

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 24: AAAI Technical Tracks

- Machine Learning I

Singapore
20-27 January 2026

ISBN: 979-8-3313-3398-0

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 I

FedGRPO: Privately Optimizing Foundation Models with Group-Relative Rewards from Domain Clients.....	19487
<i>Gongxi Zhu, Hanlin Gu, Lixin Fan, Qiang Yang, Yuxing Han</i>	
EfficientFlow: Efficient Equivariant Flow Policy Learning for Embodied AI.....	19496
<i>Jianlei Chang, Ruofeng Mei, Wei Ke, Xiangyu Xu</i>	
Automatic Channel Pruning by Searching with Structure Embedding for Hash Network.....	19505
<i>Zifan Liu, Yuan Cao, Yifan Sun, Yanwei Yu, Heng Qi</i>	
DOS: Distilling Observable Softmaps of Zipfian Prototypes for Self-Supervised Point Representation.....	19514
<i>Mohamed Abdelsamad, Michael Ulrich, Bin Yang, Miao Zhang, Yakov Miron, Abhinav Valada</i>	
Constrained Online Convex Optimization with Memory and Predictions.....	19524
<i>Mohammed Abdullah, George Iosifidis, Salah Eddine Elayoubi, Tijani Chahed</i>	
Expressive Temporal Specifications for Reward Monitoring.....	19533
<i>Omar Adalat, Francesco Belardinelli</i>	
ProbLog4Fairness: A Neurosymbolic Approach to Modeling and Mitigating Bias.....	19542
<i>Rik Adriaensen, Lucas Van Praet, Jessa Bekker, Robin Manhaeve, Pieter Delobelle, Maarten Buyt</i>	
RefiDiff: Progressive Refinement Diffusion for Efficient Missing Data Imputation.....	19551
<i>Md Atik Ahamed, Qiang Ye, Qiang Cheng</i>	
Stabilizing Policy Gradient Methods Via Reward Profiling.....	19560
<i>Shihab Ahmed, El Houcine Bergou, Yue Wang, Aritra Dutta</i>	
Expressive Power of Graph Transformers Via Logic.....	19569
<i>Veeti Ahvonen, Maurice Funk, Damian Heiman, Antti Kuusisto, Carsten Lutz</i>	
PharmaQA: Prompt-Based Molecular Representation Learning Via Pharmacophore-Oriented Question Answering.....	19580
<i>Chengwei Ai, Qiaozhen Meng, Mengwei Sun, Ruihan Dong, Hongpeng Yang, Shiqiang Ma, Xiaoyi Liu, Cheng Liang, Fei Guo</i>	
PAGE: A Unified Approach for Federated Graph Unlearning.....	19589
<i>Yuming Ai, Xunkai Li, Jiaqi Chao, Bowen Fan, Zhengyu Wu, Yinlin Zhu, Rong-Hua Li, Guoren Wang</i>	
InfoQ: Mixed-Precision Quantization Via Global Information Flow.....	19598
<i>Mehmet Emre Akbulut, Hazem Hesham Yousef Shalby, Fabrizio Pittorino, Manuel Roveri</i>	
Symmetric Aggregation of Conformity Scores for Efficient Uncertainty Sets.....	19607
<i>Nabil Alami, Jad Zakharia, Souhaib Ben Taieb</i>	

