

In Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL-HLT 2022). Seattle, Washington, USA, 2022.

Visual Cues and Error Correction for Translation Robustness

Zhenhao Li, Marek Rei, Lucia Specia

In Findings of the Association for Computational Linguistics: EMNLP 2021.

GiBERT: Enhancing BERT with Linguistic Information using a Lightweight Gated Injection Method

Nicole Peinelt, Marek Rei, Maria Liakata

In Findings of the Association for Computational Linguistics: EMNLP 2021.

Contextual Sentence Classification: Detecting Sustainability Initiatives in Company Reports

Dan Hirlea, Christopher Bryant, Maurizio Zollo, Marek Rei

ArXiv, 2021.

Zero-shot Sequence Labeling for Transformer-based Sentence Classifiers

Kamil Bujel, Helen Yannakoudakis, Marek Rei

In Proceedings of the 6th Workshop on Representation Learning for NLP (RepL4NLP 2021). Virtual Conference, 2021.

How Metaphors Impact Political Discourse: A Large-Scale Topic-Agnostic Study Using Neural Metaphor Detection

Vinodkumar Prabhakaran, Marek Rei, Ekaterina Shutova

*In Proceedings of the 15th AAAI International Conference on Web and Social Media (ICWSM 2021). Atlanta, USA, 2021.***Grammatical error detection in transcriptions of spoken English**

Andrew Caines, Christian Bentz, Kate Knill, Marek Rei, Paula Buttery

*In Proceedings of the 28th International Conference on Computational Linguistics (COLING 2020). Virtual Conference, 2020.***Verbal Multiword Expressions for Identification of Metaphor**

Omid Rohanian, Marek Rei, Shiva Taslimipour, Le An Ha

*In Proceedings of the 58th annual meeting of the Association for Computational Linguistics (ACL 2020). Seattle, USA, 2020.***Grammatical Error Correction in Low Error Density Domains: A New Benchmark and Analyses**

Simon Flachs, Ophélie Lacroix, Helen Yannakoudakis, Marek Rei, Anders Søgaard

*In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020). Virtual Conference, 2020.***Seeing Both the Forest and the Trees: Multi-head Attention for Joint Classification on Different Compositional Levels**

Miruna Pislari, Marek Rei

*In Proceedings of the 28th International Conference on Computational Linguistics (COLING 2020). Virtual Conference, 2020.***Multidirectional Associative Optimization of Function-Specific Word Representations**

Daniela Gerz, Ivan Vulić, Marek Rei, Roi Reichart, Anna Korhonen

*In Proceedings of the 58th annual meeting of the Association for Computational Linguistics (ACL 2020). Seattle, USA, 2020.***Semi-Supervised Bootstrapping of Dialogue State Trackers for Task-Oriented Modelling**

Bo-Hsiang Tseng, Marek Rei, Paweł Budzianowski, Richard Turner, Bill Byrne, Anna Korhonen

*In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP 2019). Hong Kong, China, 2019.***Modelling the interplay of metaphor and emotion through multitask learning**

Verna Dankers, Marek Rei, Martha Lewis, Ekaterina Shutova

*In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP 2019). Hong Kong, China, 2019.***Neural and FST-based approaches to grammatical error correction**

Zheng Yuan, Felix Stahlberg, Marek Rei, Bill Byrne, Helen Yannakoudakis

*In Proceedings of the 14th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2019). Florence, Italy, 2019.***Jointly Learning to Label Sentences and Tokens**

Marek Rei, Anders Søgaard

*In Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019). Honolulu, USA, 2019.***CAMsterdam at SemEval-2019 Task 6: Neural and graph-based feature extraction for the identification of offensive tweets**

Guy Aglionby, Christopher Davis, Pushkar Mishra, Andrew Caines, Helen Yannakoudakis, Marek Rei,

Ekaterina Shutova, Paula Buttery

In Proceedings of the International Workshop on Semantic Evaluation 2019 (SemEval 2019). Minneapolis, USA, 2019.

A Simple and Robust Approach to Detecting Subject-Verb Agreement Errors

Simon Flachs, Ophélie Lacroix, Marek Rei, Helen Yannakoudakis, Anders Søgaard

In Proceedings of the 17th Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2019). Minneapolis, USA, 2019.

Bad Form: Comparing Context-Based and Form-Based Few-Shot Learning in Distributional Semantic Models

Jeroen Van Hautte, Guy Emerson, Marek Rei

In Proceedings of the Second Workshop on Deep Learning for Low-Resource NLP (DeepLo 2019). Hong Kong, China, 2019.

Context is Key: Grammatical Error Detection with Contextual Word Representations

Samuel Bell, Helen Yannakoudakis, Marek Rei

In Proceedings of the 14th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2019). Florence, Italy, 2019.

