

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 25: AAAI Technical Tracks

- Machine Learning II

Singapore
20-27 January 2026

ISBN: 979-8-3313-3399-7

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 II

Diffusion Model Based Signal Recovery Under 1-Bit Quantization.....	20400
<i>Youming Chen, Zhaoqiang Liu</i>	
AV-SSAN: Audio-Visual Selective DOA Estimation Through Explicit Multi-Band Semantic-Spatial Alignment.....	20409
<i>Yu Chen, Hongxu Zhu, Jiadong Wang, Kainan Chen, Xinyuan Qian</i>	
FedCure: Mitigating Participation Bias in Semi-Asynchronous Federated Learning with Non-IID Data	20418
<i>Yue Chen, Jianfeng Lu, Shuqin Cao, Wei Wang, Gang Li, Guanghui Wen</i>	
Multi-Modal Style Transfer-Based Prompt Tuning for Efficient Federated Domain Generalization.....	20427
<i>Yuliang Chen, Xi Lin, Jun Wu, Xiangrui Cai, Qiaolun Zhang, Xichun Fan, Jiapeng Xu, Xiu Su</i>	
SIFThinker: Spatially-Aware Image Focus for Visual Reasoning	20436
<i>Zhangquan Chen, Ruihui Zhao, Chuwei Luo, Mingze Sun, Xinlei Yu, Yangyang Kang, Ruqi Huang</i>	
Gated Variational Graph Autoencoders as Experts with Competition and Consensus for Multi-View Clustering	20445
<i>Zhaoliang Chen, William K. Cheung, Hong-Ning Dai, Byron Choi, Jiming Liu</i>	
Combining LLM Semantic Reasoning with GNN Structural Modeling for Multi-View Multi-Label Feature Selection.....	20454
<i>Zhiqi Chen, Yuzhou Liu, Jiarui Liu, Wanfu Gao</i>	
INTENT: Invariance and Discrimination-Aware Noise Mitigation for Robust Composed Image Retrieval	20463
<i>Zhiwei Chen, Yupeng Hu, Zhiheng Fu, Zixu Li, Jiale Huang, Qinlei Huang, Yinwei Wei</i>	
CoGenSAM: Codebook-Interactive Generative Labeling for Adapting SAM to Crack Segmentation.....	20472
<i>Zhuangzhuang Chen, Nuo Chen, Dachong Li, Zhiliang Lin, Xingyu Feng, Yifan Zhang, Jie Chen, Jianqiang Li</i>	
Explanation-Preserving Augmentation for Semi-Supervised Graph Representation Learning	20481
<i>Zhuomin Chen, Jingchao Ni, Hojat Allah Salehi, Xu Zheng, Esteban Schafir, Farhad Shirani, Dongsheng Luo</i>	
LSHFed: Robust and Communication-Efficient Federated Learning with Locally-Sensitive Hashing Gradient Mapping.....	20490
<i>Guanjie Cheng, Mengzhen Yang, Xinkui Zhao, Shuyi Yu, Tianyu Du, Yangyang Wu, Mengying Zhu, Shuiguang Deng</i>	
Personalized Federated Learning with Bidirectional Communication Compression Via One-Bit Random Sketching	20499
<i>Jiacheng Cheng, Xu Zhang, Guanghui Qiu, Yifang Zhang, Yinchuan Li, Kaiyuan Feng</i>	
Sample-And-Search: An Effective Algorithm for Learning-Augmented k-Median Clustering in High Dimensions.....	20509
<i>Kangke Cheng, Shihong Song, Guanlin Mo, Hu Ding</i>	

