

40th AAAI Conference on Artificial Intelligence (AAAI-26), 38th Conference on Innovative Applications of Artificial Intelligence (IAAI-26), and 16th Symposium on Educational Advances in Artificial Intelligence (EAAI-26)

Volume 29: AAAI Technical Tracks

- Machine Learning VI

Singapore
20-27 January 2026

ISBN: 979-8-3313-3403-1

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2026) by Association for the Advancement of Artificial Intelligence
All rights reserved.

Printed with permission by Curran Associates, Inc. (2026)

For permission requests, please contact Association for the Advancement of Artificial Intelligence
at the address below.

Association for the Advancement of Artificial Intelligence
2275 East Bayshore Road
Suite 160
Palo Alto, California 94303
USA

Phone: 1-650-328-3123
Fax: 1-650-321-4457

<https://aaai.org/Press/press.php>

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

TABLE OF CONTENTS

AAAI TECHNICAL TRACK ON MACHINE LEARNING VI

DFDT: Dynamic Fast Decision Tree for IoT Data Stream Mining on Edge Devices.....	24053
<i>Afonso Lourenço, João Rodrigo, João Gama, Goreti Marreiros</i>	
Adapt Before Continual Learning	24061
<i>Aojun Lu, Tao Feng, Hangjie Yuan, Chunhui Ding, Yanan Sun</i>	
Towards Non-Stationary Time Series Forecasting with Temporal Stabilization and Frequency Differencing.....	24070
<i>Junkai Lu, Peng Chen, Chenjuan Guo, Yang Shu, Meng Wang, Bin Yang</i>	
Autonomous Concept Drift Threshold Determination	24079
<i>Pengqian Lu, Jie Lu, Anjin Liu, En Yu, Guangquan Zhang</i>	
ProGMLP: A Progressive Framework for GNN-To-MLP Knowledge Distillation with Efficient Trade-Offs	24088
<i>Weigang Lu, Ziyu Guan, Wei Zhao, Yaming Yang, Yujie Sun, Zheng Liang, Yibing Zhan, Dapeng Tao</i>	
STEP-Nav: Spatial-Temporal Efficient Visual Token Pruning for Vision-And-Language Navigation with Large Language Models.....	24097
<i>Yantao Lu, Shiqi Sun, Ning Liu, Bo Jiang, Ying Zhang, Jinchao Chen, Chenglie Du</i>	
Spiking-Aided Neural Architecture for Efficient and Robust WiFi Sensing	24106
<i>Yisha Lu, Liwen Jing, Jiangmao Zheng, Bowen Zhang</i>	
Beyond the Mean: Fisher-Orthogonal Projection for Natural Gradient Descent in Large Batch Training	24115
<i>Yishun Lu, Wesley Armour</i>	
SkipCat: Rank-Maximized Low-Rank Compression of Large Language Models Via Shared Projection and Block Skipping.....	24124
<i>Yu-Chen Lu, Sheng-Feng Yu, Hui-Hsien Weng, Pei-Shuo Wang, Yu-Fang Hu, Liang Hung- Chun, Hung-Yueh Chiang, Kai-Chiang Wu</i>	
Attention Retention for Continual Learning with Vision Transformers	24133
<i>Yue Lu, Xiangyu Zhou, Shizhou Zhang, Yinghui Xing, Guoqiang Liang, Wencong Zhang</i>	
Towards Effective, Stealthy, and Persistent Backdoor Attacks Targeting Graph Foundation Models.....	24142
<i>Jiayi Luo, Qingyun Sun, Lingjuan Lyu, Ziwei Zhang, Haonan Yuan, Xingcheng Fu, Jianxin Li</i>	
ViTCoP: Accelerating Large Vision-Language Models Via Visual and Textual Semantic Collaborative Pruning.....	24151
<i>Wen Luo, Peng Chen, Xiaotao Huang, Liqun Huang</i>	
QuoTA: Query-Oriented Token Assignment Via CoT Query Decouple for Long Video Comprehension.....	24160
<i>Yongdong Luo, Wang Chen, Weizhong Huang, Shukang Yin, Haojia Lin, Jinfa Huang, Chaoyou Fu, Jiayi Ji, Xiawu Zheng, Jiebo Luo</i>	

