Carsten Eickhoff, PhD

+49-7071-298-4350 | c.eickhoff@acm.org | health-nlp.com

carsteneickhoff | eickhoff | @health_nlp

Tübingen, BW - 72072, Germany

Experience

- **University of Tübingen**, *Professor* *2022 - present*
- **University Hospital Tübingen**, *Scientific Director (Medical Data Integration Center)* *2022 - present*
- **Brown University**, *Manning Assistant Professor* *2018 - 2022*
- **Harvard University**, *Visiting Fellow* *2017*
- **ETH Zurich**, *Postdoctoral Fellow* *2014 - 2017*

Education

- **Delft University of Technology** *2009 - 2014*
  Ph.D. (Computer Science) Delft, Netherlands
  - Dissertation Advisor: Arjen P. de Vries
- **The University of Edinburgh** *2008 - 2009*
  M.Sc. (Artificial Intelligence) with overall distinction Edinburgh, UK
  - Thesis Advisor: Victor Lavrenko
- **FHDW Hannover** *2005 - 2008*
  B.Sc. (Computer Science) Hanover, Germany

Metrics

- **Number of publications:** 191
- **Citations:** 7,028 (Google Scholar)
- **h-index:** 38
- **i10-index:** 104
- **Cumulative Journal Impact Factor:** 574

Publications

Peer-reviewed Journal Articles (46)

[1] Ivan Capobianco, Andrea Della Penna, André Mihaljevic, Michael Bitzer, Carsten Eickhoff, and Derna Stifini. Clinical Accuracy and Safety Concerns following GPT-5 Public Demonstration in Cancer Care. *Journal of Medical Systems*, 2025.

[2] Maya Vadiveloo, Alison Tovar, Emily Elenio, Carsten Eickhoff, John San Soucie, Sarah Feldstein Ewing, Kim Gans, and Anne Thorndike. Exploring Facilitators and Barriers for Personalized Dietary Incentives Among Online Shoppers at Cardiovascular Risk and Key Informants to Inform an Automated Shopping Platform. *Journal of Nutrition Education and Behavior*, 2025.

[3] Johanne Trippas, Shane Culpepper, Mohammad Aliannejadi, James Allan, Enrique Amigo, Jaime Arguello, Leif Azzopardi, Peter Bailey, Jamie Callan, Rob Capra, Nick Craswell, Bruce Croft, Jeff Dalton, Gianluca Demartini, Laura Dietz, Zhicheng Dou, Carsten Eickhoff, Michael Ekstrand, Nicola Ferro, Norbert Fuhr, Dorota Glowacka, Faegheh Hasibi, Danula Hettiachchi, Rosie Jones, Jaap Kamps, Noriko Kando, Sarvnaz Karimi, Makoto P Kato, Bevan Koopman, Yiqun Liu, Chenglong Ma, Joel Mackenzie, Maria Maistro, Jiaxin Mao, Dana McKay, Bhaskar Mitra, Stefano Mizzaro, Alistair Moffat, Josiane Mothe, Iadh Ounis, Lida Rashidi, Yongli Ren, Mark Sanderson, Rodrygo Santos, Falk Scholer, Chirag Shah, Laurianne Sitbon, Ian Soboroff, Damiano Spina, Paul Thomas, Julian Urbano, Arjen de Vries, Ryen White, Abby Yuan, Hamed Zamani, Oleg Zendel, Min Zhang, Justin Zobel, Shengyao Zhuang, and Guido Zuccon. Report from the Fourth Strategic Workshop on Information Retrieval in Lorne (SWIRL 2025). *SIGIR Forum*, 2025.

[4] Florian Rottach, Sebastian Schieferdecker, and Carsten Eickhoff. The Topology of Molecular Representations and Its Influence on Machine Learning Performance. *Journal of Cheminformatics*, 2025.

[5] Eliza Berman, Holly Sundberg Malek, Michael Bitzer, Nisar Malek, and Carsten Eickhoff. Retrieval Augmented Therapy Suggestion for Molecular Tumor Boards. *Journal of Medical Internet Research (JMIR)*, 2025.

[6] Friederike Holderried, Felix Eisinger, Moritz Mahling, Christian Stegemann–Philipps, Anne Herrmann–Werner, Eric Nazarenus, Alessandra Sonanini, Martina Guthoff, Carsten Eickhoff, and Martin Holderried. What's Going On with Me and How Can I Better Manage My Health? The Potential of GPT-4 to Transform Discharge Letters into Patient-Centered Letters to Enhance Patient Safety: A Prospective, Exploratory Study. *Journal of Medical Internet Research (JMIR)*, 2025.

[7] Nils Hinrichs, Alexander Meyer, Kerstin Koehler, Thomas Kaas, Meike Hiddemann, Sebastian Spethmann, Felix Balzer, Carsten Eickhoff, Volkmar Falk, Gerhard Hindricks, Nikolaos Dagres, and Friedrich Koehler. Identifying Momentary Suicidal Ideation using Machine Learning in Patients at High-Risk for Suicide. *Frontiers in Cardiovascular Medicine*, 2024.

[8] Melanie Bozzay, Christopher Hughes, Carsten Eickhoff, Heather Schatten, and Michael Armey. Identifying Momentary Suicidal Ideation using Machine Learning in Patients at High-Risk for Suicide. *Journal of Affective Disorders*, 2024.

[9] Augusto Garcia-Agundez and Carsten Eickhoff. Artificial Intelligence in Medicine: Where we Stand Today and What Lies Ahead? *Zeitschrift für Herz-,Thorax- und Gefäßchirurgie*, 2024.

[10] Nils Hinrichs, Tobias Roeschl, Pia Lanmueller, Felix Balzer, Carsten Eickhoff, Benjamin O'Brien, Volkmar Falk, and Alexander Meyer. Short-term Vital Parameter Forecasting in the Intensive Care Unit - A Benchmark Study Leveraging Data from Patients after Cardiothoracic Surgery. *PLOS Digital Health*, 2024.

[11] Tassallah Abdullahi, Laura Mercurio, Ritambhara Singh, and Carsten Eickhoff. Retrieval-Based Diagnostic Decision Support: Mixed Methods Study. *JMIR Medical Informatics*, 2024.

[12] Andrew Kalra, Preetham Bachina, Benjamin L. Shou, Jaeho Hwang, Meylakh Barshay, Shreyas Kulkarni, Isaac Sears, Carsten Eickhoff, Christian A. Bermudez, Daniel Brodie, Corey E. Ventetuolo, Glenn J. R. Whitman, Adeel Abbasi, and Sung-Min Cho. Utilizing Machine Learning to Predict Neurological Injury in Venovenous Extracorporeal Membrane Oxygenation Patients: An ELSO Registry Analysis. *JTCVS Open*, 2024.

[13] Friederike Holderried, Christian Stegemann–Philipps, Anne Herrmann–Werner, Teresa Festl–Wietek, Martin Holderried, Carsten Eickhoff, and Moritz Mahling. A Language Model-powered Simulated Patient with Automated Feedback for History Taking: Prospective Study. *JMIR Medical Education*, 2024.

[14] Andrew Kalra, Preetham Bachina, Benjamin L. Shou, Jaeho Hwang, Meylakh Barshay, Shreyas Kulkarni, Isaac Sears, Carsten Eickhoff, Christian A. Bermudez, Daniel Brodie, Corey E. Ventetuolo, Bo Soo Kim, Glenn J. R. Whitman, Adeel Abbasi, and Sung-Min Cho. Acute Brain Injury Risk Prediction Models in Venoarterial Extracorporeal Membrane Oxygenation Patients with Tree-Based Machine Learning: An ELSO Registry Analysis. *JTCVS Open*, 2024.

[15] Laura Mercurio, Augusto Garcia-Agundez, Stephanie Ruest, Susan Duffy, and Carsten Eickhoff. One Third of Alcohol Use Disorder Diagnoses are Missed by ICD Coding. *Substance Use & Addiction Journal*, 2024.

[16] Krishna Nand Keshava Murthy, Carsten Eickhoff, and Etay Ziv. Pre-operative Lung Ablation Prediction Using Deep Learning. *European Radiology*, 2024.

[17] Yuanfei Dai, Wenzhong Guo, and Carsten Eickhoff. Wasserstein Adversarial Learning based Temporal Knowledge Graph Embedding. *Information Sciences*, 2023.

[18] Tassallah Abdullahi, Ritambhara Singh, and Carsten Eickhoff. Learning to Make Rare and Complex Diagnoses with Generative AI Assistance. *JMIR Medical Education*, 2023.

[19] Adeel Abbasi, Cindy Li, Max Dekle, Christian A Bermudez, Daniel Brodie, Frank W Selke, Neel A Sodha, Corey E Ventetuolo, and Carsten Eickhoff. Interpretable Machine Learning-Based Predictive Modeling of Patient Outcomes Following Cardiac Surgery. *Journal of Thoracic and Cardiovascular Surgery*, 2023.

[20] Corey E. Ventetuolo, Navneet Singh, Carsten Eickhoff, Augusto Garcia-Agundez, Paul Bertone, Sunita S. Paudel, Dhananjay T. Tambe, Leslie A. Litzky, Katherine Cox-Flaherty, James R. Klinger, Sean F. Monaghan, Christopher J. Mullin, Mandy Pereira, Thomas Walsh, Mary Whittenhall, Troy Stevens, and Elizabeth O. Harrington. Transcriptional Profiles of Pulmonary Artery Endothelial Cells in Pulmonary Hypertension. *Nature Scientific Reports*, 2023.

[21] Augusto Garcia-Agundez, Elena García-Martín, and Carsten Eickhoff. The Potential of Machine-learning in Pharmacogenetics, Pharmacogenomics and Pharmacoepidemiology - Volume II. *Frontiers in Pharmacology*, 2023.

[22] Steffen Eickhoff, Augusto Garcia-Agundez, Daniela Haidar, Bashar Zaidat, Michael Adjei-Mosi, Peter Li, and Carsten Eickhoff. A Feasibility Study on AI-Controlled Closed-Loop Electrical Stimulation Implants. *Nature Scientific Reports*, 2023.

[23] Abdullah Ahmed, Augusto Garcia-Agundez, Ivana Petrovic, Fatemeh Radaei, James Fife, Hunter Zhou, John an Karas, Scott Moody, Jonathan Drake, Richard Jones, Carsten Eickhoff, and Michael Reznik. Delirium Detection Using Wearable Sensors and Machine Learning in Patients with Intracerebral Hemorrhage. *Frontiers in Neurology*, 2023.

[24] Claudia Meyer, Daniel Adkins, Koyena Pal, Ruggero Galici, Augusto Garcia-Agundez, and Carsten Eickhoff. Neural Text Generation in Regulatory Medical Writing. *Frontiers in Pharmacology*, 2023.