Removing Box-Free Watermarks for Image-To-Image Models Via Query-Based Reverse Engineering	19615
<i>Haonan An, Guang Hua, Hangcheng Cao, Zhengru Fang, Guowen Xu, Susanto Rahardja, Yuguang Fang</i>	
FreDN: Spectral Disentanglement for Time Series Forecasting Via Learnable Frequency Decomposition.....	19623
<i>Zhongde An, Jinhong You, Jiyanglin Li, Yiming Tang, Wen Li, Heming Du, Shouguo Du</i>	
SOSControl: Enhancing Human Motion Generation Through Saliency-Aware Symbolic Orientation and Timing Control	19632
<i>Ho Yin Au, Junkun Jiang, Jie Chen</i>	
Spectral Basis Learning for Expressive Graph Neural Networks in Link Prediction.....	19640
<i>Nilofar Azizi, Nils M. Kriege, Nicholas J. A. Harvey, Horst Bischof</i>	
Convergence of Fast Policy Iteration in Markov Games and Robust MDPs.....	19649
<i>Keith Badger, Jefferson Huang, Marek Petrik</i>	
Mechanistic Dissection of Cross-Attention Subspaces in Text-To-Image Diffusion Models.....	19657
<i>Jun-Hyun Bae, Wonyong Jo, Jaehyup Lee, Heechul Jung</i>	
Medical Vision–Language Pretraining with LLM-Guided Temporal Supervision.....	19666
<i>Liang Bai, Zhi Wang, Huimin Yan, Xian Yang</i>	
Multi-Level Domain Adaptation and Contrastive Domain Isolation with Bilinear Fusion for Patient Drug Response Prediction	19675
<i>Yuting Bai, Hanwen Lv, Wanwan Shi, Zhiyi Zou, Jiawei Luo</i>	
Collaborative Dual Representations for Semi-Supervised Partial Label Learning.....	19684
<i>Wei-Xuan Bao, Yong Rui, Min-Ling Zhang</i>	
Revisiting (Un)Fairness in Recourse by Minimizing Worst-Case Social Burden	19693
<i>Ainhize Barrainkua, Giovanni De Toni, Jose A. Lozano, Novi Quadrianto</i>	
Differentially Private Linear Programming: Reduced Sub-Optimality and Guaranteed Constraint Satisfaction	19702
<i>Alexander Benvenuti, Brendan Bialy, Miriam Dennis, Matthew Hale</i>	
Shaping Without Tearing: Controllable Diffeomorphic Deformations for Topology-Preserving 3D Point Cloud Augmentation	19711
<i>Jian Bi, Qianliang Wu, Jianjun Qian, Lei Luo, Jian Yang</i>	
Deep Clustering Based on Sparse Kolmogorov-Arnold Network and Spectral Constraint.....	19720
<i>Zixuan Bi, Yang Zhao, Ganchao Liu</i>	
FedALT: Federated Fine-Tuning Through Adaptive Local Training with Rest-Of-World LoRA	19728
<i>Jieming Bian, Lei Wang, Letian Zhang, Jie Xu</i>	
DiffOP: Reinforcement Learning of Optimization-Based Control Policies Via Implicit Policy Gradients	19737
<i>Yuexin Bian, Jie Feng, Yuanyuan Shi</i>	
MDBench: Benchmarking Data-Driven Methods for Model Discovery.....	19746
<i>Amirmohammad Ziaei Bideh, Aleksandra Georgievska, Jonathan Gryak</i>	