Supervised Dynamic Dimension Reduction with Deep Neural Network.....	24169
<i>Zhanye Luo, Yuefeng Han, Xiufan Yu</i>	
Quantum Algorithms for Spectral Sums.....	24178
<i>Alessandro Luongo, Changpeng Shao</i>	
LLMC+: Benchmarking Vision-Language Model Compression with a Plug-And-Play Toolkit.....	24189
<i>Chengtao Lv, Bilang Zhang, Yang Yong, Ruihao Gong, Yushi Huang, Shiqiao Gu, Jiajun Wu, Yumeng Shi, Jinyang Guo, Wenya Wang</i>	
Uncertainty-Guided View-Strength-Aware Feature Utilization for Multi-View Classification.....	24198
<i>Li Lv, Qian Guo, Li Zhang, Liang Du, Bingbing Jiang, Lu Chen, Xinyan Liang</i>	
Partial Fairness Awareness: Belief-Guided Strategic Mechanism for Strategic Agents	24207
<i>Xinpeng Lv, Chunyuan Zheng, Yunxin Mao, Renzhe Xu, Hao Zou, Shanzhi Gu, Liyang Xu, Huan Chen, Yuanlong Chen, Wenjing Yang, Haotian Wang</i>	
OccamVTS: Distilling Vision Models to 1% Parameters for Time Series Forecasting.....	24216
<i>Sisuo Lyu, Siru Zhong, Weilin Ruan, Qingxiang Liu, Qingsong Wen, Hui Xiong, Yuxuan Liang</i>	
Cross-Space Synergy: A Unified Framework for Multimodal Emotion Recognition in Conversation	24226
<i>Xiaosen Lyu, Jiayu Xiong, Yuren Chen, Wanlong Wang, Xiaoqing Dai, Jing Wang</i>	
TIMA: Text-Image Mutual Awareness for Balancing Zero-Shot Adversarial Robustness and Generalization Ability	24235
<i>Fengji Ma, Hei Victor Cheng, Chenxing Li, Li Liu</i>	
GraphTextack: A Realistic Black-Box Node Injection Attack on LLM-Enhanced GNNs	24244
<i>Jiaji Ma, Pujia Trivedi, Danai Koutra</i>	
MMG-Vid: Maximizing Marginal Gains at Segment-Level and Token-Level for Efficient Video LLMs.....	24253
<i>Junpeng Ma, Qizhe Zhang, Ming Lu, Zhibin Wang, Qiang Zhou, Jun Song, Shanghang Zhang</i>	
Discovering Mixture Skills for Unsupervised Reinforcement Learning	24262
<i>Nelson Ma, Junyu Xuan, Guangquan Zhang, Jie Lu</i>	
OmniScale: Scaling Any Modality Model Training with Model-Centric Distributed Recipe Zoo.....	24271
<i>Qianli Ma, Yaowei Zheng, Zhelun Shi, Zhongkai Zhao, Bin Jia, Ziyue Huang, Zhiqi Lin, Youjie Li, Jiacheng Yang, Yanghua Peng, Zhi Zhang, Xin Liu</i>	
C3RL: Rethinking the Combination of Channel-Independence and Channel-Mixing from Representation Learning.....	24281
<i>Shusen Ma, Yunbo Zhao, Yu Kang</i>	
Aligning the True Semantics: Constrained Decoupling and Distribution Sampling for Cross-Modal Alignment.....	24290
<i>Xiang Ma, Lexin Fang, Litian Xu, Caiming Zhang</i>	
ReCast: Reliability-Aware Codebook-Assisted Lightweight Time Series Forecasting	24299
<i>Xiang Ma, Taihua Chen, Pengcheng Wang, Xuemei Li, Caiming Zhang</i>	
LLM-Enhanced Energy Contrastive Learning for Out-Of-Distribution Detection in Text-Attributed Graphs	24308
<i>Xiaoxu Ma, Dong Li, Minglai Shao, Xintao Wu, Chen Zhao</i>	

