

1. Record Nr.	UNINA9910886100803321
Autore	Bifet Albert
Titolo	Machine Learning and Knowledge Discovery in Databases. Research Track : European Conference, ECML PKDD 2024, Vilnius, Lithuania, September 9–13, 2024, Proceedings, Part I // edited by Albert Bifet, Jesse Davis, Tomas Krilaviius, Meelis Kull, Eirini Ntoutsis, InDR Žliobaitis
Pubbl/distr/stampa	Cham : , : Springer Nature Switzerland : , : Imprint : Springer, , 2024
ISBN	3-031-70341-3
Edizione	[1st ed. 2024.]
Descrizione fisica	1 online resource (514 pages)
Collana	Lecture Notes in Artificial Intelligence, , 2945-9141 ; ; 14941
Altri autori (Persone)	DavisJesse KrilaviiusTomas KullMeelis NtoutsisEirini ŽliobaitisInDR
Disciplina	006.3
Soggetti	Artificial intelligence Computer engineering Computer networks Computers Image processing - Digital techniques Computer vision Software engineering Artificial Intelligence Computer Engineering and Networks Computing Milieux Computer Imaging, Vision, Pattern Recognition and Graphics Software Engineering
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Intro -- Preface -- Organization -- Invited Talks Abstracts -- The Dynamics of Memorization and Unlearning -- The Emerging Science of Benchmarks -- Enhancing User Experience with AI-Powered Search and Recommendations at Spotify -- How to Utilize (and Generate)

Player Tracking Data in Sport -- Resource-Aware Machine Learning-
A User-Oriented Approach -- Contents - Part I -- Research Track --
Adaptive Sparsity Level During Training for Efficient Time Series
Forecasting with Transformers -- 1 Introduction -- 2 Background --
2.1 Sparse Neural Networks -- 2.2 Time Series Forecasting -- 2.3
Problem Formulation and Notations -- 3 Analyzing Sparsity Effect in
Transformers for Time Series Forecasting -- 4 Proposed Methodology:
PALS -- 5 Experiments and Results -- 5.1 Experimental Settings -- 5.2
Results -- 6 Discussion -- 6.1 Performance Comparison with Pruning
and Sparse Training Algorithms -- 6.2 Hyperparameter Sensitivity -- 7
Conclusions -- References -- RumorMixer: Exploring Echo Chamber
Effect and Platform Heterogeneity for Rumor Detection -- 1
Introduction -- 2 Related Works -- 3 Methodology -- 3.1 Overview --
3.2 Echo Chamber Extraction and Representation Learning -- 3.3
Neural Architecture Search for Platform Heterogeneity -- 4 Experiments
-- 4.1 Experimental Setting -- 4.2 Performance Comparison (RQ1) --
4.3 Ablation Study (RQ2) -- 4.4 Parameter Analysis (RQ3) -- 4.5 Early
Rumor Detection (RQ4) -- 5 Conclusion -- References -- Diversified
Ensemble of Independent Sub-networks for Robust Self-supervised
Representation Learning -- 1 Introduction -- 2 Related Work -- 3
Method -- 3.1 Robust Self-supervised Learning via Independent Sub-
networks -- 3.2 Empirical Analysis of Diversity -- 3.3 Computational
Cost and Efficiency Analysis -- 4 Experimental Setup -- 5 Results and
Discussion -- 6 Ablation Study -- 7 Conclusion -- References.
Modular Debiasing of Latent User Representations in Prototype-Based
Recommender Systems -- 1 Introduction -- 2 Related Work -- 3
Methodology -- 4 Experiment Setup -- 5 Results and Analysis -- 6
Conclusion and Future Directions -- References -- A Mathematics
Framework of Artificial Shifted Population Risk and Its Further
Understanding Related to Consistency Regularization -- 1 Introduction
-- 2 Related Work -- 3 Method -- 3.1 Revisiting Data Augmentation
with Empirical Risk -- 3.2 The Augmented Neighborhood -- 3.3 The
Artificial Shifted Population Risk -- 3.4 Understanding the
Decomposition of Shifted Population Risk -- 4 Experiment -- 4.1
Experiment Implementation -- 4.2 Experimental Results -- 5
Conclusion and Discussion -- References -- Attention-Driven Dropout:
A Simple Method to Improve Self-supervised Contrastive Sentence
Embeddings -- 1 Introduction -- 2 Background and Related Work -- 3
Method -- 3.1 Attention Rollout Aggregation -- 3.2 Static Dropout Rate
-- 3.3 Dynamic Dropout Rate -- 4 Experiment -- 4.1 Datasets and
Tasks -- 4.2 Training Procedure -- 5 Result and Discussion -- 5.1
Ablation Study -- 6 Conclusion -- References -- AEMLO: AutoEncoder-
Guided Multi-label Oversampling -- 1 Introduction -- 1.1 Research
Goal -- 1.2 Motivation -- 1.3 Summary -- 2 Related Work -- 2.1 Multi-
label Classification -- 2.2 Multi-label Imbalance Learning -- 2.3 Deep
Sampling Method -- 3 Multi-label AutoEncoder Oversampling -- 3.1
Method Description and Overview -- 3.2 Loss Function -- 3.3 Generate
Instances and Post-processing -- 4 Experiments and Analysis -- 4.1
Datasets -- 4.2 Experiment Setup -- 4.3 Experimental Analysis -- 4.4
Parameter Analysis -- 4.5 Sampling Time -- 5 Conclusion --
References -- MANTRA: Temporal Betweenness Centrality
Approximation Through Sampling -- 1 Introduction -- 2 Related Work
-- 3 Preliminaries.
4 MANTRA: Temporal Betweenness Centrality Approximation Through
Sampling -- 4.1 Temporal Betweenness Estimator -- 4.2 Sample
Complexity Bounds -- 4.3 Fast Approximation of the Characteristic
Quantities -- 4.4 The MANTRA Framework -- 5 Experimental
Evaluation -- 5.1 Experimental Setting -- 5.2 Networks -- 5.3