Condensed Data Expansion Using Model Inversion for Knowledge Distillation	19755
<i>Kuluhan Binici, Shivam Aggarwal, Cihan Acar, Nam Trung Pham, Kariato Leman, Gim Hee Lee, Tulika Mitra</i>	
Asymptotic and Finite Sample Analysis of Nonexpansive Stochastic Approximations with Markovian Noise	19764
<i>Ethan Blaser, Shangtong Zhang</i>	
MechaFormer: Sequence Learning for Kinematic Mechanism Design Automation	19773
<i>Diana Bolanos, Mohammadmehdi Ataei, Pradeep Kumar Jayaraman</i>	
Cancer Survival Prediction by Cyclic Generation and Multi-Grained Alignment.....	19781
<i>Yongqi Bu, Qinggang Niu, Zhen Li, Yanyu Xu, Jun Wang, Guoxian Yu</i>	
SafeMIL: Learning Offline Safe Imitation Policy from Non-Preferred Trajectories.....	19790
<i>Returaj Burnwal, Nirav Pravinbhai Bhatt, Balaraman Ravindran</i>	
Paths Not Taken: Structure-Based Pruning in PSDD Learning and Inference	19799
<i>Cory Butz, Alejandro Santoscoy-Rivero, Camilla Lewis</i>	
Constrained Best Arm Identification with Tests for Feasibility.....	19808
<i>Ting Cai, Kirthevasan Kandasamy</i>	
HyperGOOD: Towards Out-Of-Distribution Detection in Hypergraphs.....	19817
<i>Tingyi Cai, Yunliang Jiang, Ming Li, Changqin Huang, Yujie Fang, Chengling Gao, Zhonglong Zheng</i>	
Rethinking Explanation Evaluation Under the Retraining Scheme.....	19826
<i>Yi Cai, Thibaud Ardoin, Mayank Gulati, Gerhard Wunder</i>	
HeadHunt-VAD: Hunting Robust Anomaly-Sensitive Heads in MLLM for Tuning-Free Video Anomaly Detection.....	19835
<i>Zhaolin Cai, Fan Li, Ziwei Zheng, Haixia Bi, Lijun He</i>	
Stabilizing Self-Consuming Diffusion Models with Latent Space Filtering	19844
<i>Zhongteng Cai, Yaxuan Wang, Yang Liu, Xueru Zhang</i>	
Spiking Heterogeneous Graph Attention Networks.....	19853
<i>Buqing Cao, Qian Peng, Xiang Xie, Liang Chen, Min Shi, Jianxun Liu</i>	
ProCache: Constraint-Aware Feature Caching with Selective Computation for Diffusion Transformer Acceleration	19862
<i>Fanpu Cao, Yaofo Chen, Zeng You, Wei Luo</i>	
PITE: Multi-Prototype Alignment for Individual Treatment Effect Estimation	19871
<i>Fuyuan Cao, Jiaxuan Zhang, Xiaoli Li</i>	
Causality-Aware Efficient Exploration for Cooperative Multi-Agent Reinforcement Learning	19880
<i>Hongye Cao, Tianpei Yang, Fan Feng, Hammadi Rafik Ouariachi, Yali Du, Meng Fang, Jing Huo, Yang Gao</i>	
Provably Efficient Multi-Objective Bandit Algorithms Under Preference-Centric Customization.....	19889
<i>Linfeng Cao, Ming Shi, Ness B. Shroff</i>	
PPFL: A Parameter Behavior-Driven Plug-In Personalization Engine for Federated Learning	19898
<i>Qianyue Cao, Zongwei Zhu, Zirui Lian, Rui Zhang, Boyu Li, Yi Xiong, Xuehai Zhou</i>	