Contribution-Aware Token Compression for Efficient Video Understanding Via Reinforcement Learning	24317
<i>Yinchao Ma, Qiang Zhou, Zhibin Wang, Xianing Chen, Hanqing Yang, Jun Song, Bo Zheng</i>	
Loss-Guided Auxiliary Agents for Overcoming Mode Collapse in GFlowNets	24326
<i>Idriss Malek, Aya Laajil, Abhijith Sharma, Eric Moulines, Salem Lahlou</i>	
Breaking the Dyadic Barrier: Rethinking Fairness in Link Prediction Beyond Demographic Parity.....	24335
<i>Joao Mattos, Debolina Halder Lina, Arlei Silva</i>	
Self-Adaptive Graph Mixture of Models.....	24344
<i>Mohit Meena, Yash Punjabi, Abhishek A, Vishal Sharma, Mahesh Chandran</i>	
Semantic-Aware Feature Enhancement for Partial Label Learning.....	24353
<i>Haowei Mei, Chao Zhang, Wentao Fan, Xiuyi Jia, Chunlin Chen, Huaxiong Li</i>	
Make Model Transparent: Brain Network Analysis Via Causal and Knowledge Graph Learning.....	24362
<i>Lingyuan Meng, Ke Liang, Hao Yu, Haotian Wang, Miaomiao Li, Xinwang Liu</i>	
Online Linear Regression with Paid Stochastic Features	24370
<i>Nadav Merlis, Kyoungseok Jang, Nicolò Cesa-Bianchi</i>	
Learning Branching Policies for MILPs with Proximal Policy Optimization	24379
<i>Abdelouahed Ben Mhamed, Assia Kamal Idrissi, Amal Seghrouchni</i>	
T4NMTD: Transition-Centric Reinforcement Learning for Non-Markovian Task Decomposition.....	24388
<i>Ruixuan Miao, Xu Lu, Cong Tian, Bin Yu, Zhenhua Duan</i>	
Explaining Temporal Graph Neural Network Via Quantum-Inspired Evolutionary Algorithm	24396
<i>Masahiro Mitani, Yuya Sasaki</i>	
Neural Architecture and Hyperparameter Selection Through Meta-Learning on Time Series.....	24405
<i>Erfan Moeini, Christopher Vox, Marie Anastacio, Wadie Skaf, Mitra Baratchi, Holger H. Hoos</i>	
Exact Shapley Attributions in Quadratic-Time for FANOVA Gaussian Processes.....	24414
<i>Majid Mohammadi, Krikamol Muandet, Ilaria Tiddi, Annette Ten Teije, Siu Lun Chau</i>	
Neuro-Symbolic Federated Learning Over Heterogeneous Data-Views: A Structured Approach to Distributive EHR Modelling	24422
<i>Soheila Molaei, Bahareh Fatemi, Anshul Thakur, Andrew Soltan, Fazle Rabbi, Andreas L. Opdahl, Kim Branson, Patrick Schwab, Danielle Belgrave, David A. Clifton</i>	
Bant: Byzantine Antidote Via Trial Function and Trust Scores	24431
<i>Gleb Molodtsov, Daniil Medyakov, Sergey Skorik, Nikolas Khachaturov, Shahane Tigranyan, Vladimir Aletov, Aram Avetisyan, Martin Takác, Aleksandr Beznosikov</i>	
Expandable and Differentiable Dual Memories with Orthogonal Regularization for Exemplar-Free Continual Learning.....	24441
<i>Hyung-Jun Moon, Sung-Bae Cho</i>	
Inference Offloading for Cost-Sensitive Binary Classification at the Edge	24449
<i>Vishnu Narayanan Moothedath, Umang Agarwal, Umeshraja N, James Richard Gross, Jaya Prakash Champati, Sharayu Moharir</i>	