DesireKV: Decoupling Sensitivity and Importance for Reasoning-Aware KV Cache Compression	20518
<i>Pengyu Cheng, Jiacheng Wang, Tianle Chen, Bei Liu, Xiaofeng Hou, Jiacheng Liu</i>	
Score-Based Model for Low-Rank Tensor Recovery	20527
<i>Zhengyun Cheng, Changhao Wang, Guanwen Zhang, Yi Xu, Wei Zhou, Xiangyang Ji</i>	
Best Arm Identification with Biased Contexts	20536
<i>James Cheshire, Stephan Cl��men��on</i>	
Departures: Distributional Transport for Single-Cell Perturbation Prediction with Neural Schr��dinger Bridges	20544
<i>Changxi Chi, Yufei Huang, Jun Xia, Jiangbin Zheng, Yunfan Liu, Zelin Zang, Stan Z. Li</i>	
Are Graph Transformers Necessary? Efficient Long-Range Message Passing with Fractal Nodes in MPNNs.....	20553
<i>Jeongwhan Choi, Seungjun Park, Sumin Park, Sung-Bae Cho, Noseong Park</i>	
TabGeoFlow: A Geometric Flow Matching Model for Tabular Data Synthesis.....	20562
<i>Jong In Choi</i>	
Sheaf Graph Neural Networks Via PAC-Bayes Spectral Optimization.....	20570
<i>Yoonhyuk Choi, Jiho Choi, Taewook Ko, Jongwook Kim, Chong-Kwon Kim</i>	
Let the Void Be Void: Robust Open-Set Semi-Supervised Learning Via Selective Non-Alignment.....	20579
<i>You Rim Choi, Subeom Park, Seojun Heo, Eunchung Noh, Hyung-Sin Kim</i>	
LoGIC: Multi-LoRA Guided Importance Consensus for Multi-Task Pruning in Vision Transformers.....	20588
<i>Yu-Hong Chou, Rui Fang, Hsi-Wen Chen, Ming-Syan Chen</i>	
T3Time: Tri-Modal Time Series Forecasting Via Adaptive Multi-Head Alignment and Residual Fusion.....	20597
<i>Abdul Monaf Chowdhury, Rabeya Akter, Safaeid Hossain Arib</i>	
Convex Clustering Redefined: Robust Learning with the Median of Means Estimator.....	20606
<i>Koustav Chowdhury, Bibhabasu Mandal, Sourav De, Sagar Ghosh, Swagatam Das, Debolina Paul, Saptarshi Chakraborty</i>	
Towards Robust Edge Model Adaptation Via Elastic Architecture Search	20615
<i>Xianhang Chu, Xu Yang, Kun Wei, Xi Wang</i>	
Robust Watermarking on Gradient Boosting Decision Trees	20624
<i>Jun Woo Chung, Yingjie Lao, Weijie Zhao</i>	
CATAL: Causally Disentangled Task Representation Learning for Offline Meta-Reinforcement Learning	20634
<i>Shan Cong, Chao Yu, Xiangyuan Lan</i>	
Causal Discovery from Interval-Based Event Sequences.....	20642
<i>L��na��g Cornanguer, Joscha C��ppers, Jilles Vreeken</i>	
Linear Time Algorithms for Individually Fair K-Means Via Multi-Swap Local Search	20650
<i>Beirong Cui, Qilong Feng, Junyu Huang</i>	
ShaLa: Multimodal Shared Latent Generative Modelling.....	20658
<i>Jiali Cui, Yan-Ying Chen, Yanxia Zhang, Matthew Klenk</i>	

