

1. Record Nr.	UNINA9910845480003321
Titolo	Advances in Information Retrieval : 46th European Conference on Information Retrieval, ECIR 2024, Glasgow, UK, March 24–28, 2024, Proceedings, Part III // edited by Nazli Goharian, Nicola Tonello, Yulan He, Aldo Lipani, Graham McDonald, Craig Macdonald, Iadh Ounis
Pubbl/distr/stampa	Cham : , : Springer Nature Switzerland : , : Imprint : Springer, , 2024
ISBN	3-031-56063-9
Edizione	[1st ed. 2024.]
Descrizione fisica	1 online resource (528 pages)
Collana	Lecture Notes in Computer Science, , 1611-3349 ; ; 14610
Disciplina	025.524
Soggetti	Information storage and retrieval systems Database management Data mining Machine learning Natural language processing (Computer science) Information Storage and Retrieval Database Management System Data Mining and Knowledge Discovery Machine Learning Natural Language Processing (NLP)
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di bibliografia	Includes bibliographical references and index.
Nota di contenuto	Intro -- Preface -- Organization -- Contents - Part III -- Full Papers -- Knowledge Graph Cross-View Contrastive Learning for Recommendation -- 1 Introduction -- 2 Related Work -- 3 Model Architecture -- 3.1 Problem Formulation -- 3.2 Local-Level Learning -- 3.3 Global-Level Learning -- 3.4 Cross-View Contrastive Learning -- 4 Experiments -- 4.1 Datasets and Evaluation Protocol -- 4.2 Baselines -- 4.3 Performance Comparison with Baselines (RQ1) -- 4.4 Ablation Study (RQ2) -- 4.5 Sensitivity Analysis (RQ3) -- 5 Conclusions -- References -- Mu2STS: A Multitask Multimodal Sarcasm-Humor Differential Teacher-Student Model for Sarcastic Meme Detection -- 1 Introduction -- 2 Related Works -- 3 Meme Corpus Creation -- 4

Methodology -- 4.1 Problem Formulation -- 4.2 Encoding of Meme -- 4.3 Multimodal Multisegment Multihead Fusion with Cross-Modal Understanding (M3F-CmU) Module -- 4.4 Humor-Aware Sarcasm Synced Teacher Training (HuS-STT) Module -- 4.5 Sarcasm-Illuminator Student Model Training (Sal-SMT) Module -- 5 Results and Analysis -- 5.1 Experimental Details -- 5.2 Comparison with Baselines -- 5.3 Ablation Study -- 5.4 Detailed Analysis -- 5.5 Cross-Lingual Generalization -- 5.6 Comparison with the State-of-the-Art (SOTA) Models -- 6 Error Analysis and Limitations -- 7 Conclusion and Future Work -- References

Robustness in Fairness Against Edge-Level Perturbations in GNN-Based Recommendation -- 1 Introduction -- 2 Related Work -- 2.1 Attacks and Robustness in Recommendation -- 2.2 Fairness in Recommendation -- 2.3 Robustness and Beyond-Accuracy Aspects -- 3 Methodology -- 3.1 Perturbation Task in Graph-Based Recommendation -- 3.2 Graph Perturbation Mechanism -- 3.3 Fairness Notion and Operationalization -- 4 Experimental Evaluation -- 4.1 Evaluation Setting -- 4.2 RQ1: Impact on Robustness in Fairness. 4.3 RQ2: Robustness in Fairness Under Incremental Perturbations -- 4.4 RQ3: Edge Perturbations Influence on Groups -- 5 Conclusions -- References

Is Google Getting Worse? A Longitudinal Investigation of SEO Spam in Search Engines -- 1 Introduction -- 2 Related Work -- 3 Data and Feature Extraction -- 4 SEO Spam in Product Reviews -- 5 Temporal Analysis of Product Reviews -- 6 Conclusion -- References

GLAD: Graph-Based Long-Term Attentive Dynamic Memory for Sequential Recommendation -- 1 Introduction -- 2 Methodology -- 2.1 Interaction Sequence Encoder -- 2.2 Dynamic Graph Based Memory -- 2.3 Output Gating -- 2.4 Attentive Reading of Graph Memory -- 2.5 End-to-End Training -- 3 Related Work -- 4 Experiments -- 4.1 Discussion -- 5 Conclusion -- References