CoRe-Fed: Bridging Collaborative and Representation Fairness Via Federated Embedding Distillation	24458
<i>Noorain Mukhtiar, Adnan Mahmood, Quan Z. Sheng</i>	
Unlocking Efficient Vehicle Dynamics Modeling Via Analytic World Models	24467
<i>Asen Nachkov, Danda Pani Paudel, Jan-Nico Zaech, Davide Scaramuzza, Luc Van Gool</i>	
From Decision Trees to Boolean Logic: A Fast and Unified SHAP Algorithm.....	24476
<i>Alexander Nadel, Ron Wettenstein</i>	
RankList – a Listwise Preference Learning Framework for Predicting Subjective Preferences	24486
<i>Abinay Reddy Naini, Fernando Diaz, Carlos Busso</i>	
Parameter-Free Optimal Rates for Nonlinear Semi-Norm Contractions with Applications to Q- Learning	24495
<i>Ankur Naskar, Gagan Thoppe, Vijay Gupta</i>	
Policy Search, Retrieval, and Composition Via Task Similarity in Collaborative Agentic Systems	24504
<i>Saptarshi Nath, Christos Peridis, Eseoghene Benjamin, Xinran Liu, Soheil Kolouri, Peter Kinnell, Zexin Li, Cong Liu, Shirin Dora, Andrea Soltoggio</i>	
GP-MoLFormer-Sim: Test Time Molecular Optimization Through Contextual Similarity Guidance	24513
<i>Jirí Navrátil, Jarret Ross, Payel Das, Youssef Mroueh, Samuel C Hoffman, Vijil Chenthamarakshan, Brian Belgodere</i>	
Context-Sensitive Abstractions for Reinforcement Learning with Parameterized Actions	24522
<i>Rashmeet Kaur Nayyar, Naman Shah, Siddharth Srivastava</i>	
Hexaïssa: Standing on Giants’ Shoulders—Routing the Best Chess Engines with Mixture-Of- Experts and Latent Reward Learning	24532
<i>Bach Ngo, Nguyen Hoang Khoi Do</i>	
Provably Data-Driven Projection Method for Quadratic Programming.....	24541
<i>Anh Tuan Nguyen, Viet Anh Nguyen</i>	
PKR-QA: A Benchmark for Procedural Knowledge Reasoning with Knowledge Module Learning	24549
<i>Thanh-Son Nguyen, Hong Yang, Tzeh Yuan Neoh, Hao Zhang, Ee Yeo Keat, Basura Fernando</i>	
Disentangled Generation-Based Prototypical Alignment for Few-Shot Unsupervised Domain Adaptation in Graph-Level Anomaly Detection.....	24558
<i>Zhibin Ni, Chenghao Zhang, Hai Wan, Xibin Zhao</i>	
AutoSchA: Automatic Hierarchical Music Representations Via Multi-Relational Node Isolation	24567
<i>Stephen Ni-Hahn, Rico Zhu, Jerry Yin, Yue Jiang, Cynthia Rudin, Simon Mak</i>	
Mitigating Endogenous Confirmation Bias in Noisy Label Learning for Vision-Language Models	24576
<i>Feiyang Ning, Xinyang Chen</i>	
Object-Centric World Models for Causality-Aware Reinforcement Learning.....	24585
<i>Yosuke Nishimoto, Takashi Matsubara</i>	
FlowPath: Learning Data-Driven Manifolds with Invertible Flows for Robust Irregularly-Sampled Time Series Classification	24594
<i>Yongkyung Oh, Dong-Young Lim, Sungil Kim</i>	
Task Prototype-Based Knowledge Retrieval for Multi-Task Learning from Partially Annotated Data.....	24604
<i>Youngmin Oh, Hyung-Il Kim, Jung Uk Kim</i>	

