

1. Record Nr.	UNISA996464406003316
Titolo	Advances in information retrieval : 43rd European Conference on IR Research, ECIR 2021, virtual event, March 28-April 1, 2021, proceedings, part II // Djoerd Hiemstra [and five others] (editors)
Pubbl/distr/stampa	Cham, Switzerland : , : Springer, , [2021] ©2021
ISBN	3-030-72240-6
Descrizione fisica	1 online resource (760 pages)
Collana	Lecture Notes in Computer Science ; ; 12657
Disciplina	004
Soggetti	Computers
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Intro -- Preface -- Organization -- Contents - Part II -- Contents - Part I -- Reproducibility Track Papers -- Cross-Domain Retrieval in the Legal and Patent Domains: A Reproducibility Study -- 1 Introduction -- 2 Methods -- 2.1 Task Description -- 2.2 BERT-PLI Architecture Overview -- 2.3 Cross-Domain Evaluation Approach -- 3 Experiments -- 3.1 Datasets -- 3.2 Experiment Setting -- 4 Evaluation and Analysis -- 4.1 In-Domain Evaluation for Legal Document Retrieval (RQ1) -- 4.2 In-domain Evaluation for Patent Document Retrieval (RQ2) -- 4.3 Cross-Domain Evaluation (RQ3) -- 5 Related Work -- 6 Conclusion and Future Work -- References -- A Critical Assessment of State-of-the-Art in Entity Alignment -- 1 Introduction -- 2 Datasets and Initialization -- 2.1 Datasets -- 2.2 Label-Based Initializations -- 3 Methods -- 3.1 Relation-Aware Dual-Graph Convolutional Network (RDGCN) -- 3.2 Deep Graph Matching Consensus (DGMC). -- 4 Experiments -- 4.1 Results -- 4.2 Ablation: RDGCN -- 5 Conclusion -- References -- System Effect Estimation by Sharding: A Comparison Between ANOVA Approaches to Detect Significant Differences -- 1 Introduction -- 2 Related Work -- 3 Approach -- 3.1 ANOVA Models -- 3.2 Bootstrap ANOVA (bANOVA) -- 3.3 Experimental Setup -- 4 Replicability of bANOVA -- 5 Generalizability of tANOVA and bANOVA -- 5.1 Impact of the Multiple Comparison Strategies and Bootstrapping -- 5.2 Effect of the Random Shards on the Stability of the Approaches

-- 5.3 Stability of ANOVA Models with Respect to Different Interpolation Values -- 6 Conclusions and Future Work -- References

-- Reliability Prediction for Health-Related Content: A Replicability Study -- 1 Introduction -- 2 Related Work -- 3 Dataset -- 3.1 Features -- 4 Experimental Setup -- 5 Results -- 6 Generalisation -- 6.1 Standardisation -- 6.2 New Test Datasets -- 6.3 Generalisation Conclusions.