Unpacking the Implicit Norm Dynamics of Sharpness-Aware Minimization in Tensorized Models	19907
<i>Tianxiao Cao, Kyohei Atarashi, Hisashi Kashima</i>	
LUCID: Learning-Enabled Uncertainty-Aware Certification of Stochastic Dynamical Systems	19916
<i>Ernesto Casablanca, Oliver Schön, Paolo Zuliani, Sadegh Soudjani</i>	
Covariance Scattering Transforms.....	19925
<i>Andrea Cavallo, Ayushman Raghuvanshi, Sundeep Prabhakar Chepuri, Elvin Isufi</i>	
Generalizing Fair Clustering to Multiple Groups: Algorithms and Applications	19934
<i>Diptarka Chakraborty, Kushagra Chatterjee, Debarati Das, Tien-Long Nguyen</i>	
PADiff: Predictive and Adaptive Diffusion Policies for Ad Hoc Teamwork	19943
<i>Hohei Chan, Xinzhi Zhang, Antao Xiang, Weinan Zhang, Mengchen Zhao</i>	
Enhancing Medical Large Vision-Language Models Via Alignment Distillation.....	19952
<i>Aofei Chang, Ting Wang, Fenglong Ma</i>	
D3ToM: Decider-Guided Dynamic Token Merging for Accelerating Diffusion MLLMs.....	19961
<i>Shuo Chen, Xiaofeng Zhang, Qingyang Liu, Li Niu</i>	
Benchmarking Reinforcement Learning Algorithms for ICU Ventilator Settings: An Interpretable and Probabilistic Patient Environment for Doctor Agents.....	19970
<i>Ya-Hsi Chang, Po-Chih Kuo</i>	
LoRA in LoRA: Towards Parameter-Efficient Architecture Expansion for Continual Visual Instruction Tuning	19978
<i>Chang Che, Ziqi Wang, Pengwan Yang, Cheems Wang, Hui Ma, Zenglin Shi</i>	
Dropout Prompt Learning: Towards Robust and Adaptive Vision-Language Models.....	19987
<i>Biao Chen, Lin Zuo, Mengmeng Jing, Kunbin He, Yuchen Wang</i>	
Extendable Planning Via Multiscale Diffusion.....	19996
<i>Chang Chen, Hany Hamed, Doojin Baek, Taegu Kang, Samyeul Noh, Yoshua Bengio, Sungjin Ahn</i>	
Edge Self-Adversarial Augmentation Enhances Graph Contrastive Learning Against Neighborhood Inconsistency	20005
<i>Chunchun Chen, Xing Wei, Jiayi Yang, Chenrun Wang, Yiwei Fu, Yuxing Zhang, Xin Sun, Rui Fan, Wei Ye</i>	
Cross-Domain Few-Shot Learning Via Multi-View Collaborative Optimization with Vision- Language Models	20014
<i>Dexia Chen, Wentao Zhang, Qianjie Zhu, Ping Hu, Weibing Li, Tong Zhang, Ruixuan Wang</i>	
OmniSparse: Training-Aware Fine-Grained Sparse Attention for Long-Video MLLMs	20023
<i>Feng Chen, Yefei He, Shaoxuan He, Yuanyu He, Jing Liu, Lequan Lin, Akide Liu, Zhaoyang Li, Jiyuan Zhang, Zhenbang Sun, Bohan Zhuang, Qi Wu</i>	
GINO-Q: Learning an Asymptotically Optimal Index Policy for Restless Multi-Armed Bandits	20032
<i>Gongpu Chen, Soung Chang Liew, Deniz Gündüz</i>	
VCGD: Visual Clue Guided Decoding with Caption Model for Mitigating Hallucination in Multimodal Large Language Models	20041
<i>Guoqing Chen, Fu Zhang, Bingqian Liu, Chenglong Lu, Jingwei Cheng</i>	

