

Day 1: Monday, Dec 9, 2024	Sessions/Forums	Room
8:00 - 8:30	Welcome coffee break	
8:30 – 14:00	Causal Representation Learning (CRL)	CS 3
8:30 – 16:00	Sentiment Elicitation from Natural Text for Information Retrieval and Extraction (SENTIRE)	CS 4
8:30 –11:00	Responsible AI to Increase Clinical Decision Trust: Explainability & Reliability of Machine Learning Models (TRUST)	CS 5
8:30 – 16:00	International Workshop on Data Mining for Service (DMS2024)	CS 6
8:30 – 11:00	Workshop on AI for Financial Crime Fight (AI4FCF)	CS 7
8:30 – 11:00	International Workshop on Spatial and Spatio-Temporal Data Mining (SSTDM)	CS 8
8:30 – 11:00	The 2024 Workshop on Optimization Based Techniques for Emerging Data Mining Problems (OEDM)	CS 9
11:00 – 11:30	Coffee Break	
11:30 – 16:00	Incremental Classification and Clustering, Concept Drift, Novelty Detection, Active Learning in Big/Fast data Context (IncrLearn)	CS 5
11:30 – 16:00	International Workshop on Data Mining in Finance (DMF)	CS 7

11:30 – 13:00	Workshop on Information Seeking with Big Models (BigIS)	CS 8
11:30 – 13:00	Deep Learning and Clustering (DLC)	CS 9
13:00 – 14:00	Lunch	
14:00 – 16:00	Workshop on Emerging Trends in Deep Learning for Healthcare (ETDLH)	CS 3
14:00 – 16:00	The 11th ICDM Workshop on High Dimensional Data Mining (HDM)	CS 8
14:00 – 16:00	Data Mining in Biomedical Informatics and Healthcare (DMBIH)	CS 9
16:00 – 16:30	Coffee Break	
16:30 – 19:00	Demo Track	CS 5
16:30 – 19:00	International Workshop on Multimodal Content Analysis for Social Good (MM4SG)	CS 3
16:30 – 19:00	International Workshop on Data-Centric AI (DCAI)	CS 4
16:30 – 19:00	2nd International Workshop on Adaptable, Reliable, and Responsible Learning (ARRL)	CS 6
16:30 – 18:00	Advances in AI-Driven Data Mining for Autonomous Systems (AIDM-AS)	CS 7
16:30 – 18:00	Machine Learning for Cybersecurity (MLC)	CS 8
16:30 – 18:00	International Workshop on AI for Nudging and Personalization (WAIN)	CS 9

18:00 – 19:00	The 2nd International Workshop on User Understanding from Big Data Workshop (DMU2)	CS 7
18:00 – 19:00	Neverending Machine Learning (NML)	CS 8
18:00 – 19:00	Evolutionary Data Mining and Machine Learning Workshop (EDMML)	CS 9
18:00 – 20:00	Steering Committee Meeting with ICDM 2024 & 2025 Main Organizers (By invitation only)	TBA

Conference Agenda

Day 2: Tuesday, Dec 10, 2024		08:00 – 08:45	Welcome coffee break			
		08:45 – 09:00	Opening and Welcome (Eric Xing , Conference Chairs, PC Chairs, Local Chairs)			
		09:00 – 10:00	Keynote 1: Preslav Nakov Towards Safe, Truly Open, and Factual Large Language Models Room CHB			
10:00 – 17:30	Tutorial 1 (5 hours in total) Causality and Large Models Room CHB	10:00 – 10:30	Coffee break			
		10:30 – 12:00	Session A1-1	Session A2-1	Session A3-1	Session A5-1
			Room CS 5	Room CS 7	Room CS 9	Room Hive
		12:00 – 13:30	Lunch			
		13:30 – 15:00	Session A4-1	Session A5-2	Session A6-1	Session A3-2
			Room CS 5	Room CS 7	Room CS 9	Room Hive
		15:00 – 15:30	Coffee break			
		15:30 – 17:00	Session A1-2	Session A2-2	Session A3-3	Session A6-2