Do it for HER: First-Order Temporal Logic Reward Specification in Reinforcement Learning.....	24613
<i>Pierriccardo Olivieri, Fausto Lasca, Alessandro Gianola, Matteo Papini</i>	
Computing Probabilistic Explanations for ML Models: Fixed-Parameter Algorithms	24622
<i>Sebastian Ordyniak, Mateusz Rychlicki, Stefan Szeider</i>	
Large Language Models Meet Extreme Multi-Label Classification: Scaling and Multi-Modal Framework.....	24630
<i>Diego Ortego, Marlon Rodríguez, Mario Almagro, Kunal Dahiya, David Jimenez, Juan C. Sanmiguel</i>	
The Strong Lottery Ticket Hypothesis for Multi-Head Attention Mechanisms.....	24639
<i>Hikari Otsuka, Daiki Chijiwa, Yasuyuki Okoshi, Daichi Fujiki, Susumu Takeuchi, Masato Motomura</i>	
ParaDySe: A Parallel Strategy Switching Framework for Dynamic Sequences in Transformer-Based Large Language Models	24648
<i>Zhixin Ou, Peng Liang, Linbo Qiao, Jianchen Han, Baihui Liu</i>	
DisCo DETR: Distance-Aware Multi-View Contrastive Learning for DETR Pre-Training.....	24656
<i>Chao Ouyang, Yuyang Bai, Jun Zhang, Tianlu Gao, Jun Hao, Lijun Kong, David Wenzhong Gao</i>	
Learn from Global Correlations: Enhancing Evolutionary Algorithm Via Spectral GNN	24665
<i>Kaichen Ouyang, Zong Ke, Shengwei Fu, Lingjie Liu, Puning Zhao, Dayu Hu</i>	
Policy Newton Methods for Distortion Riskmetrics.....	24674
<i>Soumen Pachal, Mizhaan Prajit Maniyar, Prashanth L. A.</i>	
OPTION: An Online Pricing Strategy for Asynchronous Federated Learning Against Free-Riding Attacks.....	24682
<i>Bangqi Pan, Jianfeng Lu, Shuqin Cao, Xiao Zhang, Gang Li, Guanghui Wen</i>	
Correcting False Alarms from Unseen: Adapting Graph Anomaly Detectors at Test Time.....	24691
<i>Junjun Pan, Yixin Liu, Chuan Zhou, Fei Xiong, Alan Wee-Chung Liew, Shirui Pan</i>	
Re-Architecting Personalized Federated Learning for Demanding Edge Environments	24700
<i>Quyuan Pan, Sheng Sun, Tingting Wi, Zhiyuan Wu, Yuwei Wang, Min Liu, Bo Gao, Jingyuan Wang</i>	
Prototype Entropy Alignment: Reinforcing Structured Uncertainty in LLM Reasoning	24709
<i>Zhengyuan Pan, Yanhao Chen, Zhongquan Jian, Wanru Zhao, Haonan Ma, Meihong Wang, Qingqiang Wu</i>	
On the Misalignment Between Data Learnability and Forgetting in Machine Unlearning	24718
<i>Zijie Pan, Zuobin Ying, Yajie Wang, Wanlei Zhou</i>	
Dynamic Sparsity: Challenging Common Sparsity Assumptions for Learning World Models in Robotic Reinforcement Learning Benchmarks.....	24727
<i>Muthukumar Pandaram, Jakob Hollenstein, David Drexel, Samuele Tosatto, Antonio Rodriguez-Sanchez, Justus Piater</i>	
Forest Vs Tree: The (N, K) Trade-Off in Reproducible ML Evaluation.....	24736
<i>Deepak Pandita, Flip Korn, Chris Welty, Christopher M Homan</i>	