Learning Systems Expansion with Efficient Heterogeneity-Aware Knowledge Transfer	20667
<i>Gaole Dai, Huatao Xu, Yifan Yang, Rui Tan, Mo Li</i>	
MemoryART: Enhancing LLMs Via Multi-Memory Models with Adaptive Resonance Theory for Healthcare Agents.....	20676
<i>Renke Dai, Hebin Hu, Jiahui Zhang, Yilin Kang, Ah-Hwee Tan</i>	
Veli: Unsupervised Method and Unified Benchmark for Low-Cost Air Quality Sensor Correction.....	20684
<i>Yahia Dalbah, Marcel Worring, Yen-Chia Hsu</i>	
Correspondence Coverage Matters for Multi-Modal Dataset Distillation.....	20693
<i>Zhuohang Dang, Minnan Luo, Chengyou Jia, Hangwei Qian, Xinyu Zhang, Xiaojun Chang, Ivor Tsang</i>	
Statistically Robust Sparse High-Order Interaction Model	20702
<i>Diptesh Das, Ichiro Takeuchi, Koji Tsuda</i>	
Intervention-Aware Time Series Modeling: Capturing and Evaluating Feature Dependencies.....	20711
<i>Ibrahim Delibasoglu, Sanjay Chakraborty, Fredrik Heintz, Mattias Tiger</i>	
FedSkeleton: Secure Multi-Party Graph Skeleton Construction for Privacy-Preserving Federated Time-Series Forecasting	20719
<i>Henggang Deng, Yuchao Tang, Wenjie Fu, Huandong Wang, Kun Chen, Tao Jiang</i>	
Efficiently Seeking Flat Minima for Better Generalization in Fine-Tuning Large Language Models and Beyond.....	20728
<i>Jiaxin Deng, Qingcheng Zhu, Junbiao Pang, Linlin Yang, Zhongqian Fu, Baochang Zhang</i>	
InfoDecom: Decomposing Information for Defending Against Privacy Leakage in Split Inference	20737
<i>Ruijun Deng, Zhihui Lu, Qiang Duan</i>	
Random Amalgamation of Adapters for Flatter Loss Landscapes: Towards Class-Incremental Learning with Better Stability	20746
<i>Yao Deng, Xiang Xiang, Jiaqi Gui</i>	
Neural Tangent Kernels Under Stochastic Data Augmentation.....	20755
<i>Joshua Deoliveira, Sajal Chakraborty, Walter Gerych, Elke Rundensteiner</i>	
Demystifying Foreground-Background Memorization in Diffusion Models	20763
<i>Jimmy Z. Di, Yiwei Lu, Yaoliang Yu, Gautam Kamath, Adam Dziedzic, Franziska Boenisch</i>	
CIP-Net: Continual Interpretable Prototype-Based Network	20772
<i>Federico Di Valerio, Michela Proietti, Alessio Ragno, Roberto Capobianco</i>	
SimDiff: Simpler Yet Better Diffusion Model for Time Series Point Forecasting.....	20781
<i>Hang Ding, Xue Wang, Tian Zhou, Tao Yao</i>	
TimeMosaic: Temporal Heterogeneity Guided Time Series Forecasting Via Adaptive Granularity Patch and Segment-Wise Decoding.....	20790
<i>Kuiye Ding, Fanda Fan, Chunyi Hou, Zheya Wang, Lei Wang, Zhengxin Yang, Jianfeng Zhan</i>	
Communication-Efficient Multi-Agent Reinforcement Learning with Spatiotemporal Information Hub.....	20799
<i>Ling Ding, Tianbai Lyu, Zhiliang Bi, Hao Wang, Shanshan Feng, Wei Yu</i>	
PrefixGPT: Prefix Adder Optimization by a Generative Pre-Trained Transformer.....	20808
<i>Ruogu Ding, Xin Ning, Ulf Schlichtmann, Weikang Qian</i>	