[25] Laura Mercurio, Sovijja Pou, Susan Duffy, and Carsten Eickhoff. Risk Factors for Pediatric Sepsis in the Emergency Department: A Machine Learning Pilot Study. *Pediatric Emergency Care*, 2023.

[26] Michal Golovanevsky, Carsten Eickhoff, and Ritambhara Singh. Multimodal Attention-based Deep Learning for Alzheimer's Disease Diagnosis. *Journal of the American Medical Informatics Association*, 2022.

[27] Augusto Garcia-Agundez, Elena García-Martín, and Carsten Eickhoff. The Potential of Machine-learning in Pharmacogenetics, Pharmacogenomics and Pharmacoepidemiology. *Frontiers in Pharmacology*, 2022.

[28] Jianhong Cheng, John Sollee, Celina Hsieh, Hailin Yue, Nicholas Vandal, Justin Shanahan, Ji-Whae Choi, Thi My Linh Tran, Kasey Halsey, Franklin Iheanacho, James Warren, Abdullah Ahmed, Carsten Eickhoff, Michael Feldman, Eduardo Barbosa, Ihab Kamel, Cheng Ting Lin, Thomas Yi, Terrance Healey, Paul Zhang, Jing Wu, Michael Atalay, Harrison X. Bai, Zhicheng Jiao, and Jianxin Wang. COVID-19 Mortality Prediction in the Intensive Care Unit with Deep Learning Based on Longitudinal Chest X-Rays and Clinical Data. *European Radiology*, 2022.

[29] Nicole M. Thomasian, Carsten Eickhoff, and Eli Y. Adashi. Advancing Health Equity with Artificial Intelligence. *Journal of Public Health Policy*, 2021.

[30] Yuanfei Dai, Chenhao Guo, Wenzhong Guo, and Carsten Eickhoff. Drug–drug interaction prediction with Wasserstein Adversarial Autoencoder-based knowledge graph embeddings. *Briefings in Bioinformatics*, 2021.

[31] Chris K Kim, Ji Whae Choi, Zhicheng Jiao, Dongcui Wang, Jing Wu, Thomas Y Yi, Kasey C Halsey, Feyisope Eweje, Thi My Linh Tran, Liu Chang, Robin Wang, John Sollee, Celina Hsieh, Ken Chang, Fang-Xue Yang, Ritambhara Singh, Jie-Lin Oue, Raymond Y Huang, Cai Feng, Michael D Feldman, Tao Liu, Ji Sheng Gong, Shaolei Lu, Carsten Eickhoff, Xue Feng, Ihab Kamel, Ronnie Sebro, Michael K Atalay, Terrance Healey, Yong Fan, Wei-Hua Liao, Jianxin Wang, and Harrison X Bai. An Automated Pipeline for Rapid Triage of COVID-19 Patients using Artificial Intelligence based on Chest Radiographs and Clinical Data. *Nature Digital Medicine*, 2021.

[32] Augusto Garcia-Agundez and Carsten Eickhoff. Towards Objective Quantification of Hand Tremors and Bradykinesia using Contactless Sensors: A Systematic Review. *Frontiers in Aging Neuroscience*, 2021.

[33] Krishna Nand Keshava Murthy, Carsten Eickhoff, and Krishna Juluru. Weakly Supervised Pneumonia Localization in Chest X-rays using Generative Adversarial Networks. *Medical Physics*, 2021.

[34] Seyed Ali Bahrainian, George Zerveas, Fabio Crestani, and Carsten Eickhoff. CATS: Customizable Abstractive Topic-based Summarization. *Transactions on Information Systems*, 2021.

[35] Adeel Abbasi, Yasmin Karasu, Cindy Li, Neel R Sodha, Carsten Eickhoff, and Corey Ventetuolo. Machine Learning to Predict Hemorrhage and Thrombosis during Extracorporeal Membrane Oxygenation. *Journal of Critical Care*, 2020.

[36] Michael Blaivas, Laura Blaivas, Gary Phillips, Roland Merchant, Mitchell Levy, Adeel Abbasi, Carsten Eickhoff, Nathan Shapiro, and Keith Corl. Development of a Deep Learning Network to Classify Inferior Vena Cava Collapse to Predict Fluid Responsiveness. *Journal of Ultrasound in Medicine*, 2020.

[37] Nina Rank, Boris Pfahringer, Jörg Kempfert, Christof Stamm, Titus Kühne, Felix Schoenrath, Volkmar Falk, Carsten Eickhoff, and Alexander Meyer. Deep-learning-based Real-time Prediction of Acute Kidney Injury Outperforms Human Predictive Performance. *Nature Digital Medicine*, 2020.

[38] Matthew Stib, Justin Vasquez, Mary Dong, Yun Ho Kim, Sumera Subzwari, Harold Triedman, Amy Wang, Charlene Wang, Anthony Yao, Mahesh Jayaraman, Jerrold Boxerman, Carsten Eickhoff, Ugur Cetintemel, Grayson Baird, and Ryan McTaggart. Detecting Large Vessel Occlusion on Multiphase CT Angiography Using a Deep Convolutional Neural Network. *Radiology*, 2020.

[39] Alexander Meyer, Mario Cypko, Carsten Eickhoff, Volkmar Falk, and Maximilian Emmert. AI-assisted Care in Medicine: A Revolution or yet Another Blunt Weapon? *European Heart Journal*, 2019.

[40] Yarden Raiskin, Carsten Eickhoff, and Patrick Beeler. Categorization of Free-text Drug Orders Using Character-level Recurrent Neural Networks. *International Journal of Medical Informatics*, 2019.

[41] Alexander Meyer, Dina Zverinski, Boris Pfahringer, Jörg Kempfert, Titus Kühne, Simon H. Sündermann, Christof Stamm, Thomas Hofmann, Volkmar Falk, and Carsten Eickhoff. Machine Learning for Real-time Prediction of Complications in Critical Care: A Retrospective Study. *The Lancet - Respiratory Medicine*, 2018.

[42] Harun Mustafa, Ingo Schilken, Mikhail Karasikov, Gunnar Rätsch, Carsten Eickhoff, and André Kahles. Dynamic compression schemes for graph coloring. *Bioinformatics*, 2018.

[43] Jing Yang and Carsten Eickhoff. Unsupervised learning of parsimonious general-purpose embeddings for user and location modelling. *Transactions on Information Systems*, 36(3), 2018.

[44] Carsten Eickhoff, Jacek Gwizdka, Claudia Hauff, and Jiyin He. Introduction to the special issue on search as learning. *Information Retrieval*, 20(5):399–402, 2017.

[45] Carsten Eickhoff and Arjen P de Vries. Increasing cheat robustness of crowdsourcing tasks. *Information Retrieval*, 16(2):121–137, 2013.

[46] Carsten Eickhoff and Victor Lavrenko. Towards role detection in virtual worlds. *Computers in Entertainment*, 6, 2012.

Peer-reviewed Conference Papers (113)

[1] Daniel León, Shrestha Ghosh, and Carsten Eickhoff. Investigating RAG-based Approaches in Clinical Trial and Patient Matching. In *Machine Learning for Health Symposium (ML4H)*, 2025.

[2] Shrestha Ghosh, Moritz Schneider, Carina Reinicke, and Carsten Eickhoff. A Survey on LLM-Assisted Clinical Trial Recruitment. In *Proceedings of the 2025 Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (AACL)*, 2025.

[3] Zerui Cheng, Stella Wohnig, Ruchika Gupta, Samiul Alam, Tassallah Abdullahi, João Alves Ribeiro, Christian Nielsen-Garcia, Saif Mir, Siran Li, Jason Orender, Seyed Ali Bahrainian, Daniel Kirste, Aaron Gokaslan, Carsten Eickhoff, and Ruben Wolff. Position: Benchmarking is Broken - Don't Let AI be Its Own Judge. In *Proceedings of the 2025 Conference on Conference on Neural Information Processing and Systems (NeurIPS)*, 2025.

[4] Michal Golovanevsky, William Rudman, Michael Lepori, Amir Bar, Ritambhara Singh, and Carsten Eickhoff. Pixels Versus Priors: Controlling Knowledge Priors in Vision-Language Models through Visual Counterfacts. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2025.

[5] Meng Lu, Catherine Chen, and Carsten Eickhoff. Pathway to Relevance: How Cross-Encoders Implement a Semantic Variant of BM25. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2025.

[6] Gregory Polyakov, Christian Hepting, Carsten Eickhoff, and Ali Bahrainian. Interpretability Analysis of Arithmetic In-Context Learning in Large Language Models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2025.

[7] Meng Lu, Ruochen Zhang, Carsten Eickhoff, and Ellie Pavlick. Paths Not Taken: Understanding and Mending the Multilingual Factual Recall Pipeline. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2025.

[8] Reza Esfandiarpoor, George Zerveas, Ruochen Zhang, Macton Mgonzo, Carsten Eickhoff, and Stephen H. Bach. Beyond Contrastive Learning: Synthetic Data Enables List-wise Training with Multiple Levels of Relevance. In *Findings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2025.

[9] Jessica Zosa Forde, Ruochen Zhang, Lintang Sutawika, Alham Fikri Aji, Samuel Cahyawijaya, Genta Indra Winata, Minghao Wu, Carsten Eickhoff, Stella Biderman, and Ellie Pavlick. Re-Evaluating Evaluation for Multilingual Summarization. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.

[10] William Rudman, Michal Golovanevsky, Amir Bar, Vedant Palit, Yann LeCun, Carsten Eickhoff, and Ritambhara Singh. Forgotten Polygons: Multimodal Large Language Models are Shape-Blind. In *Findings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, 2025.

[11] Tassallah Abdullahi, Ioanna Gemou, Nihal V. Nayak, Ghulam Murtaza, Stephen Bach, Carsten Eickhoff, and Ritambhara Singh. K-Paths: Reasoning over Graph Paths for Drug Repurposing and Drug Interaction Prediction. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2025.

[12] Gregory Polyakov, Catherine Chen, and Carsten Eickhoff. Towards Best Practices of Axiomatic Activation Patching in Information Retrieval. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 2025.

[13] Maria Heuss, Catherine Chen, Avishek Anand, Carsten Eickhoff, and Suzan Verberne. Workshop on Explainability in Information Retrieval. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 2025.

[14] Joschka Braun, Carsten Eickhoff, David Krueger, Seyed Ali Bahrainian, and Dmitrii Krasheninnikov. Understanding (Un)Reliability of Steering Vectors in Language Models. In *Proceedings of the ICLR BuildingTrust Workshop*, 2025.