7 Future Work -- 8 Conclusions -- References -- An Empirical Comparison of Web Page Segmentation Algorithms -- 1 Introduction -- 2 Related Work -- 3 Experiment Setup -- 4 Algorithms and Parameter Tuning -- 4.1 VIPS -- 4.2 HEPS -- 4.3 Cormier et al. -- 4.4 MMDetection -- 4.5 Meier et al. -- 4.6 Min-Vote@n -- 4.7 Baseline -- 5 Results of the Comparison -- 6 Conclusion -- References -- Re-assessing the "Classify and Count" Quantification Method -- 1 Introduction -- 2 "Classify and Count" and Its Variants -- 3 Quantification and Parameter Optimisation -- 3.1 Unsuitable Parameter Optimisation and Weak Baselines -- 3.2 Quantification-Oriented Parameter Optimisation -- 4 Experiments -- 4.1 Evaluation Measures -- 4.2 Data Processing -- 4.3 The Quantifiers -- 4.4 Results -- 5 Conclusions -- References -- Reproducibility, Replicability and Beyond: Assessing Production Readiness of Aspect Based Sentiment Analysis in the Wild -- 1 Introduction -- 2 Desiderata and Evaluation Rubric -- 3 Related Work -- 4 Dataset -- 5 Models Compared -- 6 Experimental Setup -- 7 Results and Discussion: Production Readiness Rubrics -- 7.1 Reproducibility and Replicability -- 7.2 Performance on the Hard and Contrast Data Slices -- 7.3 Transferability Rubric: Cross Domain Experiments -- 7.4 Summary Comparison of the Different Models Under the Production Readiness Rubrics -- 7.5 Limitations of the Present Study -- 8 Conclusion -- References -- Robustness of Meta Matrix Factorization Against Strict Privacy Constraints -- 1 Introduction -- 2 Methodology -- 2.1 Approach -- 2.2 Datasets -- 2.3 Recommendation Evaluation -- 2.4 Source Code and Materials -- 3 Reproducibility Study -- 4 Privacy-Focused Study -- 5 Conclusions and Future Work -- References -- Textual Characteristics of News Title and Body to Detect Fake News: A Reproducibility Study -- 1 Introduction. 2 Overview of the Paper by Horne and Adali -- 2.1 Approach -- 2.2 Features -- 2.3 Observation and Evaluation -- 3 Reproducibility -- 3.1 Datasets -- 3.2 Features -- 3.3 Analysis -- 3.4 Results -- 4 How to Reproduce Our Experiments -- 5 Conclusions -- References -- Federated Online Learning to Rank with Evolution Strategies: A Reproducibility Study -- 1 Introduction -- 2 Federated OLTR with Evolution Strategies -- 2.1 Federated Learning Optimization Setting -- 2.2 Gradient Estimation -- 2.3 Privatization Procedure -- 3 Experimental Settings -- 3.1 Datasets -- 3.2 Simulation -- 3.3 Evaluation Metric -- 3.4 FOLtR-ES and Comparison OLTR Methods -- 4 Results and Analysis -- 4.1 RQ1: Generalisation of FOLtR-ES Performance Beyond MQ2007/2008 -- 4.2 RQ2: Effect of Number of Clients on FOLtR-ES -- 4.3 RQ3: Comparing FOLtR-ES to State-of-the-Art OLTR Methods -- 4.4 RQ4: Extending FOLtR-ES Evaluation to Common OLTR Practice -- 5 Related Work -- 6 Conclusions -- References -- Comparing Score Aggregation Approaches for Document Retrieval with Pretrained Transformers -- 1 Introduction -- 2 Related Work -- 2.1 Passage Aggregation -- 2.2 BERT Variants -- 3 Experimental Setup -- 3.1 BERT with MaxP, FirstP, SumP, and AvgP Aggregation -- 3.2 BERT Variants -- 3.3 Deep and Shallow Sampling -- 3.4 Experimental Details -- 4 Results and Discussion -- 4.1 Reproduction and Generalization of Aggregation Approaches -- 4.2 MaxP with BERT Variants -- 4.3 Deep vs. Shallow Relevance Judgments

-- 5 Conclusion -- References -- Short Papers -- Transformer-Based Approach Towards Music Emotion Recognition from Lyrics -- 1 Introduction -- 2 Related Work -- 3 Methodology -- 3.1 Datasets -- 3.2 Lyrics Extraction -- 3.3 Proposed Architecture -- 4 Experiment and Results -- 4.1 Evaluation Measures -- 4.2 Results -- 5 Conclusion -- References.

