

1. Record Nr.	UNINA9910886089803321
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 VI // edited by Albert Bifet, Jesse Davis, Tomas Krilaviius, Meelis Kull, Eirini Ntoutsis, InDR Žilobaitis
Pubbl/distr/stampa	Cham : , : Springer Nature Switzerland : , : Imprint : Springer, , 2024
ISBN	3-031-70365-0
Edizione	[1st ed. 2024.]
Descrizione fisica	1 online resource (509 pages)
Collana	Lecture Notes in Artificial Intelligence, , 2945-9141 ; ; 14946
Altri autori (Persone)	DavisJesse KrilaviiusTomas KullMeelis NtoutsisEirini ŽilobaitisInDR
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 VI -- Research Track --
Rejection Ensembles with Online Calibration -- 1 Introduction -- 2
Notation and Related Work -- 2.1 Related Work -- 3 A Theoretical
Investigation of Rejection -- 3.1 Three Distinct Situations Can Occur
When Training the Rejector -- 3.2 Even a Perfect Rejector Will Overuse
Its Budget -- 3.3 A Rejector Should Not Trust f_s and f_b -- 4 Training a
Rejector for a Rejection Ensemble -- 5 Experiments -- 5.1 Experiments
with Deep Learning Models -- 5.2 Experiments with Decision Trees --
5.3 Conclusion from the Experiments -- 6 Conclusion -- References --
Lighter, Better, Faster Multi-source Domain Adaptation with Gaussian
Mixture Models and Optimal Transport -- 1 Introduction -- 2
Preliminaries -- 2.1 Gaussian Mixtures -- 2.2 Domain Adaptation --
2.3 Optimal Transport -- 3 Methodological Contributions -- 3.1 First
Order Analysis of MW2 -- 3.2 Supervised Mixture-Wasserstein
Distances -- 3.3 Mixture Wasserstein Barycenters -- 3.4 Multi-source
Domain Adaptation Through GMM-OT -- 4 Experiments -- 4.1 Toy
Example -- 4.2 Multi-source Domain Adaptation -- 4.3 Lighter, Better,
Faster Domain Adaptation -- 5 Conclusion -- References -- Subgraph
Retrieval Enhanced by Graph-Text Alignment for Commonsense
Question Answering -- 1 Introduction -- 2 Related Work -- 2.1
Commonsense Question Answering -- 2.2 Graph-Text Alignment -- 3
Task Formulation -- 4 Methods -- 4.1 Graph-Text Alignment -- 4.2
Subgraph Retrieval Module -- 4.3 Prediction -- 5 Experiments -- 5.1
Datasets.
5.2 Baselines -- 5.3 Implementation Details -- 5.4 Main Results -- 5.5
Ablation Study -- 5.6 Low-Resource Setting -- 5.7 Evaluation with
other GNNs -- 5.8 Hyper-parameter Analysis -- 6 Ethical
Considerations and Limitations -- 6.1 Ethical Considerations -- 6.2
Limitations -- 7 Conclusion -- References -- HetCAN: A
Heterogeneous Graph Cascade Attention Network with Dual-Level
Awareness -- 1 Introduction -- 2 Related Work -- 3 Preliminaries --
3.1 Heterogeneous Information Network -- 3.2 Graph Neural Networks
-- 3.3 Transformer-Style Architecture -- 4 The Proposed Model -- 4.1
Overall Architecture -- 4.2 Type-Aware Encoder -- 4.3 Dimension-
Aware Encoder -- 4.4 Time Complexity Analysis -- 5 Experiments --
5.1 Experimental Setups -- 5.2 Node Classification -- 5.3 Link
Prediction -- 5.4 Model Analysis -- 6 Conclusion -- References --
Interperable Target-Feature Aggregation for Multi-task Learning Based
on Bias-Variance Analysis -- 1 Introduction -- 2 Preliminaries -- 2.1
Related Works: Dimensionality Reduction, Multi-task Learning -- 3
Bias-Variance Analysis: Theoretical Results -- 4 Multi-task Learning via
Aggregations: Algorithms -- 5 Experimental Validation -- 5.1 Synthetic
Experiments and Ablation Study -- 5.2 Real World Datasets -- 6
Conclusions and Future Developments -- References -- The Simpler
The Better: An Entropy-Based Importance Metric to Reduce Neural
Networks' Depth -- 1 Introduction -- 2 Related Works -- 3 Method --
3.1 How Layers Can Degenerate -- 3.2 Entropy for Rectifier Activations
-- 3.3 EASIER -- 4 Experiments -- 4.1 Experimental Setup -- 4.2
Results -- 4.3 Ablation Study -- 4.4 Limitations and Future Work -- 5
Conclusion -- References -- Towards Few-Shot Self-explaining Graph
Neural Networks -- 1 Introduction -- 2 Problem Definition -- 3 The
Proposed MSE-GNN -- 3.1 Architecture of MSE-GNN -- 3.2
Optimization Objective.