Experimental Results -- 6 Conclusions -- References --

Dimensionality-Induced Information Loss of Outliers in Deep Neural Networks -- 1 Introduction -- 2 Problem Setting and Related Work -- 2.1 Stable Rank of the Matrix -- 2.2 Feature-Based Detection -- 2.3 Projection-Based Detection -- 2.4 Similarity of DNN Representations -- 2.5 Noise Sensitivity in the DNN -- 3 Results -- 3.1 Overview of the Experiments and a Possible Picture -- 3.2 Observation of Dimensionality via Stable Ranks -- 3.3 Transition of OOD Detection Performance -- 3.4 Block Structure of CKA -- 3.5 Instability of OOD Samples to Noise Injection -- 3.6 Dataset Bias-Induced Imbalanced Inference -- 3.7 Quantitative Comparison of OOD Detection Performance -- 4 Discussion -- 5 Summary and Conclusion --

References -- Towards Open-World Cross-Domain Sequential Recommendation: A Model-Agnostic Contrastive Denoising Approach -- 1 Introduction -- 2 Methodology -- 2.1 Problem Formulation -- 2.2 Embedding Encoder -- 2.3 Denoising Interest-Aware Network -- 2.4 Fusion Gate Unit -- 2.5 Model Training -- 2.6 Inductive Representation Generator -- 3 Experiments -- 3.1 Datasets -- 3.2 Experiment Setting -- 3.3 Performance Comparisons (RQ1) -- 3.4 Ablation Study (RQ2) -- 3.5 Online Evaluation (RQ3) -- 3.6 Model Analyses (RQ4) -- 3.7 Parameter Sensitivity (RQ5) -- 4 Related Work -- 5 Conclusion --

References -- MixerFlow: MLP-Mixer Meets Normalising Flows -- 1 Introduction -- 2 Related Works -- 3 Preliminaries -- 4 MixerFlow Architecture and Its Components -- 5 Experiments. 5.1 Density Estimation on 3232 Datasets -- 5.2 Density Estimation on 6464 Datasets -- 5.3 Enhancing MAF with the MixerFlow -- 5.4 Datasets with Specific Permutations -- 5.5 Hybrid Modelling -- 5.6 Integration of Powerful Architecture -- 6 Conclusion and Limitations -- 7 Future Work and Broader Impact --

References -- Handling Delayed Feedback in Distributed Online Optimization: A Projection-Free Approach -- 1 Introduction -- 1.1 Our Contribution -- 1.2 Related Work -- 2 Projection-Free Algorithms Under Delayed Feedback -- 2.1 Preliminaries -- 2.2 Centralized Algorithm -- 2.3 Distributed Algorithm -- 3 Numerical Experiments -- 4 Concluding Remarks --