[15] Ruochen Zhang, Qinan Yu, Matianyu Zang, Carsten Eickhoff, and Ellie Pavlick. The Same but Different: Structural Similarities and Differences in Multilingual Language Modeling. In *Proceedings of the 13th International Conference on Learning Representations (ICLR)*, 2025.

[16] Michal Golovanevsky, William Rudman, Vedant Palit, Carsten Eickhoff, and Ritambhara Singh. What Do VLMs NOTICE? A Mechanistic Interpretability Pipeline for Gaussian-Noise-free Text-Image Corruption and Evaluation. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2025.

[17] Andrew Parry, Catherine Chen, Carsten Eickhoff, and Sean MacAvaney. MechIR: A Mechanistic Interpretability Framework for Information Retrieval. In *European Conference on Information Retrieval (ECIR)*, 2025.

[18] Siran Li, Linus Stenzel, Carsten Eickhoff, and Seyed Ali Bahrainian. Enhancing Retrieval-Augmented Generation: A Study of Best Practices. In *International Conference on Computational Linguistics (COLING)*, 2025.

[19] Michal Golovanevsky, Eva Schiller, Akira A Nair, Ritambhara Singh, and Carsten Eickhoff. One-Versus-Others Attention: Scalable Multimodal Integration for Biomedical Data. In *Pacific Symposium on Biocomputing (PSB)*, 2025.

[20] Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. Talking Heads: Understanding Inter-Layer Communication in Transformer Language Models. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

[21] Michal Golovanevsky, Eva Schiller, Akira A Nair, Ritambhara Singh, and Carsten Eickhoff. One-Versus-Others Attention: Scalable Multimodal Integration for Biomedical Data. In *ICML Workshop on Accessible and Efficient Foundation Models for Biological Discovery (AccMLBio)*, 2024.

[22] Tassallah Abdullahi, Ritambhara Singh, and Carsten Eickhoff. Retrieval Augmented Zero-Shot Text Classification. In *ACM SIGIR International Conference on the Theory of Information Retrieval (ICTIR)*, 2024.

[23] Seyed Ali Bahrainian, Jonathan Dou, and Carsten and Eickhoff. Text Simplification via Adaptive Teaching. In *Findings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.

[24] Catherine Chen, Jack Merullo, and Carsten Eickhoff. Axiomatic Causal Interventions for Reverse Engineering Relevance Computation in Neural Retrieval Models. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. ACM, 2024.

[25] Catherine Chen and Carsten Eickhoff. Evaluating Search System Explainability with Psychometrics and Crowdsourcing. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. ACM, 2024.

[26] Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. Language Models Implement Simple Vector Arithmetic. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*. ACL, 2024.

[27] William Rudman and Carsten Eickhoff. Stable Anisotropic Regularization. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*, 2024.

[28] Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. Circuit Component Reuse Across Tasks in Transformer Language Models. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*, 2024.

[29] George Zerveas, Navid Rekabsaz, and Carsten Eickhoff. Enhancing the Ranking Context of Dense Retrieval through Reciprocal Nearest Neighbors. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 2023.

[30] William Rudman, Catherine Chen, and Carsten Eickhoff. Outlier Dimensions Encode Task-Specific Knowledge. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 2023.

[31] Maria Heuss, Daniel Cohen, Masoud Mansoury, Maarten de Rijke, and Carsten Eickhoff. Predictive Uncertainty-based Bias Mitigation in Ranking. In *Proceedings of the 32nd ACM International on Conference on Information and Knowledge Management*. ACM, 2023.

[32] Sofia Blinova, Xinyu Zhou, Martin Jaggi, Carsten Eickhoff, and Seyed Ali Bahrainian. SIMSUM: Document-level Text Simplification via Simultaneous Summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.

[33] Jack Merullo, Louis Castricato, Carsten Eickhoff, and Ellie Pavlick. Linearly Mapping from Image to Text Space. In *Proceedings of the Eleventh International Conference on Learning Representations (ICLR)*, 2023.

[34] Deepak Kumar, Oleg Lesota, George Zerveas, Daniel Cohen, Carsten Eickhoff, and Navid Rekabsaz. Parameter-efficient On-demand Bias Mitigation via AdapterFusion. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*. ACL, 2023.

[35] George Zerveas, Navid Rekabsaz, Daniel Cohen, and Carsten Eickhoff. CODER: An efficient framework for improving retrieval through COntextual Document Embedding Reranking. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 2022.

[36] İlkay Yıldız Potter, George Zerveas, Carsten Eickhoff, and Dominique Duncan. Unsupervised Multivariate Time-Series Transformers for Seizure Identification on EEG. In *Proceedings of the 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*. IEEE, 2022.

[37] Augusto Garcia-Agundez and Carsten Eickhoff. When BERT Fails - The Limits of EHR Classification. In *Proceedings of the AMIA Annual Symposium*. American Medical Informatics Association, 2022.

[38] Jack Merullo, Dylan Ebert, Carsten Eickhoff, and Ellie Pavlick. Pretraining on Interactions for Learning Grounded Affordance Representations. In *Proceedings of the 11th Joint Conference on Lexical and Computational Semantics (*SEM)*. ACL, 2022.

[39] George Zerveas, Navid Rekabsaz, Daniel Cohen, and Carsten Eickhoff. Mitigating Bias in Search Results through Set-based Document Reranking and Neutrality Regularization. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. ACM, 2022.

[40] Daniel Cohen, Kevin Du, Bhaskar Mitra, Laura Mercurio, Navid Rekabsaz, and Carsten Eickhoff. Inconsistent Ranking Assumptions in Medical Search and their Downstream Consequences. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. ACM, 2022.

[41] William Rudman, Nate Gillman, Taylor Rayne, and Carsten Eickhoff. IsoScore: Measuring the Uniformity of Embedding Space Utilization. In *Findings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*. ACL, 2022.

[42] Seyed Ali Bahrainian, Sheridan Feucht, and Carsten Eickhoff. NEWTS: A Corpus for News Topic-Focused Summarization. In *Findings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*. ACL, 2022.

[43] Seyed Ali Bahrainian, Martin Jaggi, and Carsten Eickhoff. Self-Supervised Neural Topic Modeling. In *Findings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 2021.

[44] Babak Hemmatian, Sheridan Feucht, Rachel Avram, Alexander Wey, Muskaan Garg, Kate Spitalnic, Carsten Eickhoff, Ellie Pavlick, Bjorn Sandstede, and Steven A. Sloman. A Novel Corpus of Discourse Structure in Humans and Computers. In *Proceedings of the EMNLP Workshop on Computational Approaches to Discourse (CODI)*. ACL, 2021.

[45] Oleg Lesota, Navid Rekabsaz, Daniel Cohen, Klaus Antonius Grasserbauer, Carsten Eickhoff, and Markus Schedl. A Modern Perspective on Query Likelihood with Deep Generative Retrieval Models. In *Proceedings of the 7th ACM SIGIR International Conference on the Theory of Information Retrieval (ICTIR)*. ACM, 2021.

[46] George Zerveas, Dhaval Jayaraman, Srideepika Patel, Anuradha Bhamidipaty, and Carsten Eickhoff. Representation Learning of Multivariate Time Series using a Transformer Framework. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*. ACM, 2021.

[47] Daniel Cohen, Bhaskar Mitra, Oleg Lesota, Navid Rekabsaz, and Carsten Eickhoff. Not All Relevance Scores are Equal: Efficient Uncertainty and Calibration Modeling for Deep Retrieval Models. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. ACM, 2021.

[48] Navid Rekabsaz, Oleg Lesota, Markus Schedl, Jon Brassey, and Carsten Eickhoff. TripClick: The Log Files of a Large Health Web Search Engine. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. ACM, 2021.

[49] Ruochen Zhang and Carsten Eickhoff. SOCCER: An Information-Sparse Discourse State Tracking Collection in the Sports Commentary Domain. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*. ACL, 2021.

[50] Abdullah Ahmed, Adeel Abbasi, and Carsten Eickhoff. Benchmarking Modern Named Entity Recognition Techniquesfor Free-text Health Record Deidentification. In *Proceedings of the AMIA Informatics Summit*. AMIA, 2021.

[51] Albert Webson, Zhizhong Chen, Carsten Eickhoff, and Ellie Pavlick. Are "Undocumented Workers" the Same as "Illegal Aliens"? Disentangling Denotation and Connotation in Vector Spaces. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 2020.

[52] Aaron S. Eisman, Nishant R. Shah, Carsten Eickhoff, George Zerveas, Elizabeth S. Chen, Wen-Chih Wu, and Indra Neil Sarkar. Extracting Angina Symptoms from Clinical Notes Using Pre-Trained Transformer Architectures. In *AMIA Informatics Summit*. AMIA, 2020.

[53] Jerome Ramos and Carsten Eickhoff. Search Result Explanations Improve Efficiency and Trust. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. ACM, 2020.

[54] Cindy Li, Elizabeth Chen, Guergana Savova, Hamish Fraser, and Carsten Eickhoff. Mining Misdiagnosis Patterns from Biomedical Literature. In *Proceedings of the AMIA Informatics Summit*. AMIA, 2020.

[55] Gil Alon, Elizabeth Chen, Guergana Savova, and Carsten Eickhoff. Diagnosis Prevalence vs. Efficacy in Machine-learning Based Diagnostic Decision Support. In *Proceedings of the AMIA Informatics Summit*. AMIA, 2020.

[56] Carsten Eickhoff, Yubin Kim, and Ryen W. White. Overview of the Health Search and Data Mining (HSDM 2020) Workshop. In *Proceedings of the 13th ACM International Conference on Web Search and Data Mining (WSDM)*. ACM, 2020.

[57] Leandro von Werra, Marcel Schöngens, D. Ece Uzun, and Carsten Eickhoff. Generative Adversarial Networks in Precision Oncology. In *Proceedings of the 5th ACM SIGIR International Conference on the Theory of Information Retrieval (ICTIR)*. ACM, 2019.

[58] Carsten Eickhoff, Floran Gmehlin, Anu Patel, Jocelyn Boullier, and Hamish Fraser. DC^3 – A Diagnostic Case Challenge Collection. In *Proceedings of the 5th ACM SIGIR International Conference on the Theory of Information Retrieval (ICTIR)*. ACM, 2019.

[59] Sebastian Hofstätter, Navid Rekabsaz, Carsten Eickhoff, and Allan Hanbury. On the Effect of Low-Frequency Terms on Neural-IR Models. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. ACM, 2019.

[60] Sebastian Hofstätter, Navid Rekabsaz, Mihai Lupu, Carsten Eickhoff, and Allan Hanbury. Enriching Word Embeddings for Patent Retrieval with Global Context. In *Proceedings of the 41st European Conference on Information Retrieval (ECIR)*. Springer, 2019.