BiGBERT: Classifying Educational Web Resources for Kindergarten-12th Grades -- 1 Introduction -- 2 BiGBERT -- 3 Experiments and Discussion -- 4 Conclusion and Future Work -- References -- How Do Users Revise Zero-Hit Product Search Queries? -- 1 Introduction -- 2 Related Work -- 3 Collecting Zero-Hit and Successful Queries with Similar Intents -- 4 Four Revision Strategies for Zero-Hit Product Queries -- 5 Coverage of the Four Revision Strategies -- 6 Conclusions and Future Research -- References -- Query Performance Prediction Through Retrieval Coherency -- 1 Introduction -- 2 Proposed Approach -- 3 Experiments -- 4 Concluding Remarks -- References -- From the Beatles to Billie Eilish: Connecting Provider Representativeness and Exposure in Session-Based Recommender Systems -- 1 Introduction -- 2 Metrics and Algorithms -- 3 Experiments -- 3.1 Data and Algorithms -- 3.2 Results -- 4 Conclusions -- References -- Bayesian System Inference on Shallow Pools -- 1 Introduction -- 2 Related Work -- 3 Statistical Models -- 4 Analysis -- 5 Conclusion -- References -- Exploring Gender Biases in Information Retrieval Relevance Judgement Datasets -- 1 Introduction -- 2 Methodology -- 2.1 Datasets -- 2.2 Query Gender Identification and Labeling -- 2.3 Quantifying Psychological Characteristics -- 3 Findings -- 4 Concluding Remarks -- References -- Assessing the Benefits of Model Ensembles in Neural Re-ranking for Passage Retrieval -- 1 Introduction -- 2 Passage Re-ranking with Neural Ensembles -- 2.1 Fast Geometric Model Ensembling -- 2.2 Rank Fusion Methods -- 3 Neural Ranking Models -- 4 Experimental Evaluation -- 5 Conclusions and Future Work -- References -- Event Detection with Entity Markers -- 1 Introduction -- 2 Related Work -- 3 Approach -- 4 Experiments and Results -- 5 Conclusions and Perspectives -- References.

Simplified TinyBERT: Knowledge Distillation for Document Retrieval -- 1 Introduction -- 2 Background -- 3 Simplified TinyBERT for Ranking -- 3.1 Method -- 3.2 Implementation Details -- 4 Experiments -- 4.1 Experimental Setup -- 4.2 Results -- 5 Conclusion -- References -- Improving Cold-Start Recommendation via Multi-prior Meta-learning -- 1 Introduction -- 2 Methodology -- 2.1 Framework of MPML -- 2.2 Accuracy-Based Task Clustering -- 2.3 Meta Update -- 2.4 Complexity Analysis -- 3 Experiment -- 3.1 Experimental Setup -- 3.2 Experimental Results -- 4 Conclusion -- References -- A White Box Analysis of ColBERT -- 1 Introduction -- 2 Experimental Setting -- 3 ColBERT Term Importance -- 4 Analysis of Exact and Soft Matches -- 5 Conclusion -- References -- Diversity Aware Relevance Learning for Argument Search -- 1 Introduction -- 2 Preliminaries -- 3 Our Approach for Candidate Refinement -- 3.1 Relevance Filter -- 3.2 Premise Ranker -- 4 Evaluation -- 5 Conclusion -- References -- SQE-GAN: A Supervised Query Expansion Scheme via GAN -- 1 Introduction -- 2 Framework -- 3 Applying SQE-GAN in the Query -- 4 Experiments -- 5 Conclusions -- References -- Rethink Training of BERT Rerankers in Multi-stage Retrieval Pipeline -- 1 Introduction -- 2 Background -- 3 Methodologies -- 4 Experiment Methodologies -- 5 Document Ranking Performance -- 6 Analysis -- 7 Conclusion -- References -- Should I Visit This Place? Inclusion and Exclusion Phrase Mining from Reviews -- 1 Introduction -- 2 Related Work -- 3 Proposed Approach -- 3.1 Dataset -- 3.2 Methods -- 4 Experiments -- 4.1 Results -- 4.2

Error Analysis -- 5 Conclusion -- References -- Dynamic Cross-Sentential Context Representation for Event Detection -- 1 Introduction -- 2 Model -- 2.1 Intra-sentential GCN -- 2.2 Cross-Sentential Context Representation -- 3 Experiments -- 3.1 Data and Preprocessing.
3.2 Study of Model's Parameters.