Effective Adhoc Retrieval Through Traversal of a Query-Document Graph -- 1 Introduction -- 2 Related Work -- 3 Query-Document Graph -- 3.1 Query Population -- 3.2 Ranking Neighbours -- 4 QDG in an Adaptive Retrieval Framework -- 4.1 Reverted Adaptive Retrieval -- 4.2 Resource Selection -- 5 Experiments -- 5.1 Experimental Setup -- 5.2 Results -- 6 Qualitative Analysis -- 7 Conclusions -- References

How to Forget Clients in Federated Online Learning to Rank? -- 1 Introduction -- 2 Related Work -- 3 Methodology -- 3.1 Preliminary -- 3.2 Unlearning in FOLTR -- 3.3 Efficiency Analysis -- 3.4 Evaluating Unlearning -- 4 Experimental Setup -- 5 Results and Analysis -- 5.1 Validation of Evaluation Methodology -- 5.2 Effectiveness of Unlearning -- 5.3 Hyper-parameters Analysis -- 6 Conclusion -- References

An Adaptive Feature Selection Method for Learning-to-Enumerate Problem -- 1 Introduction -- 2 Related Work -- 3 Proposed Method -- 4 Experiments -- 4.1 Experimental Result -- 4.2 Analysis on Why AdaFeaSE Works -- 4.3 Lasso with Initial Weight Values Being 1 -- 4.4 Scalability of the Algorithm. 5 Conclusions -- References

A Transformer-Based Object-Centric Approach for Date Estimation of Historical Photographs -- 1 Introduction -- 2 Related Work -- 3 DEW-B Dataset -- 4 Method -- 4.1 Single Image Date Estimation Model -- 4.2 Leveraging Object Detection for Date Estimation -- 4.3 DEXPERT Architecture -- 5 Experiments -- 6 Conclusions -- A Implementation details -- References

Shallow Cross-Encoders for Low-Latency Retrieval -- 1 Introduction -- 2 Efficiency/Effectiveness Tradeoffs in Cross-Encoders -- 3 Training Shallow Cross-Encoders with gBCE -- 4 Experiments -- 4.1 Experimental Setup -- 4.2 RQ1. Efficiency/Effectiveness Tradeoffs -- 4.3 RQ2. Effect of the Training Scheme -- 4.4 RQ3. CPU Inference -- 5 Conclusion -- A Effect of gBCE Training Scheme on Tiny BERT-Based