[61] Xing Wei and Carsten Eickhoff. Distant supervision in clinical information retrieval. In *AMIA Annual Symposium Proceedings*. American Medical Informatics Association, 2018.

[62] Alba García Seco de Herrera, Carsten Eickhoff, Vincent Andrearczyk, and Henning Müller. Overview of the ImageCLEF 2018 caption prediction tasks. In *CLEF 2018 Working Notes*, CEUR Workshop Proceedings, Avignon, France, September 10-14 2018. CEUR-WS.org.

[63] Bogdan Ionescu, Henning Müller, Mauricio Villegas, Alba García Seco de Herrera, Carsten Eickhoff, Vincent Andrearczyk, Yashin Dicente Cid, Vitali Liauchuk, Vassili Kovalev, Sadid A. Hasan, Yuan Ling, Oladimeji Farri, Joey Liu, Matthew Lungren, Duc-Tien Dang-Nguyen, Luca Piras, Michael Riegler, Liting Zhou, Mathias Lux, and Cathal Gurrin. Overview of ImageCLEF 2018: Challenges, datasets and evaluation. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, Proceedings of the Ninth International Conference of the CLEF Association (CLEF 2018), Avignon, France, September 10-14 2018. LNCS Lecture Notes in Computer Science, Springer.

[64] Carsten Eickhoff. Cognitive biases in crowdsourcing. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 162–170. ACM, 2018.

[65] Ferenc Galkó and Carsten Eickhoff. Biomedical question answering via weighted neural network passage retrieval. In *Proceedings of the 40th European Conference on Information Retrieval (ECIR)*. Springer, 2018.

[66] Thijs Vogels, Octavian-Eugen Ganea, and Carsten Eickhoff. Web2text: Deep structured boilerplate removal. In *Proceedings of the 40th European Conference on Information Retrieval (ECIR)*. Springer, 2018.

[67] Carsten Eickhoff, Immanuel Schwall, Alba Garcıa Seco de Herrera, and Henning Müller. Overview of image-clefcaption 2017-image caption prediction and concept detection for biomedical images. In *Proceedings of the International Conference of the Cross-Language Evaluation Forum for European Languages (CLEF)*, 2017.

[68] Noah Hollmann and Carsten Eickhoff. Ranking and feedback-based stopping for recall-centric document retrieval. In *Proceedings of the International Conference of the Cross-Language Evaluation Forum for European Languages (CLEF)*, 2017.

[69] Bogdan Ionescu, Henning Müller, Mauricio Villegas, Helbert Arenas, Giulia Boato, Duc-Tien Dang-Nguyen, Yashin Dicente Cid, Carsten Eickhoff, Alba G Seco de Herrera, Cathal Gurrin, et al. Overview of imageclef 2017: Information extraction from images. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 315–337. Springer, 2017.

[70] Rolf Jagerman, Carsten Eickhoff, and Maarten de Rijke. Computing web-scale topic models using an asynchronous parameter server. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1337–1340. ACM, 2017.

[71] Zsolt Mezei and Carsten Eickhoff. Evaluating music recommender systems for groups. In *Proceedings of the ACM RecSys Workshop on Value-Aware Multi-Stakeholder Recommendation (VAMS)*, 2017.

[72] Piyush Bansal, Carsten Eickhoff, and Thomas Hofmann. Active content-based crowdsourcing task selection. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pages 529–538. ACM, 2016.

[73] Octavian-Eugen Ganea, Marina Ganea, Aurelien Lucchi, Carsten Eickhoff, and Thomas Hofmann. Probabilistic bag-of-hyperlinks model for entity linking. In *Proceedings of the 25th International Conference on World Wide Web*, pages 927–938. International World Wide Web Conferences Steering Committee, 2016.

[74] Maria Han Veiga and Carsten Eickhoff. A cross-platform collection of social network profiles. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 665–668. ACM, 2016.

[75] Wen Li, Carsten Eickhoff, and Arjen P de Vries. Probabilistic local expert retrieval. In *European Conference on Information Retrieval*, pages 227–239. Springer, 2016.

[76] Piyush Bansal, Andrey Bârsan, Martin Davtyan, and Carsten Eickhoff. Content-aware crowdsourcing vote aggregation. In *Proceedings of the 15th Dutch-Belgian Information Retrieval Workshop (DIR)*. DIR 2016, 2016.

[77] Paulina Grnarova, Florian Schmidt, Stephanie L Hyland, and Carsten Eickhoff. Neural document embeddings for intensive care patient mortality prediction. In *Proceedings of the NIPS 2016 Workshop on Machine Learning for Health*, 2016.

[78] Maria Han Veiga and Carsten Eickhoff. Privacy leakage through innocent content sharing in online social networks. In *Proceedings of the ACM SIGIR Workhop on Privacy Preserving Information Retrieval*. ACM, 2016.

[79] Lorenz Kuhn and Carsten Eickhoff. Implicit negative feedback in clinical information retrieval. In *Proceedings of the ACM SIGIR Medical Information Retrieval Workshop*, 2016.

[80] Jared Niederhauser, Carsten Eickhoff, Aurelien Lucchi, and Thomas Hofmann. Efficient web search diversification via approximate graph coverage. In *Proceedings of the ACM CIKM Workshop on Big Network Analytics*. ACM, 2016.

[81] Jeroen Vuurens, Carsten Eickhoff, and Arjen P de Vries. Efficient parallel learning of word2vec. In *Proceedings of the 3rd ICML Machine Learning Systems Workshop*, 2016.

[82] Nino Weingart and Carsten Eickhoff. Retrieval techniques for contextual learning. In *Proceedings of the ACM SIGIR Search as Learning Workshop*, 2016.

[83] Martin Davtyan, Carsten Eickhoff, and Thomas Hofmann. Exploiting document content for efficient aggregation of crowdsourcing votes. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, pages 783–790. ACM, 2015.

[84] Carsten Eickhoff, Sebastian Dungs, and Vu Tran. An eye-tracking study of query reformulation. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pages 13–22. ACM, 2015.

[85] Carsten Eickhoff, Arjen P de Vries, and Thomas Hofmann. Modelling term dependence with copulas. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pages 783–786. ACM, 2015.

[86] Carsten Eickhoff, Jaime Teevan, Ryen White, and Susan Dumais. Lessons from the journey: a query log analysis of within-session learning. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 223–232. ACM, 2014.

[87] Carsten Eickhoff and Arjen P de Vries. Modelling complex relevance spaces with copulas. In *Proceedings of the 23rd acm international conference on conference on information and knowledge management*, pages 1831–1834. ACM, 2014.

[88] Wen Li, Carsten Eickhoff, and Arjen P de Vries. Geo-spatial domain expertise in microblogs. In *European Conference on Information Retrieval*, pages 487–492. Springer, Cham, 2014.

[89] Wen Li, Carsten Eickhoff, and Arjen P De Vries. Interactive summarization of social media. In *Proceedings of the 5th Information Interaction in Context Symposium*, pages 312–315. ACM, 2014.

[90] Carsten Eickhoff. Crowd-powered experts: Helping surgeons interpret breast cancer images. In *Proceedings of the First International Workshop on Gamification for Information Retrieval*, pages 53–56. ACM, 2014.

[91] Carsten Eickhoff, Arjen P de Vries, and Kevyn Collins-Thompson. Copulas for information retrieval. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*, pages 663–672. ACM, 2013.

[92] Carsten Eickhoff, Kevyn Collins-Thompson, Paul Bennett, and Susan Dumais. Designing human-readable user profiles for search evaluation. In *European Conference on Information Retrieval*, pages 701–705. Springer, Berlin, Heidelberg, 2013.

[93] Carsten Eickhoff, Wen Li, and Arjen P De Vries. Exploiting user comments for audio-visual content indexing and retrieval. In *European Conference on Information Retrieval*, pages 38–49. Springer, Berlin, Heidelberg, 2013.

[94] Carsten Eickhoff, Kevyn Collins-Thompson, Paul N Bennett, and Susan Dumais. Personalizing atypical web search sessions. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 285–294. ACM, 2013.

[95] Leif Azzopardi, Doug Dowie, Sergio Duarte, Carsten Eickhoff, Richard Glassey, Karl Gyllstrom, Djoerd Hiemstra, Franciska De Jong, Frea Kruisinga, Kelly Marshall, et al. Emse: supporting children's information needs within a hospital environment. In *European Conference on Information Retrieval*, pages 578–580. Springer, 2012.

[96] Carsten Eickhoff, Leif Azzopardi, Djoerd Hiemstra, Franciska MG de Jong, Arjen P de Vries, Doug Dowie, Sergio Duarte Torres, Richard Glassey, Karl Gyllstrom, Frea Kruisinga, et al. Emse: initial evaluation of a child-friendly medical search system. In *Proceedings of the 4th Information Interaction in Context Symposium*. ACM, 2012.

[97] Carsten Eickhoff, Christopher G Harris, Arjen P de Vries, and Padmini Srinivasan. Quality through flow and immersion: gamifying crowdsourced relevance assessments. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 871–880. ACM, 2012.

[98] Carsten Eickhoff. Relevance as a subjective and situational multidimensional concept. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 998–998. ACM, 2012.

[99] Carsten Eickhoff, Pieter Dekker, and Arjen P De Vries. Supporting children's web search in school environments. In *Proceedings of the 4th Information Interaction in Context Symposium*, pages 129–137. ACM, 2012.

[100] Karl Gyllstrom, Carsten Eickhoff, Arjen P de Vries, and Marie-Francine Moens. The downside of markup: examining the harmful effects of css and javascript on indexing today's web. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, pages 1990–1994. ACM, 2012.

[101] Gabriella Kazai, Monica Landoni, Carsten Eickhoff, and Peter Brusilovsky. Booksonline'12: 5th workshop on online books, complementary social media and their impact. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, pages 2764–2765. ACM, 2012.

[102] Anton Leuski, Carsten Eickhoff, James Ganis, and Victor Lavrenko. The blademistress corpus: From talk to action in virtual worlds. In *LREC*, pages 4060–4067, 2012.

[103] Wen Li, Carsten Eickhoff, and Arjen P de Vries. Want a coffee?: predicting users' trails. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 1171–1172. ACM, 2012.

[104] Carsten Eickhoff, Pavel Serdyukov, and Arjen P De Vries. A combined topical/non-topical approach to identifying web sites for children. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 505–514. ACM, 2011.

[105] Carsten Eickhoff, Tamara Polajnar, Karl Gyllstrom, Sergio Duarte Torres, and Richard Glassey. Web search query assistance functionality for young audiences. In *European Conference on Information Retrieval*, pages 776–779. Springer, Berlin, Heidelberg, 2011.