Variable Typing: Assigning Meaning to Variables in Mathematical Text

Yiannos Stathopoulos, Simon Baker, Marek Rei, Simone Teufel

In Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2018). New Orleans, United States, 2018.

Scoring Lexical Entailment with a Supervised Directional Similarity Network

Marek Rei, Daniela Gerz, Ivan Vulić

In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018). Melbourne, Australia, 2018.

Zero-shot Sequence Labeling: Transferring Knowledge from Sentences to Tokens

Marek Rei, Anders Søgaard

In Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2018). New Orleans, United States, 2018.

Neural Multi-task Learning in Automated Assessment

Ronan Cummins, Marek Rei

arXiv:1801.06830, 2018.

Advance Prediction of Ventricular Tachyarrhythmias using Patient Metadata and Multi-Task Networks

Marek Rei, Josh Oppenheimer, Marek Sirendi

In Proceedings of the NeurIPS Workshop on Machine Learning for Health (ML4H 2018). Montreal, Canada, 2018.

Sequence classification with human attention

Maria Barrett, Joachim Bingel, Nora Hollenstein, Marek Rei, Anders Søgaard

In Proceedings of the SIGNLL Conference on Computational Natural Language Learning (CoNLL 2018). Brussels, Belgium, 2018.

****Special award for best paper****

Grasping the Finer Point: A Supervised Similarity Network for Metaphor Detection

Marek Rei, Luana Bulat, Douwe Kiela, Ekaterina Shutova

In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP-2017). Copenhagen, Denmark, 2017.

An Error-Oriented Approach to Word Embedding Pre-Training

Youmna Farag, Marek Rei, Ted Briscoe

*In Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications (BEA-2017). Copenhagen, Denmark, 2017.***Detecting Off-topic Responses to Visual Prompts**

Marek Rei

*In Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications (BEA-2017). Copenhagen, Denmark, 2017.***Auxiliary Objectives for Neural Error Detection Models**

Marek Rei, Helen Yannakoudakis

*In Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications (BEA-2017). Copenhagen, Denmark, 2017.***Artificial Error Generation with Machine Translation and Syntactic Patterns**

Marek Rei, Mariano Felice, Zheng Yuan, Ted Briscoe

*In Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications (BEA-2017). Copenhagen, Denmark, 2017.***Semi-supervised Multitask Learning for Sequence Labeling**

Marek Rei

*In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL-2017). Vancouver, Canada, 2017.***Neural Sequence-Labeling Models for Grammatical Error Correction**

Helen Yannakoudakis, Marek Rei, Øistein E. Andersen, Zheng Yuan

*In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP-2017). Copenhagen, Denmark, 2017.***Attending to characters in neural sequence labeling models**

Marek Rei, Gamal K.O. Crichton, Sampo Pyysalo

*In Proceedings of the 26th International Conference on Computational Linguistics (COLING-2016). Osaka, Japan, 2016.***Compositional Sequence Labeling Models for Error Detection in Learner Writing**

Marek Rei, Helen Yannakoudakis

*In Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL-2016). Berlin, Germany, 2016.***Automatic Text Scoring Using Neural Networks**

Dimitrios Alikaniotis, Helen Yannakoudakis, Marek Rei

*In Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL-2016). Berlin, Germany, 2016.***A Joint Model for Word Embedding and Word Morphology**

Kris Cao, Marek Rei

*In Proceedings of the 1st Workshop on Representation Learning for NLP (RepL4NLP-2016). Berlin, Germany, 2016.***Sentence Similarity Measures for Fine-Grained Estimation of Topical Relevance in Learner Essays**

Marek Rei, Ronan Cummins

In Proceedings of the 11th Workshop on Innovative Use of NLP for Building Educational Applications (BEA). San Diego, United States, 2016.

Online Representation Learning in Recurrent Neural Language Models

Marek Rei

*In Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP). Lisbon, Portugal, 2015.***Looking for hyponyms in vector space**

Marek Rei, Ted Briscoe

*In Proceedings of the Eighteenth Conference on Computational Natural Language Learning (CoNLL-14). Baltimore, Maryland, United States, 2014.***Parser lexicalisation through self-learning**

Marek Rei, Ted Briscoe

*In Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2013). Atlanta, United States, 2013.***Unsupervised Entailment Detection between Dependency Graph Fragments**

Marek Rei, Ted Briscoe

*In Proceedings of the 2011 Workshop on Biomedical Natural Language Processing (BioNLP-11). Portland, United States, 2011.***Intelligent Information Access from Scientific Papers**

Ted Briscoe, Karl Harrison, Andrew Naish-Guzman, Andy Parker, Marek Rei, Advait Siddharthan, David Sinclair, Mark Slater, Rebecca Watson

*Current Challenges in Patent Information Retrieval, edited by Mihai Lupu, Katja Mayer, John Tait and Anthony J. Trippe. Springer, Dordrecht, 2011.***Combining Manual Rules and Supervised Learning for Hedge Cue and Scope Detection**

Marek Rei, Ted Briscoe

The 14th Conference on Natural Language Learning (CoNLL-10). Uppsala, Sweden, 2010.