Causality-Aligned Semantic Recovery for Incomplete Cross-Modal Retrieval.....	20050
<i>Haipeng Chen, Yu Liu, Xun Yang, Yuheng Liang, Yingda Lyu</i>	
MS-PPO: Mean Standard Deviation Proximal Policy Optimization for Reliable Parking Space Search in Structured Environments.....	20059
<i>Haoming Chen, Hongliang Guo</i>	
Decentralized Non-Convex Stochastic Optimization with Heterogeneous Variance.....	20067
<i>Hongxu Chen, Ke Wei, Luo Luo</i>	
PROMISE: Prompt-Attentive Hierarchical Contrastive Learning for Robust Cross-Modal Representation with Missing Modalities.....	20076
<i>Jiajun Chen, Sai Cheng, Yuan Yutao, Yirui Zhang, Haitao Yuan, Peng Peng, Yi Zhong</i>	
Do We Truly Need So Many Samples? Multi-LLM Repeated Sampling Efficiently Scales Test- Time Compute.....	20083
<i>Jianhao Chen, Zishuo Xun, Bocheng Zhou, Han Qi, Hangfan Zhang, Qiaosheng Zhang, Yang Chen, Wei Hu, Yuzhong Qu, Shuyue Hu</i>	
MARS: A Meta-Adaptive Reinforcement Learning Framework for Risk-Aware Multi-Agent Portfolio Management.....	20092
<i>Jiayi Chen, Jing Li, Guiling Wang</i>	
Towards Multiple Missing Values-Resistant Unsupervised Graph Anomaly Detection.....	20100
<i>Jiazhen Chen, Xiuqin Liang, Sichao Fu, Zheng Ma, Weihua Ou</i>	
Conditional Distribution Learning for Graph Classification.....	20109
<i>Jie Chen, Hua Mao, Chuanbin Liu, Zhu Wang, Xi Peng</i>	
Offline Fictitious Self-Play for Competitive Games.....	20118
<i>Jingxiao Chen, Weiji Xie, Weinan Zhang, Yong Yu, Ying Wen</i>	
DC-SPAN: A Dual Contrastive Attention Network for Multi-View Clustering.....	20127
<i>Jingyi Chen, Zhibin Dong, Tiejun Li, Yibo Han</i>	
MathSE: Improving Multimodal Mathematical Reasoning Via Self-Evolving Iterative Reflection and Reward-Guided Fine-Tuning.....	20136
<i>Jinhao Chen, Zhen Yang, Jianxin Shi, Tianyu Wo, Jie Tang</i>	
Pseudo Multi-View K-Means Clustering.....	20145
<i>Jinqian Chen, Jihua Zhu, Haoyu Tang, Qinghai Zheng</i>	
Sample-Specific Modality Diagnosis and Cross-Modal Enhancement for Incomplete Multimodal Representations.....	20154
<i>Junsong Chen, Jiyuan Liu, Suyuan Liu, Wei Zhang, Ao Li, En Zhu, Xinwang Liu</i>	
Zero-Shot Recommendation: Towards Class Semantic Relation Learning for Inferring Labels of Unseen Micro-Videos.....	20163
<i>Junyang Chen, Huan Wang, Yirui Wu, Qiuzhen Lin, Yunfeng Diao, Junkai Ji</i>	
OmniDPO: A Preference Optimization Framework to Address Omni-Modal Hallucination.....	20172
<i>Junzhe Chen, Tianshu Zhang, Shiyu Huang, Yuwei Niu, Chao Sun, Rongzhou Zhang, Guanyu Zhou, Lijie Wen</i>	
SAPo: Self-Adaptive Process Optimization Makes Small Reasoners Stronger.....	20181
<i>Kaiyuan Chen, Guangmin Zheng, Jin Wang, Xiaobing Zhou, Xuejie Zhang</i>	