[106] Gabriella Kazai, Carsten Eickhoff, and Peter Brusilovsky. Booksonline'11: 4th workshop on online books, complementary social media, and crowdsourcing. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, pages 2619–2620. ACM, 2011.

[107] Wen Li, Pavel Serdyukov, Arjen P de Vries, Carsten Eickhoff, and Martha Larson. The where in the tweet. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, pages 2473–2476. ACM, 2011.

[108] Carsten Eickhoff and Arjen de Vries. How crowdsourcable is your task. In *Proceedings of the ACM WSDM workshop on crowdsourcing for search and data mining*, pages 11–14, 2011.

[109] Carsten Eickhoff, Christopher G Harris, Padmini Srinivasan, and Arjen P de Vries. Geann – games for engaging annotations. In *Proceedings of the SIGIR 2011 Workshop on Crowdsourcing for Information Retrieval (CIR)*, 2011.

[110] Raynor Vliegendhart, Martha Larson, Christoph Kofler, Carsten Eickhoff, and Johan Pouwelse. Investigating factors influencing crowdsourcing tasks with high imaginative load. In *Proceedings of the ACM WSDM Workshop on Crowdsourcing for Search and Data Mining*, pages 27–30, 2011.

[111] Jeroen Vuurens, Arjen P de Vries, and Carsten Eickhoff. How much spam can you take? an analysis of crowdsourcing results to increase accuracy. In *Proceedings of the ACM SIGIR Workshop on Crowdsourcing for Information Retrieval*, pages 21–26, 2011.

[112] Carsten Eickhoff and Arjen P de Vries. Identifying suitable youtube videos for children. *3rd Networked and electronic media summit (NEM)*, 2010.

[113] Carsten Eickhoff, Pavel Serdyukov, and Arjen P de Vries. Web page classification on child suitability. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, pages 1425–1428. ACM, 2010.

Edited Volumes (5)

[1] Avi Arampatzis, Evangelos Kanoulas, Theodora Tsikrika, Stefanos Vrochidis, Hideo Joho, Christina Lioma, Carsten Eickhoff, Aurélie Névéol, Linda Cappellato, and Nicola Ferro, editors. *Experimental IR Meets Multilinguality, Multimodality, and Interaction - Proceedings of the Eleventh International Conference of the CLEF Association (CLEF 2020)*, volume 12260. Springer, 2020.

[2] Linda Cappellato, Carsten Eickhoff, Nicola Ferro, and Aurélie Névéol, editors. *Working Notes of CLEF 2020 - Conference and Labs of the Evaluation Forum*, volume 2552. CEUR WS, 2020.

[3] Carsten Eickhoff, Yubin Kim, and Ryen White, editors. *Proceedings of the 1st ACM WSDM Health Search and Data Mining Workshop*, volume 2551. CEUR WS, 2020.

[4] Carsten Eickhoff, Jacek Gwizdka, Claudia Hauff, and Jiyin He, editors. *Information Retrieval*, volume 20. Springer, 2017.

[5] Carsten Eickhoff and Arjen de Vries, editors. *Proceedings of the Dutch-Belgian Workshop on Information Retrieval*. Number 13. CEUR Workshop Proceedings, 2013.

Conference Abstracts (17)

[1] Joschka Braun, Carsten Eickhoff, and Seyed Ali Bahrainian. Beyond Multiple Choice: Evaluating Steering Vectors for Adaptive Free-Form Summarization. In *ICML Workshop on Test-Time Adaptation*, 2025.

[2] Melanie Bozzay, Christopher Hughes, Carsten Eickhoff, Heather Schatten, and Michael Armey. Identifying Momentary Ideation via Machine Learning Models in Patients at High-risk for Suicide. In *European Symposium on Suicide and Suicidal Behavior*, 2024.

[3] Adeel Abbasi, Isaac Sears, George Zerveas, Corey Ventetuolo, Neel Sodha, and Carsten Eickhoff. Transfer Learning Boosts Performance Of Deep Learning Model Prediction of In-Hospital Mortality In Patients Treated With Extracorporeal Membrane Oxygenation. In *Proceedings of the 34th Annual ELSO Conference*. ELSO, 2023.

[4] Stephanie Ruest, William Rudman, Carsten Eickhoff, and Susan Duffy. Characteristics of Pediatric Emergency Department Encounters for Fractures Concerning for Abuse. In *Proceedings of the Injury Free Coalition for Kids Conference*, 2023.

[5] Isaac Sears, Augusto Garcia-Agundez, George Zerveas, William Rudman, Laura Mercurio, Adeel Abbasi, and Carsten Eickhoff. Leveraging Unlabeled Electroencephalographic Data to Predict Neurologic Recovery After Cardiac Arrest. In *Proceedings of the 50th Computing in Cardiology Conference (CinC)*. Physionet, 2023.

[6] Isaac Sears, Mitchell Levy, Corey E Ventetuolo, Carsten Eickhoff, and Adeel Abbasi. Time-Series Machine Learning Approach to Sepsis Prediction in the Intensive Care Unit. In *Proceedings of the Annual Meeting of the American Thoracic Society (ATS)*, 2023.

[7] Katherine Zhong, Elizabeth Chen, Carsten Eickhoff, and Paul Greenberg. Using Machine Learning to Predict Hospitalization and Mortality of COVID-19 Patients with Diabetic Retinopathy. In *Proceedings of the Annual Meeting of the Association for Research in Vision and Ophthalmology (ARVO)*, 2023.

[8] Michal Golovanevsky, Carsten Eickhoff, and Ritambhara Singh. Multimodal Attention-based Deep Learning for Alzheimer's Disease Diagnosis. In *Proceedings of the Machine Learning in Computational Biology Conference (MLCB)*, 2022.

[9] Stephen Schmit, Rebecca Ortiz, Philip Caffery, Christopher T. Tucci, Carsten Eickhoff, Elias Hyams, and Gyan Pareek. Artificial Intelligence in Robotic-Assisted Laparoscopic Prostatectomy: Developing Predictive Models for Positive Surgical Margins. In *Proceedings of the American Urological Association's Annual Conference*. AUA, 2022.

[10] Babak Hemmatian, Aryan Srivastava, Nathaniel Goodman, Jonathan Lee, Carsten Eickhoff, and Steven A. Sloman. Arguments and Stories: A Big Data Analysis of the Emerging Consensus on Marijuana Legalization (2008-2019). In *Society for Personality and Social Psychology Convention*. SPSP, 2021.

[11] Babak Hemmatian, Nathaniel Goodman, Jonathan Lee, Sabina J. Sloman, Uriel Cohen Priva, Carsten Eickhoff, and Steven A. Sloman. Narratives of Consensus in a Decade of Reddit Discourse about Hot-Button Issues. In *Proceedings of the International Conference on Thinking (ICT)*. ICT, 2021.

[12] Babak Hemmatian, Rachel Avram, Sheridan Feucht, Alexander Wey, Kate Spitalnic, Muskaan Garg, Carsten Eickhoff, Björn Sandstede, Ellie Pavlick, and Steven A. Sloman. The Anatomy of Marijuana Legalization Discourse: Narrative and Argument Quality in Human and Computer-generated Texts. In *Proceedings of the International Conference on Thinking (ICT)*. ICT, 2021.

[13] Alex Wey, Babak Hemmatian, Sheridan Feucht, Kate Spitalnic, Muskaan Garg, Carsten Eickhoff, Ellie Pavlick, Björn Sandstede, and Steven Sloman. Can computers tell a story? Discourse Structure in Computer-generated Text and Humans. In *Annual Conference of the Cognitive Science Society (CogSci)*. CSS, 2021.

[14] Sheridan Feucht, Babak Hemmatian, Alex Wey, Kate Spitalnic, Muskaan Garg, Carsten Eickhoff, Ellie Pavlick, Björn Sandstede, and Steven Sloman. The Anatomy of Discourse: Linguistic Predictors of Narrative and Argument Quality. In *Annual Conference of the Cognitive Science Society (CogSci)*. CSS, 2021.

[15] Babak Hemmatian, Sabina Sloman, Uriel Cohen Priva, Carsten Eickhoff, and Steven Sloman. Evidence-responsiveness in a Decade of Reddit Discourse about Hot-button Topics. In *Society for Personality and Social Psychology Convention*. SPSP, 2020.

[16] Ruochen Zhang, Zenjiang Shen, Carsten Eickhoff, and Ellie Pavlick. UndNet: A Mutual Understanding Maximization Framework for Neural Dialogue Generation. In *Natural Language, Dialogue and Speech Symposium (NDS)*. The New York Academy of Sciences, 2019.

[17] Benjamin Insley, Syed Rizvi, Jonathon Cahill, Joshua Stone, and Carsten Eickhoff. Deep Neural Net Forecasting of Multiple Sclerosis Disease Severity. In *Proceedings of the 71st Annual Meeting of the American Academy of Neurology (AAN)*, 2019.

Other Publications (15)

[1] Isaac Sears, Augusto Garcia-Agundez, George Zerveas, William Rudman, Laura Mercurio, Corey E. Ventetuolo, Adeel Abbasi, and Carsten Eickhoff. Leveraging Unlabeled Electroencephalographic Data to Predict Neurological Recovery for Comatose Patients Following Cardiac Arrest. 2023.

[2] William Rudman, Jack Merullo, Laura Mercurio, and Carsten Eickhoff. ACQuA: Arrhythmia Classification with Quasi-Attractors. In *Proceedings of the 49th Computing in Cardiology Conference (CinC)*. Physionet, 2022.

[3] George Zerveas, Ruochen Zhang, Leila Kim, and Carsten Eickhoff. Brown university at trec deep learning 2019. In *Proceedings of the 28th Text Retrieval Conference (TREC)*. NIST, 2019.

[4] Abdullah Ahmed, Bashar Zaidat, Hwai-Liang Tung, Isaac Nathoo, Charles Wang, and Carsten Eickhoff. Brown university at trec precision medicine 2019. In *Proceedings of the 28th Text Retrieval Conference (TREC)*. NIST, 2019.

[5] Xing Wei and Carsten Eickhoff. Embedding electronic health records for clinical information retrieval. In *https://arxiv.org/abs/1811.05402*, 2018.

[6] Prakrit Baruah, Riya Dulepet, Kyle Qian, and Carsten Eickhoff. Brown university at trec precision medicine 2018. In *Proceedings of the 27th Text Retrieval Conference (TREC)*, 2018.

[7] Negar Foroutan Eghlidi, Jannick Griner, Nicolas Mesot, Leandro von Werra, and Carsten Eickhoff. Eth zurich at trec precision medicine 2017. In *Proceedings of the 26th Text Retrieval Conference (TREC)*, 2017.

[8] Carsten Eickhoff and Arjen P de Vries. Robust statistical methods in web retrieval. *ACM SIGWEB Newsletter*, (Winter):4, 2016.