Note 1: In the areas of ML and NLP, most top-level publishing happens through conferences instead of journals. ML/NLP conferences only publish full-length peer-reviewed papers.

Note 2: Between 2012-2014 I was working as a researcher in SwiftKey. Research from that period was not published due to confidentiality restrictions, but a section of it has been made available as a patent.

Book Chapters

Intelligent Information Access from Scientific Papers

Ted Briscoe, Karl Harrison, Andrew Naish-Guzman, Andy Parker, Marek Rei, Advait Siddharthan, David Sinclair, Mark Slater and Rebecca Watson

Current Challenges in Patent Information Retrieval, edited by Mihai Lupu, Katja Mayer, John Tait and Anthony J. Trippe, 2011.

Patents

Patent applications for which I am a named inventor:

WO2024165844A1 Methods and apparatus for classifying a telecommunication call

Appella Ai Limited

Machine learning methods for classifying a telecommunication call, based on my collaboration with

Appella Ai / Locai Labs.
Filed 2024. Published 2024.

US20210353166A1 Analysis of cardiac data

Transformative AI Ltd

Machine learning system for predicting cardiac events in patients and using it to guide their medical care, based on my research for Transformative AI.

Filed 2019. Published 2021.

US11497430B2 Analysis of cardiac data

Transformative AI Ltd

Machine learning system for predicting cardiac events in patients and using it to guide their medical care, based on my research for Transformative AI.

Filed 2018. Published 2022. Granted 2022.

US11797822B2 Neural network having input and hidden layers of equal units

Microsoft Technology Licensing, Llc

Customised neural network structure for increasing the efficiency of language models, based on my research in SwiftKey.

Filed 2016. Published 2023. Granted 2023.

Theses

Minimally Supervised Dependency-based Methods for Natural Language Processing

PhD thesis, University of Cambridge

Supervised by Ted Briscoe

Cambridge, United Kingdom, 2013

Adaptive Interactive Information Extraction

MPhil thesis, University of Cambridge

Supervised by Ted Briscoe

Cambridge, United Kingdom, 2009

Audio-visual Speech Synthesis

Bachelor's thesis, Tallinn University of Technology

Supervised by Einar Meister

Tallinn, Estonia, 2008

Invited Talks

AI for Language Teaching and Assessment

Royal Society Yusuf Hamied Workshop for India and the UK. November 2025

World Models

Pactum AI, coding agents workshop. June 2025

From ML to AI Assistants

LHV Bank. January 2025

Planning and Reasoning with Language Models

Estonian Academy of Sciences. January 2025

What Makes AI Tick: A Look Inside Large Language Models*Baltic FinReg Summit. June 2024***Sequential Modelling of Patient Records***Speinshart Symposium on Large Health Models. March 2024***Interpretable Architectures and Guided Attention for Neural Language Models***University of Oxford. November 2023***Encoders: The Art of Packing Text into Vectors***ESSCaSS Tartu Summer School. August 2022***Memorisation versus Generalisation in Language Models***University of Cambridge, UK. May 2022***Memorisation versus Generalisation in Language Models***Polish Natural Language Processing Group, UK. Nov 2021***Multidirectional Associative Optimization of Function-Specific Word Representations***University of Cambridge, UK. June 2020***Inductive Transfer Learning for Neural Models***Ajou University, South Korea. July 2019***Application of Deep Learning in NLP***RANLP Summer School on Deep Learning in NLP, Bulgaria. August 2019***How to Pay Attention: Learning to Transfer Knowledge between Sentences and Tokens***University of Cambridge, UK. May 2019***Inductive Transfer Learning for Neural Models of Natural Language***Imperial College London, UK. March 2019***Semi-supervised and Zero-shot Methods for Sequence Labeling***TU Darmstadt, Germany. November 2018***Language Modelling and Machine Translation***University of Makerere, Uganda. June 2018***Human Interpretability of Machine Learning Models***University of Tartu, Estonia. May 2018***Understanding and Assessing Language with Neural Network Models***Grammarly, Ukraine. March 2018***Neural Architectures for Sequence Labelling***Grammarly, Ukraine. March 2018***Auxiliary Objectives for Neural Sequence Labelling Models***University of Copenhagen, Denmark. November 2017***Sequence Labeling Models for Error Detection in Learner Writing***Machine Learning for Further Language Learning, summer school in Chania, Greece. July 2017***Neural Architectures for Sequence Labelling***Apple Cambridge. June 2017*