Learning Time in Static Classifiers	20816
<i>Xi Ding, Lei Wang, Piotr Koniusz, Yongsheng Gao</i>	
Sliding-Window Merging for Compacting Patch-Redundant Layers in LLMs.....	20826
<i>Xuan Ding, Rui Sun, Yunjian Zhang, Xiu Yan, Yueqi Zhou, Kaihao Huang, Suzhong Fu, Angelica I Aviles-Rivero, Chuanlong Xie, Yao Zhu</i>	
RAC-DMVC: Reliability-Aware Contrastive Deep Multi-View Clustering Under Multi-Source Noise.....	20835
<i>Shihao Dong, Yue Liu, Xiaotong Zhou, Yuhui Zheng, Huiying Xu, Xinzhong Zhu</i>	
Accelerating LLM Inference Throughput Via Asynchronous KV Cache Prefetching.....	20844
<i>Yanhao Dong, Yubo Miao, Weinan Li, Xiao Zheng, Chao Wang, Jiasheng Wu, Feng Lyu</i>	
EnViT: Enhancing the Performance of Early-Exit Vision Transformers Via Exit-Aware Structured Dropout-Enabled Self-Distillation.....	20852
<i>Yonghao Dong, Qiang He, Penghong Rui, Zhenzhe Zheng, Zhao Li, Feifei Chen, Hai Jin, Yun Yang</i>	
Learning Dynamics as Feedback: An Adaptive Entropy Flow Dynamics Framework for Long-Tailed Human Action Recognition	20861
<i>Yuan Dong, Zhe Zhao, Liheng Yu, Di Wu, Pengkun Wang</i>	
Constrained Particle Seeking: Solving Diffusion Inverse Problems with Just Forward Passes	20870
<i>Hongkun Dou, Zike Chen, Zeyu Li, Hongjue Li, Lijun Yang, Yue Deng</i>	
DSCF: Dual-Source Counterfactual Fusion for High-Dimensional Combinatorial Interventions	20879
<i>Jitong Dou, Lingrui Luo, Bing Zhu, Hengliang Luo, Mingjun Zhong, Yurong Cheng</i>	
Closer to Biological Mechanism: Drug-Drug Interaction Prediction from the Perspective of Pharmacophore.....	20887
<i>Mingliang Dou, Linfeng Wen, Jinyang Xie, Jijun Tang, Shiqiang Ma, Fei Guo</i>	
Predicting the Future by Retrieving the Past	20896
<i>Dazhao Du, Tao Han, Song Guo</i>	
GCIB: Causal Intervention Guided Graph Information Bottleneck Framework	20905
<i>Hangyuan Du, Rong Wang, Lixin Cui, Gaoxia Jiang, Liang Bai, Wenjian Wang</i>	
LGAN: An Efficient High-Order Graph Neural Network Via the Line Graph Aggregation	20914
<i>Lin Du, Lu Bai, Jincheng Li, Lixin Cui, Hangyuan Du, Lichi Zhang, Yuting Chen, Zhao Li</i>	
Forget What Has Seen: Selective Concept Unlearning in Segmentation Foundation Models.....	20923
<i>Miaozeng Du, Jiaqi Li, Sirui Pan, Yi Zhan, Guilin Qi, Yuxin Zhang, Rihui Jin, Yinjia Shu, Qianshan Wei</i>	
Flexible Concept Bottleneck Model	20932
<i>Xingbo Du, Qiantong Dou, Lei Fan, Rui Zhang</i>	
Deep Incomplete Multi-View Clustering Via Hierarchical Imputation and Alignment.....	20941
<i>Yiming Du, Ziyu Wang, Jian Li, Rui Ning, Lusi Li</i>	
Learning Intrinsic Hierarchy for Generalized Category Discovery.....	20950
<i>Yu Duan, Junzhi He, Zhanxuan Hu, Mengda Ji, Rong Wang, Quanxue Gao</i>	
K-ProtoDiff: Key Prototypes-Guided Diffusion for Time Series Generation.....	20959
<i>Yuhang Duan, Lin Lin, Xiaoshuai Wu</i>	