[9] Simon Greuter, Philip Junker, Lorenz Kuhn, Felix Mance, Virgile Mermet, Angela Rellstab, and Carsten Eickhoff. Eth zurich at trec 2016 clinical decision support. In *NIST Special Publication 500-321: The Twenty-Fifth Text REtrieval Conference Proceedings (TREC 2016)*. National Institute of Standards and Technology, 2016.

[10] Carsten Eickhoff. Human intelligence in search and retrieval. In *Proceedings of the 2nd ECIR Workshop on Gamification in Information Retrieval*, page 3, 2015.

[11] Carsten Eickhoff. *Contextual Multidimensional Relevance Models*. PhD thesis, Delft University of Technology, 10 2014.

[12] Gabriella Kazai, Carsten Eickhoff, and Peter Brusilovsky. Report on booksonline'11: 4th workshop on online books, complementary social media, and crowdsourcing. In *ACM SIGIR Forum*, volume 46, pages 43–50. ACM, 2012.

[13] Carsten Eickhoff, Christopher G Harris, Padmini Srinivasan, and Arjen P de Vries. Geann at the trec 2011 crowdsourcing track. In *TREC*, 2011.

[14] Wen Li, Carsten Eickhoff, and Arjen P De Vries. Dmir on microblog track 2011. In *TREC*, 2011.

[15] Jeroen BP Vuurens, Carsten Eickhoff, and Arjen P de Vries. Managing the quality of large-scale crowdsourcing. In *TREC*. NIST, 2011.

AWARDS

- **Computer Science Teaching Award (University of Tübingen)** *2024*
- **Best Paper Award (IEEE ICMLA)** *2022*
- **Manning Endowed Chair of Medical Science (Brown University)** *2022*
- **Best Short Paper Award (ECIR)** *2019*
- **Karen T. Romer Undergraduate Teaching and Research Award (Brown University)** *2019*
- **DHZK Research Paper of the Month** *2018*
- **Winning entry to the BioASQ Biomedical Question Answering Challenge** *2017*
- **Best Abstract Award (AIMed)** *2016*
- **Best Paper Honorable Mention (ACM SIGIR)** *2015*
- **Best Reviewer Award (ECIR)** *2014*
- **Microsoft Bing Most Innovative Paper Award** *2011*

GRANTS AS PI (€ 8,648,958)

- **UNC-Tübingen Collaborative Seed Grant (PI)** *2025 - 2026*
  Explainable Clinical Decision Support Systems. **€10,000**
  ◦ University of Tübingen
- **ProfilPLUS (PI)** *2023 - 2025*
  Artificial Intelligence in Biomedicine. **€75,600**
  ◦ University of Tübingen
- **Covid-19 Relief Grant (PI)** *2022 - 2023*
  Explaining COVID-19 Outcome Disparities via Natural Language Processing. **$100,000**
  ◦ Peter G. Peterson Foundation Pandemic Response Policy Research Fund
- **NSF/III Medium (PI)** *2020 - 2024*
  Generative Neural Information Retrieval Models. **$999,761**
  ◦ National Science Foundation

- **OVPR Seed Award (PI)** *2020 - 2021*
  Finding Social Narratives in Big Data. **$71,000**
  ◦ Brown University
- **DARPA Seed (Co-PI)** *2019 - 2022*
  Cognitively-Motivated Word Learning in Embodied Virtual Agents. **$954,509**
  ◦ Defense Advanced Research Projects Agency
- **IARPA BETTER (Co-PI)** *2019 - 2023*
  Task-and-User Aware Representation Learning for Fine-Grained Cross-Lingual Information Retrieval. **$6,100,000**
  ◦ Intelligence Advanced Research Projects Activity
- **NSF Workshop Support (PI)** *2018 - 2019*
  Student Travel Support for ACM WSDM 2019. **$18,450**
  ◦ National Science Foundation
- **OVPR Big Data Seed Grant (PI)** *2018 - 2019*
  Forecasting Patterns of Delirium and Early Recovery After Acute Stroke. **$25,000**
  ◦ Brown University
- **NSF Workshop Support (PI)** *2018 - 2019*
  Doctoral Consortium for HCOMP 2018. **$24,009**
  ◦ National Science Foundation
- **SNSF Ambizione (PI)** *2017 - 2019*
  Representation Learning for Clinical Artificial Intelligence. **CHF 743,444**
  ◦ Swiss National Science Foundation
- **Microsoft Azure for Research Grant (PI)** *2013 - 2014*
  Statistically Robust Multi-criteria Relevance Modelling. **$47,000**
  ◦ Microsoft

Grants as co-I (€ 66,737,557)

- **DFG Cluster of Excellence (Co-I)** *2026 - 2032*
  Machine Learning: New Perspectives for Science. **€50,000,000**
  ◦ German Research Foundation
- **NIH/AHRQ R01 (Co-I)** *2023 - 2028*
  Evaluating Diagnostic Decision Support Systems for Patients Requiring Urgent Primary or Emergency Care or with Stroke. **$1,929,759**
  ◦ National Institutes of Health
- **NIH/NIGMS COBRE (Co-I)** *2021 - 2026*
  Rhode Island Hospital Injury Control Center for Biomedical Research Excellence. **$14,807,798**
  ◦ National Institutes of Health

Grants Mentored (€ 2,207,641)

- **NIH NHLBI K01** *2023 - 2028*
  Maya Vadiveloo **$861,550**
- **ACM SIGHPC Computational & Data Science Fellowship** *2022 - 2026*
  Tassallah Abdullahi **$60,000**
- **Fulbright Research Award** *2022 - 2023*
  Kevin Du **$24,054**
- **NIH AdvanceCTR Pilot** *2022 - 2023*
  Joshua Kemp **$37,430**
- **NIH NIMH K23** *2021 - 2026*
  Taylor Burke **$978,240**
- **ERC Marie Skłodowska-Curie Fellowship** *2022 - 2025*
  Augusto Garcia-Agundez **€245,732**
- **SNSF Early Postdoc Mobility Fellowship** *2020 - 2021*
  Seyed Ali Bahreinian **CHF 75,450**
- **CSC Chinese Government Scholarship** *2019 - 2020*
  Yuanfei Dai **$30,000**

Invited Talks

[1] "Juggling 1.8T Balls - The Frontier of LM Interpretation". **Charite, Germany**, *Invited Talk*, 11-24-2025.

[2] "Pixels Versus Priors: Controlling Knowledge Flow in Vision-Language Models". **RIKEN IMS, Japan**, *Invited Talk*, 11-04-2025.

[3] "Juggling 1.8T Balls - The Frontier of LM Interpretation". **Bloomberg, UK**, *Invited Talk*, 10-09-2025.

[4] "Scientific Communication". **Machine Learning Conference, Tübingen, Germany**, *Invited Panelist*, 09-30-2025.

[5] "KI in der Medizin – Chancen, Grenzen und Risiken". **Köhler Campus, Germany**, *Invited Talk*, 08-19-2025.

[6] "Generative AI in Medicine". **Tübingen Workshop for Blood Cellular Therapies, Germany**, *Keynote Address*, 06-26-2025.

[7] "Juggling 1.8T Balls - The Frontier of LM Interpretation". **University of North Carolina, Chapel Hill NC**, *Invited Talk*, 06-18-2025.

[8] "Juggling 1.8T Balls - The Frontier of LM Interpretation". **University of Melbourne, Australia**, *Invited Talk*, 02-13-2025.

[9] "Juggling 1.8T Balls - The Frontier of LM Interpretation". **CZS Institute for AI and Law, Germany**, *Invited Talk*, 02-05-2025.

[10] "What is Going on in Language Models?". **TU Munich, Germany**, *Invited Talk*, 12-05-2024.

[11] "Juggling 1.8 Trillion Balls - Explainable Natural Language Processing". **RIKEN IMS, Japan**, *Keynote Address*, 10-16-2024.

[12] "AI and Large Language Models - Revolutionizing Healthcare". **Eurospine, Austria**, *Keynote Address*, 10-04-2024.

[13] "Mechanistic Interpretability". **NII Shonan, Japan**, *Invited Talk*, 10-02-2024.

[14] "Artificial Intelligence in Medicine". **Köhler Campus, Germany**, *Invited Talk*, 09-10-2024.

[15] "Artificial Intelligence - A Glance under the Hood". **University of Halle, Germany**, *Keynote Address*, 09-06-2024.

[16] "A Machine Learning View on Algorithmic Collusion". **AI MEETS LAW, University of Tübingen, Germany**, *Invited Talk*, 04-19-2024.

[17] "Generative AI in Medicine". **Heart and Diabetes Center NRW, Germany**, *Invited Talk*, 03-21-2024.

[18] "LLMs - A Glimpse under the Hood". **ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR), UK**, *Invited Talk*, 03-12-2024.

[19] "Masterclass: Data Justice". **University of Tübingen, Germany**, *Keynote Address*, 02-20-2024.

[20] "Large Language Models in Clinical Use". **German Society for Thoracic and Cardiac Surgery, Germany**, *Keynote Address*, 02-17-2024.

[21] "Transparent AI Methods in Clinical Medicine". **Integrata Foundation, Germany**, *Keynote Address*, 01-22-2024.

[22] "Diagnostic Decision Support via Zero Shot Text Classification". **University of Tübingen PhD Program Experimental Medicine, Germany**, *Invited Talk*, 11-29-2023.

[23] "ChatGPT & Friends - An Introduction to Generative Language Models". **State Library of Württemberg, Germany**, *Keynote Address*, 11-21-2023.

[24] "Knowledge from Data". **German Society for Thoracic and Cardiac Surgery, Germany**, *Keynote Address*, 10-28-2023.

[25] "Who controls the AI?". **Tübingen Museum Society and SWR, Germany**, *Panel Discussion*, 10-20-2023.

[26] "Data Science & Health Career Panel". **National Student Data Corp, New York**, *Panel Discussion*, 09-29-2023.

[27] "The Internet and How to Wield it". **International Max Planck Research School, Germany**, *Tutorial*, 09-15-2023.

[28] "Medicine in Artificially Intelligent Times". **University of Tübingen, Medical School Graduation Ceremony, Germany**, *Keynote Address*, 07-21-2023.

[29] "Retrieval-powered Zero-shot Text Classification". **University of Applied Sciences Cologne, Germany**, *Invited Talk*, 05-05-2023.

[30] "Retrieval-powered Zero-shot Text Classification". **Max Planck Institute for Intelligent Systems, Germany**, *Invited Talk*, 04-05-2023.

[31] "Uncertainty in Information Retrieval?". **University of Massachusetts, Massachusetts**, *Invited Talk*, 02-17-2023.