3.3 Meta Training -- 4 Experiments -- 4.1 Datasets and Experimental
Setup -- 5 Related Works -- 6 Conclusion -- References -- Uplift
Modeling Under Limited Supervision -- 1 Introduction -- 2 Related
Work -- 3 Proposed Methodology -- 3.1 Uplift Modeling with Graph

Neural Networks (UMGNet) -- 3.2 Active Learning for Uplift GNNs (UMGNet-AL) -- 4 Experimental Evaluation -- 4.1 Datasets -- 4.2 Benchmark Models -- 4.3 Experiments -- 5 Conclusion -- References -- Self-supervised Spatial-Temporal Normality Learning for Time Series Anomaly Detection -- 1 Introduction -- 2 Related Work -- 3 STEN: Spatial-Temporal Normality Learning -- 3.1 Problem Statement -- 3.2 Overview of The Proposed Approach -- 3.3 OTN: Order Prediction-Based Temporal Normality Learning -- 3.4 DSN: Distance Prediction-Based Spatial Normality Learning -- 3.5 Training and Inference -- 4 Experiments -- 4.1 Experimental Setup -- 4.2 Main Results -- 4.3 Ablation Study -- 4.4 Qualitative Analysis -- 4.5 Sensitivity Analysis -- 4.6 Time Efficiency -- 5 Conclusion -- References -- Modeling Text-Label Alignment for Hierarchical Text Classification -- 1 Introduction -- 2 Related Work -- 3 Methodology -- 3.1 Text Encoder -- 3.2 Graph Encoder -- 3.3 Generation of Composite Representation -- 3.4 Loss Functions -- 4 Experiments -- 4.1 Datasets and Evaluation Metrics -- 4.2 Implementation Details -- 4.3 Experimental Results -- 4.4 Analysis -- 5 Conclusion -- A Details of Statistical Test -- B Performance Analysis on Additional Datasets -- References -- Secure Aggregation Is Not Private Against Membership Inference Attacks -- 1 Introduction -- 2 Related Work -- 3 Preliminaries -- 4 Privacy Analysis of Secure Aggregation -- 4.1 Threat Model -- 4.2 SecAgg as a Noiseless LDP Mechanism -- 4.3 Asymptotic Privacy Guarantee -- 4.4 Upper Bounding $M()$ via Dominating Pairs of Distributions. 4.5 Lower Bounding $M()$ and Upper Bounding $fM()$ via Privacy Auditing -- 5 Experiments and Discussion -- 6 Conclusions -- A Correlated Gaussian Mechanism -- A.1 Optimal LDP Curve: Proof of Theorem 2 -- A.2 The Case $S_d = \{x \in \mathbb{R}^d : \|x\|_2 \leq r\}$ -- A.3 Trade-Off Function: Proof of Proposition 1 -- B LDP Analysis of the Mechanism (1) in a Special Case: Proof of Theorem 3 -- References -- Evaluating Negation with Multi-way Joins Accelerates Class Expression Learning -- 1 Introduction -- 2 Preliminaries -- 2.1 The Description Logic ALC -- 2.2 Class Expression Learning -- 2.3 Semantics and Properties of SPARQL -- 2.4 Worst-Case Optimal Multi-way Join Algorithms -- 3 Mapping ALC Class Expressions to SPARQL Queries -- 4 Negation in Multi-way Joins -- 4.1 Rewriting Rule for Negation and UNION Normal Form -- 4.2 Multi-way Join Algorithm -- 4.3 Implementation -- 5 Experimental Results -- 5.1 Systems, Setup and Execution -- 5.2 Datasets