Decoupling Shared and Personalized Knowledge: A Dual-Branch Federated Learning Framework for Multi-Domain with Non-IID Data	24745
<i>Yiran Pang, Zhen Ni, Xiangnan Zhong</i>	
MoCHA: Advanced Vision-Language Reasoning with MoE Connector and Hierarchical Group Attention	24755
<i>Yuqi Pang, Bowen Yang, Yun Cao, Rong Fan, Xiaoyu Li, Chen He</i>	
Transferable Model-Agnostic Vision-Language Model Adaptation for Efficient Weak-To-Strong Generalization	24764
<i>Jihwan Park, Taehoon Song, Sanghyeok Lee, Miso Choi, Hyunwoo J. Kim</i>	
Fair Model-Based Clustering.....	24773
<i>Jinwon Park, Kunwoong Kim, Jihu Lee, Yongdai Kim</i>	
ProLoG: Hybrid Prompt and LoRA Based Adaptation of Vision-Language Models for OOD Generalization	24782
<i>Jungwuk Park, Dong-Jun Han, Jaekyun Moon</i>	
How Many Experts Are Enough? Towards Optimal Semantic Specialization for Mixture-Of-Experts.....	24792
<i>Sumin Park, Noseong Park</i>	
What Makes a Good Generated Image? Investigating Human and Multimodal LLM Image Preference Alignment.....	24801
<i>Rishab Parthasarathy, Jasmine Collins, Cory Stephenson</i>	
Semantic-Consistent Bidirectional Contrastive Hashing for Noisy Multi-Label Cross-Modal Retrieval	24811
<i>Likang Peng, Chao Su, Wenyuan Wu, Yuan Sun, Dezhong Peng, Xi Peng, Xu Wang</i>	
Surrogate as Teacher: Distillation-Guided Graph Poisoning Attack.....	24820
<i>Xingyu Peng, Ke Xu</i>	
A Novel Sliced Fused Gromov-Wasserstein Distance.....	24828
<i>Moritz Piening, Robert Beinert</i>	
ReBoot: Encrypted Training of Deep Neural Networks with CKKS Bootstrapping.....	24837
<i>Alberto Pirillo, Luca Colombo</i>	
On Stealing Graph Neural Network Models	24846
<i>Marcin Podhajski, Jan Dubinski, Franziska Boenisch, Adam Dziedzic, Agnieszka Pregowska, Tomasz Pawel Michalak</i>	
Online Multi-LLM Selection Via Contextual Bandits Under Unstructured Context Evolution.....	24855
<i>Manhin Poon, Xiangxiang Dai, Xutong Liu, Fang Kong, John C. S. Lui, Jinhang Zuo</i>	
Learning Fair Representations with Kolmogorov-Arnold Networks	24864
<i>Amisha Priyadarshini, Sergio Gago-Masague</i>	
Graph Contrastive Learning with Balanced Hard Negatives and Fine-Grained Semantic-Aware Positives	24873
<i>Hongshan Pu, Haoxu Zhang, Ye Liu, Hongmin Cai</i>	

CoCo-MILP: Inter-Variable Contrastive and Intra-Constraint Competitive MILP Solution Prediction	24882
<i>Tianle Pu, Jianing Li, Yingying Gao, Shixuan Liu, Zijie Geng, Haoyang Liu, Chao Chen, Changjun Fan</i>	
Revisiting the Data Sampling in Multimodal Post-Training from a Difficulty-Distinguish View.....	24891
<i>Jianyu Qi, Ding Zou, Wenrui Yan, Rui Ma, Jiaxu Li, Zhijie Zheng, Zhiguo Yang, Rongchang Zhao</i>	
Continuous Vision-Language-Action Co-Learning with Semantic-Physical Alignment for Behavioral Cloning.....	24900
<i>Xiuxiu Qi, Yu Yang, Jiannong Cao, Luyao Bai, Chongshan Fan, Chengtai Cao, Hongpeng Wang</i>	
Next Generation Active Learning: Mixture of LLMs in the Loop	24909
<i>Yuanyuan Qi, Xiaohao Yang, Jueqing Lu, Guoxiang Guo, Joanne Enticott, Gang Liu, Lan Du</i>	
Full-Atom Peptide Design Via Riemannian–Euclidean Bayesian Flow Networks	24918
<i>Hao Qian, Shikui Tu, Lei Xu</i>	
CR ³ : Boosting Compositional Reasoning in MLLMs Through Rule-Based Reinforcement Learning	24927
<i>Shun Qian, Bingquan Liu, Chengjie Sun, Peijin Xie, Zhen Xu, Baoxun Wang</i>	
Beyond Binary Erasure: Soft-Weighted Unlearning for Fairness and Robustness	24936
<i>Xinbao Qiao, Ningning Ding, Yushi Cheng, Meng Zhang</i>	
Beyond Missing Data Imputation: Information-Theoretic Coupling of Missingness and Class Imbalance for Optimal Irregular Time Series Classification	24945
<i>Xin Qin, Mengna Liu, Wenjie Wang, Shuxin Li, Tianjiao Li, Xiufeng Liu, Xu Cheng</i>	
Multi-View Learning Via Trusted Pairwise Entity Energy	24954
<i>Yalan Qin, Guorui Feng, Xinpeng Zhang</i>	

Author Index