FlorE: Integrating Full Lorentz Group and Directional Offsets for Effective Knowledge Graph Embedding.....	20968
<i>Zehua Duo, Jiang Li, Xiangdong Su, Guanglai Gao</i>	
Simulation-Driven Railway Delay Prediction: An Imitation Learning Approach	20977
<i>Clément Elliker, Jesse Read, Sonia Vanier, Albert Bifet</i>	
Decomposing and Composing: Towards Efficient Vision-Language Continual Learning Via Rank-1 Expert Pool in a Single LoRA	20985
<i>Zhan Fa, Yue Duan, Jian Zhang, Lei Qi, Wanqi Yang, Yinghuan Shi</i>	
Efficient Multimodal Large Language Model Via Dynamic KV Cache Quantization	20994
<i>Jiahao Fan, Chien-Ming Chen</i>	
Online Cross-Modal Hashing with Expanding Label Space	21002
<i>Wentao Fan, Chao Zhang, Chunlin Chen, Huaxiong Li</i>	
MoETTA: Test-Time Adaptation Under Mixed Distribution Shifts with MoE-LayerNorm	21011
<i>Xiao Fan, Jingyan Jiang, Zhaoru Chen, Fanding Huang, Xiao Chen, Qinting Jiang, Bowen Zhang, Xing Tang, Zhi Wang</i>	
TFRank: Think-Free Reasoning Enables Practical Pointwise LLM Ranking	21020
<i>Yongqi Fan, Xiaoyang Chen, Dezhi Ye, Jie Liu, Haijin Liang, Jin Ma, Ben He, Yingfei Sun, Tong Ruan</i>	
Hierarchical Structure-Property Alignment for Data-Efficient Molecular Generation and Editing	21029
<i>Ziyu Fan, Zhijian Huang, Yahan Li, Xiaowen Hu, Siyuan Shen, Yunliang Wang, Zeyu Zhong, Shuhong Liu, Shuning Yang, Shangqian Wu, Min Wu, Lei Deng</i>	
Towards Illumination-Aware Restoration of Metalens-Captured Images: A New Dataset and a Strong Baseline.....	21038
<i>Fen Fang, Xinan Liang, Muli Yang, Jinghong Zheng, Tobias Mass, Ying Sun, Xulei Yang, Xuewu Xu, Zhengguo Li</i>	
Graph Domain Adaptation Via Homophily-Agnostic Reconstructing Structure	21047
<i>Ruiyi Fang, Shuo Wang, Ruizhi Pu, Qiuhao Zeng, Hao Zheng, Ziyang Wang, Jiale Cai, Zhimin Mei, Song Tang, Charles Ling, Boyu Wang</i>	
To Align Or Not to Align: Strategic Multimodal Representation Alignment for Optimal Performance.....	21056
<i>Wanlong Fang, Tianle Zhang, Alvin Chan</i>	
Goal-Oriented Time-Series Forecasting: Foundation Framework Design.....	21065
<i>Luca-Andrei Fechete, Mohamed Sana, Fadhel Ayed, Nicola Piovesan, Wenjie Li, Antonio De Domenico, Tareq Si Salem</i>	
WIET: Harmonizing Group-Aware Model Weighting and Worker Allocation for Ensemble Temporal Prediction MaaS	21074
<i>Binbin Feng, Shikun He, Yingxin Wang, Pengwei Wang, Xiang Gao, Zhijun Ding</i>	
Federated Incomplete Multi-View Clustering with Tensorized Low-Rank Constraint	21083
<i>Wei Feng, Danting Liu, Qianqian Wang, Mengping Jiang, Bin Liu</i>	
Behavior Tokens Speak Louder: Disentangled Explainable Recommendation with Behavior Vocabulary.....	21092
<i>Xinshun Feng, Mingzhe Liu, Yi Qiao, Tongyu Zhu, Leilei Sun, Shuai Wang</i>	