Cross-Encoder on the MS MARCO Dev Set -- References -- Asking Questions Framework for Oral History Archives -- 1 Introduction -- 2 Speech Recognition for Oral Archives -- 3 Asking Questions Framework -- 3.1 Asking Questions Model -- 3.2 Semantic Continuity Model -- 4 Experimental Evaluation -- 4.1 Semantic Continuity Model Evaluation -- 4.2 AQ Models Comparison -- 5 Asking Questions Demonstrator -- 6 Conclusion and Future Work -- References -- Improved Learned Sparse Retrieval with Corpus-Specific Vocabularies -- 1 Introduction -- 2 Background and Related Work -- 3 Corpus-Specific Vocabularies -- 3.1 Vocabulary Selection -- 3.2 Pre-training Objectives -- 3.3 Sparse Retrieval Models -- 3.4 Document Expansion -- 3.5 Distillation-Based Training -- 4 Experiments -- 4.1 Setup -- 4.2 Vocabulary Selection and Index Statistics -- 4.3 Retrieval Quality -- 4.4 Query Latency -- 4.5 Pre-training Cost and Model Size -- 5 Conclusion and Future Work -- References -- Interactive Topic Tagging in Community Question Answering Platforms -- 1 Introduction -- 2 Related Work. 3 Proposed Method -- 3.1 Problem Definition -- 3.2 Question and Tag Embedding -- 3.3 Question Tagging Model -- 3.4 Feedback-Aware Question Tagging -- 4 Experiments -- 4.1 Dataset -- 4.2 Baselines -- 4.3 Evaluation Strategy -- 4.4 Results and Findings -- 5 Concluding Remarks -- References -- Yes, This Is What I Was Looking For! Towards Multi-modal Medical Consultation Concern Summary Generation -- 1 Introduction -- 2 Related Works -- 3 Dataset -- 3.1 Data Collection and Annotation -- 3.2 Qualitative Aspects -- 4 Methodology -- 4.1 Multi-modal Feature Extraction -- 4.2 Contextualized M-Modality Fusion -- 4.3 MMCS Generation -- 4.4 Experimental Details -- 5 Result and Discussion -- 5.1 Experimental Results -- 5.2 Findings to Research Questions -- 6 Analysis -- 7 Conclusion -- 8 Ethical Consideration -- References -- Mitigating Data Sparsity via Neuro-Symbolic Knowledge Transfer -- 1 Introduction -- 2 Related Works -- 3 Background -- 3.1 Notation -- 3.2 Matrix Factorization -- 3.3 Factorization Machine -- 3.4 Logic Tensor Networks -- 4 Method -- 4.1 Real Logic Knowledge Base -- 4.2 Grounding of Symbols -- 4.3 Learning of the LTN -- 5 Experiments -- 5.1 Datasets -- 5.2 Experimental Setting -- 5.3 Training Details -- 6 Results -- 7 Conclusions and Future Works -- References -- InDi: Informative and Diverse Sampling for Dense Retrieval -- 1 Introduction -- 2 Preliminaries and Notation -- 3 Methods -- 3.1 Informative and Diverse Sampling -- 3.2 In-Batch Deduplication -- 4 Experimental Setup -- 5 Results -- 5.1 Main Results -- 5.2 Fine-Grained Topical Understanding -- 5.3 Additional Dataset -- 5.4 Diversity and Informativeness Importance -- 5.5 Diversity and Informativeness Measures -- 5.6 Ablation -- 5.7 Cross Encoder Filtering with InDi -- 5.8 Computational Surplus -- 6 Related Work -- 6.1 Retrieval -- 6.2 Negative Selection for Dense Retrieval. 6.3 Active Learning -- 7 Conclusions -- References -- Short Papers -- Enhancing Legal Named Entity Recognition Using RoBERTa-GCN with CRF: A Nuanced Approach for Fine-Grained Entity Recognition -- 1 Introduction -- 2 Related Work -- 3 Dataset -- 4 Methodology -- 5 Experimental Setup -- 6 Results -- 7 Conclusion -- References -- A Novel Multi-Stage Prompting Approach for Language Agnostic MCQ Generation Using GPT -- 1 Introduction and Background -- 2 Method -- 3 Experiments -- 4 Conclusion and Future Work -- References -- An Adaptive Framework of Geographical Group-Specific Network on O2O Recommendation -- 1 Introduction -- 2 Method -- 2.1 Framework -- 2.2 Implementation of Group-Specification -- 2.3 Algorithm -- 3 Experiment -- 3.1 Experimental Configurations -- 3.2 Offline Experiment and Sensitivity Analysis -- 3.3 Online A/B Test and Ablation Test -- 4 Conclusion -- References -- A Study on Hierarchical

Text Classification as a Seq2seq Task -- 1 Introduction -- 2 Analysis Framework -- 2.1 Sequence Modeling -- 2.2 Decoding Strategies -- 2.3 Proposed Metrics -- 3 Experiment and Datasets -- 3.1 Datasets and Implementation Details -- 3.2 Results -- 4 Conclusion -- References -- Improving the Robustness of Dense Retrievers Against Typos via Multi-Positive Contrastive Learning -- 1 Introduction -- 2 Methodology -- 2.1 Dense Retriever with Self-supervised Contrastive Learning -- 2.2 Dense Retriever with Dual Learning -- 2.3 Dense Retriever with Dual Learning and Self-Teaching -- 3 Experimental Setup -- 4 Results -- 5 Conclusions -- References -- Learning-to-Rank with Nested Feedback -- 1 Introduction and Related Work -- 2 Modeling Nested Ranking Signals -- 3 Experimental Validation -- 3.1 Offline Training and Experiment Results -- 3.2 Online A/B Experimentation -- 4 Discussion and Future Work -- References. MMCRec: Towards Multi-modal Generative AI in Conversational Recommendation.

Sommario/riassunto

The six-volume set LNCS 14608, 14609, 14609, 14610, 14611, 14612 and 14613 constitutes the refereed proceedings of the 46th European Conference on IR Research, ECIR 2024, held in Glasgow, UK, during March 24–28, 2024. The 57 full papers, 18 finding papers, 36 short papers, 26 IR4Good papers, 18 demonstration papers, 9 reproducibility papers, 8 doctoral consortium papers, and 15 invited CLEF papers were carefully reviewed and selected from 578 submissions. The accepted papers cover the state of the art in information retrieval focusing on user aspects, system and foundational aspects, machine learning, applications, evaluation, new social and technical challenges, and other topics of direct or indirect relevance to search. .