References -- Secure Dataset Condensation for Privacy-Preserving and Efficient Vertical Federated Learning -- 1 Introduction -- 2 Related Work -- 2.1 Vertical Federated Learning -- 2.2 Privacy Protection in VFL -- 2.3 Dataset Size Reduction in FL -- 3 Preliminaries -- 3.1 Problem Formulation -- 3.2 Dataset Condensation -- 3.3 Secure Aggregation -- 3.4 Differential Privacy -- 4 Proposed Approach -- 4.1 Overview -- 4.2 Class-Wise Secure Aggregation -- 4.3 VFDC Algorithm -- 4.4 Privacy Analysis -- 5 Experimental Study -- 5.1 Experimental Setup -- 5.2 Visualization of Condensed Dataset -- 5.3 Performance Comparison -- 5.4 Efficiency Improvement -- 5.5 Impact of Hyperparameters -- 6 Conclusion and Future Directions --

References -- Neighborhood Component Feature Selection for Multiple Instance Learning Paradigm -- 1 Introduction -- 2 Methods -- 2.1 The Lazy Learning Approach for Multiple Instance Learning Setting -- 2.2 Neighborhood Component Feature Selection for Single Instance Learning Setting -- 2.3 Our Proposal: Neighborhood Component Feature Selection for the Multiple Instance Learning Setting -- 3 Datasets -- 3.1 Musk Dataset -- 3.2 DEAP Dataset -- 4 Experimental Procedure -- 5 Experimental Results. 5.1 Musk Dataset -- 5.2 DEAP Dataset -- 5.3 Comparison with State-of-the-Art -- 5.4 Statistical Significance -- 5.5 Computational Complexity -- 6 Conclusions --

References -- MESS: Coarse-Grained Modular Two-Way Dialogue Entity Linking Framework -- 1 Introduction -- 2 Related Work -- 2.1 Mention-to-Entities -- 2.2 Transferred EL -- 3 Our MESS Framework -- 3.1 M2E Module -- 3.2 E2M Module -- 3.3

SS Module -- 3.4 Dialogue Module -- 4 Experiments -- 4.1 Setting -- 4.2 Results -- 4.3 Ablation Studies -- 5 Conclusion -- References -- Session Target Pair: User Intent Perceiving Networks for Session-Based Recommendation -- 1 Introduction -- 2 Related Work -- 3 Methodology -- 3.1 Problem Statement -- 3.2 Session-Level Intent Representation Module -- 3.3 Target-Level Intent Representation Module -- 3.4 Intent Alignment Mechanism Module -- 3.5 Prediction and Training -- 4 Experiments -- 4.1 Experiment Setups -- 4.2 Overall Performance -- 4.3 Model Analysis and Discussion -- 5 Conclusion -- References -- Hierarchical Fine-Grained Visual Classification Leveraging Consistent Hierarchical Knowledge -- 1 Introduction -- 2 Related Work -- 2.1 Fine-Grained Visual Classification -- 2.2 Hierarchical Multi-granularity Classification -- 2.3 Graph Representation Learning -- 3 Approach -- 3.1 Problem Setting -- 3.2 Multi-granularity Graph Convolutional Neural Network -- 3.3 Hierarchy-Aware Conditional Supervised Learning -- 3.4 Loss Function -- 3.5 Tree-Structured Granularity Consistency Rate -- 4 Experiments -- 4.1 Datasets -- 4.2 Experimental Settings -- 4.3 Compared Methods -- 4.4 Ablation Study -- 4.5 Comparison with State-of-the-Art Method -- 4.6 Qualitative Analysis -- 5 Conclusion -- References -- Backdoor Attacks with Input-Unique Triggers in NLP -- 1 Introduction -- 2 Related Work -- 3 Method -- 3.1 Problem Formulation -- 3.2 NURA: Input-Unique Backdoor Attack. 3.3 Model Training.

Sommario/riassunto

This multi-volume set, LNAI 14941 to LNAI 14950, constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2024, held in Vilnius, Lithuania, in September 2024. The papers presented in these proceedings are from the following three conference tracks: - Research Track: The 202 full papers presented here, from this track, were carefully reviewed and selected from 826 submissions. These papers are present in the following volumes: Part I, II, III, IV, V, VI, VII, VIII. Demo Track: The 14 papers presented here, from this track, were selected from 30 submissions. These papers are present in the following volume: Part VIII. Applied Data Science Track: The 56 full papers presented here, from this track, were carefully reviewed and selected from 224 submissions. These papers are present in the following volumes: Part IX and Part X.