Neural Architectures for Sequence Labelling*University College London. May 2017***Neural Architectures for Sequence Labelling***University of Cambridge. May 2017***Character-based and Multi-task Learning for Neural Sequence Labeling***University of Sheffield. March 2017*

Program Committee

I am part of the program committee for the following conferences and workshops:

The Conference on Neural Information Processing Systems (NeurIPS 2025)

The 2025 International Conference on Learning Representations (ICLR 2025)

The Conference on Neural Information Processing Systems (NeurIPS 2024)

The 41st International Conference on Machine Learning (ICML 2024)

The 2024 International Conference on Learning Representations (ICLR 2024)

The Conference on Neural Information Processing Systems (NeurIPS 2023)

The 2023 International Conference on Learning Representations (ICLR 2023)

The Third Workshop on Generation, Evaluation & Metrics (GEM 2023)

The 39th International Conference on Machine Learning (ICML 2022)

The Conference on Neural Information Processing Systems (NeurIPS 2022)

The 2022 International Conference on Learning Representations (ICLR 2022)

Joint CDT Conference on AI for Healthcare (CAI4H 2022)

The 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP 2021)

The Conference on Neural Information Processing Systems (NeurIPS 2021)

The Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP 2021)

The 3rd Workshop on NLP for Conversational AI (NLP4CONVAI 2021)

The 2021 International Conference on Learning Representations (ICLR 2021)

The 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)

The Conference on Neural Information Processing Systems (NeurIPS 2020)

The Second Workshop on Knowledge Extraction and Integration for Deep Learning Architectures (DeeLIO 2020)

The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)

The 37th International Conference on Machine Learning (ICML 2020)

The SIGNLL Conference on Computational Natural Language Learning (CoNLL 2020)
The Conference on Neural Information Processing Systems (NeurIPS 2019)
The 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP-IJCNLP 2019)
The SIGNLL Conference on Computational Natural Language Learning (CoNLL 2019)
The 14th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2019)
The 1st Workshop on NLP for ConvAI (NLP4CONVAI 2019)
The 5th Workshop on Noisy User-generated Text (W-NUT 2019)
The 22nd Nordic Conference on Computational Linguistics (NoDaLiDa 2019)
The 36th International Conference on Machine Learning (ICML 2019)
The 2nd Learning from Limited Labeled Data Workshop (LLD 2019)
The 3rd Workshop on Evaluating Vector Space Representations for NLP (RepEval 2019)
The Second Workshop on Fact Extraction and Verification (FEVER 2019)
The Workshop on Structured Prediction for NLP (SPNLP 2019)
The Workshop on Deep Learning Approaches for Low Resource Natural Language Processing (DeepLo 2019)
The 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018)
The SIGNLL Conference on Computational Natural Language Learning (CoNLL 2018)
The Conference on Computational Linguistics (COLING 2018)
The 13th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2018)
The Seventh Joint Conference on Lexical and Computational Semantics (*SEM 2018)
The 3rd Workshop on Representation Learning for NLP (RepL4NLP 2018)
The First Workshop on Fact Extraction and Verification (FEVER 2018)
The Workshop on Figurative Language Processing (FigLang 2018)
Workshop on Deep Learning Approaches for Low Resource Natural Language Processing (DeepLo 2018)
Workshop on Subword and Character Level Models in NLP 2018 (SCLeM 2018)
Workshop on Figurative Language Processing 2018 (FigLang 2018)
The Sixth Joint Conference on Lexical and Computational Semantics (*SEM 2017)
The 12th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2017)
Workshop on Evaluating Vector Space Representations for NLP 2017 (RepEval 2017)
Workshop on Subword and Character Level Models in NLP 2017 (SCLeM 2017)
Workshop on Evaluating Vector Space Representations for NLP 2016 (RepEval 2016)
The 11th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2016)

Additional Reviewing

In addition to the program committees above, I review for the following journals and publishers:

ACL Rolling Review

MIT Press

The AAAI Conference on Artificial Intelligence (AAAI)

Journal of Natural Language Engineering (JNLE)

Special Issue of the Natural Language Engineering: Informing Neural Architectures for NLP with Linguistic and Background Knowledge

Artificial Intelligence Journal

International Journal of Artificial Intelligence in Education (IJAIED)

Cambridge Language Sciences Symposium

Programming Skills

Main experience: Python, PyTorch, Tensorflow, Theano, Java, JavaScript, PHP, SQL, HTML, CSS

Limited experience: C++, C, Android, Clojure, Perl, JSP, Lisp, Prolog, Haskell