[32] "Retrieval-powered Zero-shot Text Classification". **Oracle, Switzerland**, *Invited Talk*, 02-15-2023.

[33] "Retrieval-powered Zero-shot Text Classification". **University of Augsburg, Germany**, *Invited Talk*, 01-16-2023.

[34] "Retrieval-powered Zero-shot Text Classification". **University of Tübingen, Germany**, *Invited Talk*, 12-12-2022.

[35] "An Introduction to Deep Learning". **Advance-CTR Machine Learning for Health Seminar, Rhode Island**, *Tutorial*, 10-07-2022.

[36] "TripClick: The Log Files of a Large Health Web Search Engine". **University of Amsterdam, Netherlands**, *Invited Talk*, 05-10-2022.

[37] "Uncertainty and Calibration in Search". **Search Engines Amsterdam, Netherlands**, *Invited Talk*, 04-29-2022.

[38] "Uncertainty and Calibration in Search". **Radboud University Nijmegen, Netherlands**, *Invited Talk*, 04-08-2022.

[39] "Uncertainty and Explainability in Search". **University of Amsterdam, Netherlands**, *Invited Talk*, 12-03-2021.

[40] "An Introduction to Machine Learning". **Brown University Data Science Initiative and Advance-CTR, Rhode Island**, *Invited Talk*, 10-01-2021.

[41] "Health Search and Text Understanding". **Kaiser Permanente, New York**, *Invited Talk*, 09-24-2021.

[42] "Keynote - AI in Sales and Marketing". **Merck, Austria**, *Keynote Address*, 09-01-2021.

[43] "Uncertainty Estimation and Explainability in Clinical Machine Learning". **University of Munich, Germany**, *Invited Talk*, 07-28-2021.

[44] "Uncertainty Estimation and Explainability in Information Retrieval". **University of Hanover, Germany**, *Invited Talk*, 07-01-2021.

[45] "Uncertainty Estimation and Explainability in Clinical Machine Learning". **University of Tübingen, Germany**, *Invited Talk*, 06-28-2021.

[46] "AI in Personalized Medicine". **Austrian Health Forum, Austria**, *Keynote Address*, 06-24-2021.

[47] "Processing Clinical Notes with Transformers". **Radboud UMC Diagnostic Image Analysis Group (DIAG), Netherlands**, *Invited Talk*, 03-10-2021.

[48] "Deep-learning-based Real-time Prediction of Acute Kidney Injury". **Foundation for the National Institutes of Health (FNIH), New York**, *Invited Talk*, 02-23-2021.

[49] "Machine Learning in Medicine - A Primer on Methods for Physicians". **German Center for Cardiovascular Research (DHZK), Germany**, *Tutorial*, 02-19-2021.

[50] "Artificial Intelligence in Medicine". **Rhode Island Neurological Society, RI**, *Keynote Address*, 11-18-2020.

[51] "Is This It? AI for Differential Diagnosis". **Slater Technology Fund, RI**, *Invited Talk*, 10-28-2020.

[52] "Cognitive Bias in Crowdsourcing". **Hack@Home, RI**, *Invited Talk*, 10-10-2020.

[53] "State of the Art and Open Challenges in Biomedical AI". **New England Medical Innovation Center (NEMIC) Summit, RI**, *Invited Talk*, 09-25-2020.

[54] "Prevalence vs. Efficacy in Diagnostic Decision Support". **Bioinformatics 4All, NY**, *Invited Talk*, 08-16-2020.

[55] "Improving the Efficiency of Medical Writing with AI". **DIA Medical Affairs and Scientific Communications Forum, NY**, *Invited Talk*, 03-25-2020.

[56] "Predicting Adverse Events in Intensive Care". **Pfizer Inc., NY**, *Invited Talk*, 11-04-2019.

[57] "Abstractive Summarization of Electronic Health Records". **Google Improving Healthcare with AI Workshop, Israel**, *Invited Talk*, 10-23-2019.

[58] "Clinical Natural Language Processing". **IDIAP Research institute, Switzerland**, *Invited Talk*, 08-20-2019.

[59] "Interactive Summarization of Electronic Health Records". **Microsoft Faculty Summit, WA**, *Invited Talk*, 07-19-2019.

[60] "Personalized Text Summarization". **Spotify Inc., MA**, *Invited Talk*, 06-25-2019.

[61] "AI in Healthcare Panel". **Ocean Innovation Technology Fest, RI**, *Panel Discussion*, 04-27-2019.

[62] "Clinical Text Understanding and Decision Support". **TU Munich, Germany**, *Invited Talk*, 06-29-2017.

[63] "Introduction to Deep Learning". **TUI Machine Learning Group, Spain**, *Invited Talk*, 03-30-2017.

[64] "Clinical Text Understanding and Decision Support". **University of Waterloo, Canada**, *Invited Talk*, 03-23-2017.

[65] "Clinical Text Understanding and Decision Support". **Harvard University, MA**, *Invited Talk*, 02-17-2017.

[66] "Clinical Text Understanding and Decision Support". **University of Massachusetts Amherst, MA**, *Invited Talk*, 02-10-2017.

[67] "Clinical Text Understanding and Decision Support". **University of Illinois, IL**, *Invited Talk*, 02-06-2017.

[68] "Clinical Text Understanding and Decision Support". **Rice University, TX**, *Invited Talk*, 01-26-2017.

[69] "Introduction to Deep Learning". **University of Zurich, Switzerland**, *Invited Talk*, 11-22-2016.

[70] "Clinical Text Understanding and Decision Support". **EPFL Lausanne, Switzerland**, *Invited Talk*, 11-09-2016.

[71] "Patient-centric Literature Retrieval". **Swiss eHealth Summit, Switzerland**, *Invited Talk*, 09-21-2016.

[72] "Clinical Information Retrieval". **University of Zurich, Switzerland**, *Invited Talk*, 04-20-2016.

[73] "Human Intelligence in Search and Retrieval". **ECIR GamifIR Gamification Workshop, Austria**, *Keynote Address*, 04-01-2015.

[74] "Introduction to Information Retrieval". **Amsterdam Polytechnic University, Netherlands**, *Invited Talk*, 03-21-2013.

[75] "Exploiting User Comments for Content Indexing". **Dutch Hadoop Meetup, Netherlands**, *Invited Talk*, 11-16-2012.

[76] "The Role of the User in IR". **University of Hannover, Germany**, *Invited Talk*, 05-15-2012.

[77] "Information Retrieval for Children". **Free University Amsterdam, Netherlands**, *Invited Talk*, 12-05-2011.

[78] "Filtering Shared Video Content for Children". **ICT.Open, Netherlands**, *Invited Talk*, 11-14-2011.

[79] "Information Retrieval for Children". **Radboud University Nijmegen, Netherlands**, *Invited Talk*, 04-06-2011.

[80] "Introduction to Information Extraction". **Twente University, Netherlands**, *Invited Talk*, 11-23-2010.

Teaching

- **Modern Search Engines**, University of Tübingen *Summer 2023 - 2025*
- **Natural Language Processing**, University of Tübingen *Summer 2023 - 2025*
- **Understanding Language Models**, University of Tübingen *Summer 2024 - 2025*
- **Introduction to Python Programming**, University of Tübingen *Winter 2023 - 2025*
- **AI in Medicine**, University of Tübingen *Summer 2023*
- **Introduction to Artificial Intelligence**, University of Tübingen *Winter 2023*
- **Artificial Intelligence in Biomedicine**, Brown University *Spring 2019 - 2021*
- **Advanced Topics in IR & NLP**, ETH Zürich *Spring 2016 - 2017*
- **Machine Learning**, ETH Zürich *Fall 2016*
- **Research in Computer Science**, ETH Zürich *Fall 2014 - 2016*

Advising

Junior Faculty (6)
Harrisen Scells (Assistant Professor at University of Tübingen), **Kristof Meding** (Assistant Professor at University of Tübingen), **Taylor Burke** (Assistant Professor at Harvard University), **Adeel Abbasi** (Assistant Professor at Brown University), **Maya Vadiveloo** (Associate Professor at the University of Rhode Island), **Laura Mercurio** (Assistant Professor at Brown University)

Postdoc (4)
Shrestha Ghosh (ongoing), **Seyed Ali Bahrainian** (ongoing), **Augusto Garcia-Agundez** (Postdoc at UCSF), **Daniel Cohen** (Research Scientist at DataMinr)

Ph.D. (14)
Miriam Rateike (ongoing), **Siran Li** (ongoing), **Gregory Polyakov** (ongoing), **Florian Rottach** (ongoing), **Michal Golovanevsky** (ongoing), **Catherine Chen** (ongoing), **Tasallah Abdullahi** (ongoing), **Ruochen Zhang** (ongoing), **William Rudman Jr.** (PhD 2025, Postdoc at UT Austin), **John Merullo** (PhD 2025, Research Scientist at Goodfire), **Georgios Zerveas** (PhD 2023, Senior Applied Scientist at Microsoft Turing), **Krishna Nand Keshava Murthy** (PhD 2022, Lead Data Scientist at Memorial Sloan Kettering Cancer Center), **Zhizhong Chen** (PhD 2021, Software Engineer at Twitter), **Wen Li** (PhD 2015, Software Engineer at Google)

M.Sc. (45)
Nianzi Yi, **Selina Ücker**, **Leon Scharwächter**, **Michelle Schlicher**, **Felix Strothmann**, **Michael Giang**, **Simon Frank**, **Dominik Hildebrand**, (PhD candidate at University of Tübingen), **Kora Helm**, **Marc Seelmann**, **David Bernhardt**, **Koyena Pal**, (PhD candidate at Northeastern University), **Babak Hemmatian Borujeni**, (Assistant Professor at University of Nebraska-Lincoln), **Ge Zhu**, **Priscilla Hernandez**, **Kunyu Zhang**, **Li Zhu**, **Heiki Riesenkampf**, **Jannick Griner**, **Julian Fuchs**, **Ingo Schilken**, **Karl Ruszics** (PhD candidate at ETH Zurich), **Floran Gmehlin** (Data Scientist at Pricehubble AG), **Xing Wei** (Developer at Google), **Lennart van der Goten** (Data Scientist at SAP), **David Frank** (Developer at Google), **Ferenc Galko** (Developer at Google), **Rik Melis** (Developer at AutoForm Engineering GmbH), **Arturo Vivas**, **Dina Zverinski** (Developer at Google), **Yarden Raiskin** (Developer at IBM), **Felix Mance** (Developer at Palantir), **Immanuel Schwall** (PhD candidate at ETH Zurich), **Simon Greuter**, **Zsolt Mezei** (Co-Founder at Gleblu AG), **Jing Yang** (PhD candidate at ETH Zurich), **Faiza Alsaied** (Developer at AutoForm Engineering GmbH), **Rolf Jagerman** (Developer at Google), **Jan Wilken Dörrie** (Developer at Facebook), **Lukas Elmer** (Co-Founder at Renuo