and Queries -- 5.3 Results and Discussion -- 6 Related Work -- 7 Conclusion And Future Work -- References -- LayeredLiNGAM: A Practical and Fast Method for Learning a Linear Non-gaussian Structural Equation Model -- 1 Introduction -- 2 Related Work -- 3 Preliminaries -- 3.1 LiNGAM -- 3.2 DirectLiNGAM -- 4 LayeredLiNGAM -- 4.1 Generalization of Lemma 2 -- 4.2 Algorithm -- 4.3 Adaptive Thresholding -- 5 Experiments -- 5.1 Datasets and Evaluation Metrics -- 5.2 Determining Threshold Parameters -- 5.3 Results on Synthetic Datasets -- 5.4 Results on Real-World Datasets -- 6 Conclusion -- References -- Enhanced Bayesian Optimization via Preferential Modeling of Abstract Properties -- 1 Introduction -- 2 Background -- 2.1 Bayesian Optimization -- 2.2 Rank GP Distributions -- 3 Framework -- 3.1 Expert Preferential Inputs on Abstract Properties -- 3.2 Augmented GP with Abstract Property Preferences -- 3.3 Overcoming Inaccurate Expert Inputs. 4 Convergence Remarks -- 5 Experiments -- 5.1 Synthetic Experiments -- 5.2 Real-World Experiments -- 6 Conclusion -- References -- Enhancing LLM's Reliability by Iterative Verification Attributions with Keyword Fronting -- 1 Introduction -- 2 Related Work -- 2.1 Retrieval-Augmented Generation -- 2.2 Text Generation Attribution -- 3 Methodology -- 3.1 Task Formalization -- 3.2 Overall Framework --

3.3 Keyword Fronting -- 3.4 Attribution Verification -- 3.5 Iterative Optimization -- 4 Experiments -- 4.1 Experimental Setup -- 4.2 Main Results -- 4.3 Ablation Studies -- 4.4 Impact of Hyperparameters -- 4.5 The Performance of the Iteration -- 5 Conclusion -- References -- Reconstructing the Unseen: GRIOT for Attributed Graph Imputation with Optimal Transport -- 1 Introduction -- 2 Related Works -- 3 Multi-view Optimal Transport Loss for Attribute Imputation -- 3.1 Notations -- 3.2 Optimal Transport and Wasserstein Distance -- 3.3 Definition of the $\text{MultiW Loss Function}$ -- 3.4 Instantiation of MultiW Loss with Attributes and Structure -- 4 Imputing Missing Attributes with MultiW Loss -- 4.1 Architecture of GRIOT -- 4.2 Accelerating the Imputation -- 5 Experimental Analysis -- 5.1 Experimental Protocol -- 5.2 Imputation Quality v.s. Node Classification Accuracy -- 5.3 Imputing Missing Values for Unseen Nodes -- 5.4 Time Complexity -- 6 Conclusion and Perspectives -- References -- Introducing Total Harmonic Resistance for Graph Robustness Under Edge Deletions -- 1 Introduction -- 2 Problem Statement and a New Robustness Measure -- 2.1 Problem Statement and Notation -- 2.2 Robustness Measures -- 3 Related Work -- 4 Comparison of Exact Solutions -- 5 Greedy Heuristic for k-GRoDel -- 5.1 Total Harmonic Resistance Loss After Deleting an Edge. 5.2 Forest Index Loss After Deleting an Edge.

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.