FIRM-MoE: Fine-Grained Expert Decomposition for Resource-Adaptive MoE Inference.....	20190
<i>Keyu Chen, Qihang Zhou, Bin Qian, Zhenyu Wen, Wenchao Meng, Shibo He</i>	
ChartEditor: A Reinforcement Learning Framework for Robust Chart Editing	20199
<i>Liangyu Chen, Yichen Xu, Jianzhe Ma, Yuqi Liu, Donglu Yang, Liang Zhang, Zihao Yue, Wenxuan Wang, Qin Jin</i>	
Generating In-Distribution Counterfactual Explanation for Graph Neural Networks	20208
<i>Linmao Chen, Chaobo He, Junwei Cheng, Chunying Li, Quanlong Guan</i>	
Spatial-Frequency Spiking Neural Network for Underwater Object Detection	20217
<i>Long Chen, Wei Miao, Xin Gao, Yunzhi Zhuge, Hongming Xu, Yaxin Li, Qi Xu</i>	
Connectivity-Guided Sparsification of 2-FWL GNNs: Preserving Full Expressivity with Improved Efficiency	20226
<i>Rongqin Chen, Fan Mo, Pak Lon Ip, Shenghui Zhang, Dan Wu, Ye Li, Leong Hou U</i>	
Extracting Multimodal Learnings in CLIP: Unveiling the Multimodal Generalizable Knowledge	20235
<i>Ruiming Chen, Junming Yang, Shiyu Xia, Xu Yang, Xin Geng</i>	
OTARo: Once Tuning for All Precisions Toward Robust On-Device LLMs	20244
<i>Shaoyuan Chen, Zhixuan Chen, Dawei Yang, Zhihang Yuan, Qiang Wu</i>	
FedMerge: Federated Model Merging for Personalization.....	20253
<i>Shutong Chen, Tianyi Zhou, Guodong Long, Jing Jiang, Chengqi Zhang</i>	
When Top-Ranked Recommendations Fail: Modeling Multi-Granular Negative Feedback for Explainable and Robust Video Recommendation.....	20262
<i>Siran Chen, Boyu Chen, Chenyun Yu, Yi Ouyang, Lei Cheng, Chengxiang Zhuo, Zang Li, Yali Wang</i>	
Piercing the Fog: Disentangling Key Features for Vision Models in Multi-Degradation Scenarios	20271
<i>Siyu Chen, Shiqiang Ma, Fei Guo</i>	
Horizontal and Vertical Federated Causal Structure Learning Via Higher-Order Cumulants	20280
<i>Wei Chen, Wanyang Gu, Linjun Peng, Ting Yan, Ruichu Cai, Zhifeng Hao, Kun Zhang</i>	
Geometry-Aware Variational Information Maximization for Deep Incomplete Multi-View Clustering	20289
<i>Wenlan Chen, Lu Gao, Daoyuan Wang, Fei Guo, Cheng Liang</i>	
GCA: Geometry-Aware Conditional Alignment for Partial Domain Adaptation with Coding Rate Reduction	20298
<i>Xiaohui Chen, Chuan-Xian Ren</i>	
G-IR: Geometric Image Representation for Learning	20307
<i>Xin Chen, Qi Zhao, Wei Zeng, Zongben Xu</i>	
Prune&Comp: Free Lunch for Layer-Pruned LLMs Via Iterative Pruning with Magnitude Compensation.....	20316
<i>Xinrui Chen, Hongxing Zhang, Fanyi Zeng, Yongxian Wei, Yizhi Wang, Xitong Ling, Guanghao Li, Chun Yuan</i>	
XLlinear: A Lightweight and Accurate MLP-Based Model for Long-Term Time Series Forecasting with Exogenous Inputs	20325
<i>Xinyang Chen, Huidong Jin, Yu Huang, Zaiwen Feng</i>	

Scaling Law Analysis in Federated Learning: How to Select the Optimal Model Size?.....	20336
<i>Xuanyu Chen, Nan Yang, Shuai Wang, Dong Yuan</i>	
F2SST: Frequency-To-Spatial Semantic Transfer for Few-Shot Image Classification.....	20345
<i>Xueyi Chen, Bangjun Wang, Jiaqing Fan, Li Zhang, Fanzhang Li</i>	
Beyond Sharpness: A Flatness Decomposition Framework for Efficient Continual Learning.....	20354
<i>Yanan Chen, Tieliang Gong, Yunjiao Zhang, Wen Wen</i>	
Active Multi-Source Domain Adaptation for Multimodal Fake News Detection	20363
<i>Yanping Chen, Weijie Shi, Mengze Li, Yue Cui, Jiaming Li, Ruiyuan Zhang, Hao Chen, Hanghui Guo, Shimin Di, Ziyi Liu, Jia Zhu, Jiajie Xu</i>	
SchellingFormer: Laplacian Matrix-Guided Geometric Transformer for Robust Schelling Point Detection	20372
<i>Yihao Chen, Haobo Jiang, Liang Yu, Jianmin Zheng</i>	
Detecting Unobserved Confounders: A Kernelized Regression Approach.....	20381
<i>Yikai Chen, Yunxin Mao, Chunyuan Zheng, Hao Zou, Shanzhi Gu, Shixuan Liu, Yang Shi, Wenjing Yang, Kun Kuang, Haotian Wang</i>	
Enhancing Kernel Power \mathcal{K} -Means: Scalable and Robust Clustering with Random Fourier Features and Possibilistic Method	20390
<i>Yixi Chen, Weixuan Liang, Tianrui Liu, Jun-Jie Huang, Ao Li, Xueling Zhu, Xinwang Liu</i>	

Author Index