AG), **Maria Han Veiga** (Assistant Professor at Ohio State University), **Jared Niederhauser** (Developer at GetYourGuide), **Martin Davtyan** (Data Scientist at Episteme Capital), **Siddharth Sarda** (Data Scientist at Booking.com), **Marina Horlescu** (Developer at Google), **Florian Schmidt** (PhD candidate at ETH Zurich), **Pieter Dekker** (Developer at Apprentice Uitgevers)

B.Sc. (21)
Eliza Gilbert, **Marco Madeja**, **Saman Sang**, **Koyena Pal**, **Abdullah Ahmed**, **Karen Tu**, **Prakrit Baruah**, **Erica Guo**, **Jerome Ramos**, **Nicolas Mesot** (MSc at ETH Zurich), **Samuel Müller** (MSc at ETH Zurich), **Panuya Balasuntharam** (MSc at ETH Zurich), **Kevin Klein** (MSc at ETH Zurich), **Noah Hollmann** (MSc at ETH Zurich), **Angela Rellstab** (MSc at ETH Zurich), **Lorenz Kuhn** (MSc at ETH Zurich), **Philip Junker** (MSc at ETH Zurich), **Tobias Verhulst** (MSc at ETH Zurich), **Nino Weingart** (MSc at ETH Zurich), **Nina Zinsli** (MSc at ETH Zurich), **Kieran Nirkko** (Intern at Google)

Doctoral Committees (32)
Ehsan Doostmohammadi (University of Linköping, Sweden, PhD, *expected*), **Dominik Langner** (University of Tübingen, Germany, PhD, *expected*), **Stefan Kraft** (University of Tübingen, Germany, PhD, *expected*), **Ahmad-Reza Ehyaei** (University of Tübingen, Germany, PhD, *expected*), **Katrin Renz** (University of Tübingen, Germany, PhD, *expected*), **Paul Fischer** (University of Tübingen, Germany, PhD, *expected*), **Denise Junger** (University of Tübingen, Germany, PhD, *expected*), **Johann Frei** (University of Augsburg, Germany, PhD, *expected*), **Junjie He** (University of Stuttgart, Germany, PhD, *expected*), **Prithivi Jung Thapa** (University of Tübingen, Germany, PhD, *expected*), **Tom Sühr** (University of Tübingen, Germany, PhD, *expected*), **Polina Tsvilodub** (University of Tübingen, Germany, PhD, *expected*), **Meghal Dani** (University of Tübingen, Germany, PhD, *expected*), **Raphael Scheible** (TU Munich, Germany, PhD, *expected*), **Rita Gonzalez Marquez** (University of Tübingen, Germany, PhD, *expected*), **Diogo Nunes** (University of Lisbon, Portugal, PhD, 2025), **Shuay Wang** (University of Queensland, Australia, PhD, 2025), **Maria Heuss** (University of Amsterdam, The Netherlands, PhD, 2025), **Isabelle Loster** (University of Tübingen, Germany, MD, 2025), **David La Barbera** (University of Udine, Italy, PhD, 2025), **Mihai Horia Popescu** (University of Udine, Italy, PhD, 2025), **Hassan Shahmohammadi** (University of Tübingen, Germany, PhD, 2024), **Nirmal Roy** (TU Delft, The Netherlands, PhD, 2024), **Jiechen Xu** (University of Queensland, Australia, PhD, 2024), **Tim Draws** (TU Delft, The Netherlands, PhD, 2023), **Alexander Bondarenko** (University of Jena, Germany, PhD, 2023), **Austin R. Ward** (University of North Carolina at Chapel Hill, USA, PhD, 2023), **Amir Ilkhechi** (Brown University, USA, PhD, 2023), **Babak Hemmatian Borujeni** (Brown University, USA, PhD, 2021), **Rolf Jagerman** (University of Amsterdam, The Netherlands, PhD, 2020), **Anton van der Vegt** (University of Queensland, Australia, PhD, 2020), **Navid Rekabsaz** (TU Vienna, Austria, PhD, 2018)

Service to Community

Board Positions

ACM ICTIR, Steering Committee Member	2025 - 2028
ACM SIGIR, Treasurer	2022 - 2025
ECIR, Award Selection Committee	2023
Cross Language Eval. Forum (CLEF), Steering Committee Member	2020 - present
Northeast Big Data Innovation Hub, Steering Committee Member	2020 - present

Chair

SIGIR Explainability in IR Workshop, Co-Chair	2025
ECIR, Reproducibility Track Co-Chair	2022
CLEF Labs Organizing Committee, Co-Chair	2020
WSDM Health Search and Data Mining Workshop, Co-Chair	2020
ACM WSDM, Doctoral Consortium Chair	2019
AAAI HCOMP, Doctoral Consortium Chair	2018
ACM SIGIR, Proceedings Chair	2018
ImageCLEF Biomedical Captioning Track, Co-Chair	2017 - 2018
Dutch-Belgian Information Retrieval Workshop (DIR), General Chair	2013
CIKM BooksOnline Workshop, Co-Chair	2011 - 2012

Editor

Information Retrieval, Associate Editor	2023 - present
Frontiers in Public Health, Associate Editor	2020 - present
Frontiers in Computer Science, Review Editor	2020 - present
Frontiers in Computer Science, Guest Editor	2022
Frontiers in Pharmacology, Guest Editor	2021
MDPI Sensors, Guest Editor	2022
Springer Information Retrieval Journal, Guest Editor	2017

Senior Area Chair
European Chapter of the Association for Computational Linguistics (EACL)

Senior Program Committee Member
ACM Conference on Research & Development in IR (SIGIR)
Association for Computational Linguistics Rolling Review (ACL)
The Web Conference (WWW)
ACM International Conference on Web Search & Data Mining (WSDM)
European Conference on Information Retrieval (ECIR)

Reviewer (selected)
Journal of the American Medical Association (JAMA), *International Conference on Machine Learning (ICML)*, *Neural Information Processing Systems (NeurIPS)*, *The Web Conference (WebConf)*, *ACM Conference on Knowledge Discovery & Data Mining (KDD)*, *AAAI Conference on Artificial Intelligence*, *Empirical Methods in Natural Language Processing (EMNLP)*, *Journal of Medical Internet Research (JMIR)*, *Journal of Biomedical Informatics (JBI)*, *Communications of the ACM (CACM)*, *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, *ACM Conference on Information & Knowledge Management (CIKM)*, *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, *ACM Computing Surveys (CSUR)*, *Methods in Medical Informatics*, *ACM Transactions on the Web (TWEB)*, *ACM Transactions on Information Systems (TOIS)*, *PLOS ONE*, *Cancer Medicine*, *AMIA Annual Symposium*, *AMIA Informatics Summit*, *IEEE ICHI Workshop on Health Natural Language Processing*, *ACM Transactions on Information Systems Technology (TIST)*, *Elsevier Information Processing & Management (IPM)*, *Elsevier Information Systems*, *IEEE International Symposium on Computer Based Medical Systems (CBMS)*, *International Workshop on Health Natural Language Processing (HealthNLP)*, *ACM SIGIR Workshop on Neural Information Retrieval (Neu-IR)*, *Asia Information Retrieval Societies Conference (AIRS)*, *Semantic Web Journal*, *IEEE Transactions on Affective Computing (TAC)*, *Symposium on String Processing and Information Retrieval (SPIRE)*, *Living Labs for Information Retrieval Evaluation at CIKM*, *Russian Information Retrieval Workshop (RussIR)*, *Workshop on Gamification for Information Retrieval (GamifIR)*, *Workshop on Crowdsourcing and Gamification in the Cloud (CGCloud)*, *Workshop on Searching 4 Fun at IIiX*, *Workshop on Context-awareness in Retrieval and Recommendation (CARR)*, *ECIR Workshop on Semantic Indexing and Information Retrieval for Health*

Review Panels
German Research Foundation (DFG)
National Science Foundation (NSF)
Swiss National Science Foundation (SNSF)
European Commission (EC)
Belgian National Science Foundation (BelSPO)
Dutch Organization for Scientific Research (NWO)
UK Research and Innovation (UKRI)

SERVICE TO UNIVERSITY

Board Positions

Data Science Initiative, Campus Advisory Board Member	2021 - 2024
Center for Digital Health, Research Lead for Big Data and Analytics	2020 - 2023
Hack for STEM Diversity, Steering Committee Member	2021

Chair

University of Tübingen, Data Use and Access Committee	2022 - present
Center for Biomedical Informatics, Events Committee	2019–2020

Committees

AdvanceCTR Pilot Projects	2021
Brown University, AI/ML in Medicine CoP	2021
Brown University, UTRA Selection Committee	2020
OVPR Seed Award Reviewer Committee	2019–2021
Center for Biomedical Informatics, Curriculum/Admissions Committee	2019–2020
Center for Biomedical Informatics, CTR Oversight Committee	2019–2020
Computer Science Department, PhD Admission Committee	2018–2022
Center for Computational Molecular Biology, PhD Admission Committee	2022

MEDIA

Frankfurter Allgemeine Zeitung, Piotr Heller.
"*Warum Chatbots keine Ärzte sind*" 10-15-2025
Frankfurter Allgemeine Zeitung, Piotr Heller.
"*Der Chatarzt*" 02-03-2025
Reutlinger Generalanzeiger, Miriam Steinrücken.
"*Chat-GPT as a Giant Opportunity*" 04-29-2023
SWR Aktuell, Kay-Uwe Hennig.
"*Automatic Text Generation with Chat-GPT. [German Original]*" 02-01-2023

Kronenzeitung, Erich Vogel.
"*Corona and the Healthcare System in Turmoil. [German Original]*" 09-21-2021
Tiroler Tageszeitung, Nicole Strozzi.
"*How Machines Change Medicine [German Original]*" 07-01-2021
E-HEALTH-COM, Philipp Grätzel.
"*Algorithms Predict Bleeding and Kidney Failure [German Original]*" 10-09-2018
Slate Magazine, Michael Agger.
"*Google Kids: The Sequel*" 06-14-2011

Professional Memberships

- **Association for Computational Linguistics (ACL)**, Professional Member *2019 - Present*
- **Swiss Society for Medical Informatics (SGMI)**, Full Member *2016 - Present*
- **Inst. Electrical and Electronics Engineers (IEEE)**, Professional Member *2015 - Present*
- **ACM Special Interest Group on IR (SIGIR)**, Full Member *2012 - Present*
- **Association for Computing Machinery (ACM)**, Professional Member *2011 - Present*