Adaptive and Asymptotic Mean-Based Subclass Discriminant Analysis	21101
<i>Yuzhe Feng, Yunlong Gao, Feiping Nie</i>	
EM-KD: Distilling Efficient Multimodal Large Language Model with Unbalanced Vision Tokens	21111
<i>Ze Feng, Sen Yang, Boqiang Duan, Wankou Yang, Jingdong Wang</i>	
Statistical Learning Theory for Distributional Classification	21120
<i>Christian Fiedler</i>	
DeToNATION: Decoupled Torch Network-Aware Training on Interlinked Online Nodes	21128
<i>Mogens Henrik From, Jacob Nielsen, Lukas Galke, Peter Schneider-Kamp</i>	
ENHash: Error Notebook-Guided Fine-Grained Learning for Unsupervised Cross-Modal Hashing	21136
<i>Hao Fu, Zebing Yao, Chuangchuang Tan, Guanghua Gu</i>	
Beyond Fixed Tasks: Unsupervised Environment Design for Task-Level Pairs	21145
<i>Daniel Furelos-Blanco, Charles Pert, Frederik Kelbel, Alex F. Spies, Alessandra Russo, Michael Dennis</i>	
BIQ: Bisection Interval Quantization for Communication-Efficient Federated Learning	21154
<i>Luyang Gai, Shusen Yang, Xuebin Ren, Zihao Zhou</i>	
Fairness-Aware Design for Contextual Experiments: Guaranteeing Reliability and Equity in Heterogeneous Subgroups	21163
<i>Guangyan Gan, Ling Zhang, Yanhua Cheng, Yongxiang Tang, Kaiyuan Li, Xialong Liu, Peng Jiang</i>	
Learning Whom to Align With: Progressive Anomaly Combination Detection for Partially View-Aligned Clustering	21171
<i>Hang Gao, Zuosong Cai, Yuze Li, Cheng Liu, Gaoyang Li, Ying Li, Wei Du, You Zhou</i>	
MARE: Multimodal Analogical Reasoning for Disease Evolution-Aware Radiology Report Generation	21180
<i>Qingqing Gao, Tengfei Liu, Xiaoyan Li, Xiaodan Zhang, Zhongfan Sun, Boyue Wang, Baocai Yin, Zhaohui Liu</i>	
DiAPR: Dimensionally-Allocated Prototype Refinement for Non-Exemplar Class Incremental Learning	21189
<i>Ruixuan Gao, Qijun Zhao, Keren Fu</i>	
CMedBench: A Comprehensive Benchmark for Efficient Medical Large Language Models	21198
<i>Shengbo Gao, Jinyang Guo, Lixian Su, Yifu Ding, Shiqiao Gu, Aishan Liu, Yuqing Ma, Zhiwang Zhang, Xianglong Liu</i>	
The Semantic Architect: How FEAML Bridges Structured Data and LLMs for Multi-Label Tasks	21207
<i>Wanfu Gao, Zebin He, Jun Gao</i>	
Talon: Breaking the Synchronization Barrier in Speculative Decoding with Hybrid Model-Based and Retrieve-Based Drafting	21216
<i>Xiangxiang Gao, Weisheng Xie, Lixin , Xuwei Fang, Chen Hang, Changqun Li, Yuhan Lin, Xiaolong Xu</i>	
Reconcile Gradient Modulation for Harmony Multimodal Learning	21225
<i>Xiyuan Gao, Bing Cao, Baoquan Gong, Pengfei Zhu</i>	
Visual Bridge: Universal Visual Perception Representations Generating	21234
<i>Yilin Gao, Shuguang Dou, Junzhou Li, Zhiheng Yu, Yin Li, Dongsheng Jiang, Shugong Xu</i>	

Adaptive Momentum and EMA-Weighted Modeling for Imbalanced Label Distribution Learning.....	21243
<i>Yongbiao Gao, Xiangcheng Sun, Chao Tan, Chunyu Hu, Guohua Lv</i>	
Unified Structural Factors for Transfer Learning Generalization with PAC-Bayesian Guarantees.....	21252
<i>Ziqi Gao</i>	
TGCD: A Framework for Generalized Category Discovery in Time-Series Data.....	21260
<i>Chandan Gautam, Lew Choon Hean, Ankit Das, Xiaoli Li, Savitha Ramasamy</i>	
TLGC: Taylor Linear Attention-Guided Graph Convolutions for Revealing Spatial Domains in Spatial Multi-Omics Data.....	21269
<i>Aoyun Geng, Chunyan Cui, Yunyun Su, Zhenjie Luo, Feifei Cui, Zilong Zhang</i>	
ORVIT: Near-Optimal Online Distributionally Robust Reinforcement Learning	21278
<i>Debamita Ghosh, George K. Atia, Yue Wang</i>	
SDE-HARL: Scalable Distributed Policy Execution for Heterogeneous-Agent Reinforcement Learning	21287
<i>Toan D. Gian, Mohammad Abdi, Nathaniel D. Bastian, Francesco Restuccia</i>	
HAMLET4Fairness: Enhancing Fairness in AI Pipelines Through Human-Centered AutoML and Argumentation.....	21296
<i>Joseph Giovanelli, Giuseppe Pisano, Roberta Calegari</i>	
Rethinking Membership Inference Attacks for CLIP	21305
<i>Lluis Gomez</i>	

Author Index