			Room CS 5	Room CS 7	Room CS 9	Room Hive
		18:00 – 20:00	Welcome Reception (Location: ADNEC)			

Day 3: Wednesday, Dec 11, 2024		08:00 – 09:00	Welcome coffee break			
		09:00 – 10:00	Session A1-3	Session A3-4	Session A5-3	Session A6-3
			Room CS 5	Room CS 7	Room CS 9	Room Hive
		10:00 – 10:30	Coffee break			
10:30 – 12:00	Women Forum Room Hive	10:30 – 12:00	Session A5-4	Session A6-4	Session A1-4	Session A2-3
			Room CS 5	Room CS 7	Room CS 9	Room CHB
		12:00 – 13:30	Lunch			
		13:30 – 14:30	Keynote 2: Bernhard Schölkopf Towards causal world models and digital twins Room CHB			
		14:30 – 18:00	Organised desert trip			
		18:30 – 20:30	Banquet and Award Ceremony (Location: Desert)			

Day 4: Thursday, Dec 12, 2024		08:00 – 09:00	Welcome coffee break			
		09:00 – 10:00	Keynote 3: Claudia Plant Clustering: Balancing Abstraction and Representation Room CHB			
10:00 – 12:00	Tutorial 2 Hypergraph Neural Networks: An In-Depth and Step-by-Step Guide Room CHB	10:00 – 10:30	Coffee break			
		10:30 – 12:00	Session A3-5	Session A5-5	Session A6-5	Session A1-5
			Room CS 5	Room CS 7	Room CS 9	Room Hive
		12:00 – 13:30	Lunch			
13:30 – 15:30	Tutorial 3 Uncertain Boundaries: Multidisciplinary Approaches to Copyright Issues in Generative AI Room CHB	13:30 – 15:00	Session A1-6	Session A2-5	Session A3-6	Session A2-4
			Room CS 5	Room CS 7	Room CS 9	Room Hive
		15:00 – 15:30	Coffee break			
		15:30 – 17:30	Panel discussion: TBA Room Room CHB			
		17:30	Conference concluding remarks			

Conference paper presentations

Keynote Lecture: 60 minutes (about 45 minutes for talk and 15 minutes for Q and A)

Main conference regular paper (R): 20 minutes (about 15 minutes for talk and 5 minutes for Q and A)

Main conference short paper (S): 15 minutes (about 10 minutes for talk and 5 minutes for Q and A)

Day 2: December 10, 2024

Session A1-1 Foundations, algorithms, models, and theory of data mining

Room CS 5, 10:30-11:40

Session Chair: Xingquan (Hill) Zhu, Florida Atlantic University, xzhu3@fau.edu

10:30	DM306	Efficient Network Embedding by Approximate Equitable Partitions	R
		Giuseppe Squillace, Mirco Tribastone, Max Tschaikowski, and Andrea Vandin	
10:50	DM319	ADOD: Adaptive Density Outlier Detection	R
		Li Qian, Jing Qian, Xin Sun, Wengang Guo, and Christian Böhm	
11:10	DM227	Matrix Profile for Anomaly Detection on Multidimensional Time Series	S
		Chin-Chia Michael Yeh, Audrey Der, Uday Singh Saini, Vivian Lai, Yan Zheng, Junpeng Wang, Xin Dai, Zhongfang Zhuang, Yujie Fan, Huiyuan Chen, Prince Aboagye, Liang Wang, Wei Zhang, and Eamonn Keogh	
11:25	DM271	CL4CO: A Curriculum Training Framework for Graph-based Neural Combinatorial Optimization	S
		Yang Liu, Chuan Zhou, Peng Zhang, Zhao Li, Shuai Zhang, Xixun Lin, and Xindong Wu	

Session A2-1 Deep learning and statistical methods for data mining

Room CS 7, 10:30-11:45

Session Chair: Flavio Giobergia, Politecnico di Torino, flavio.giobergia@polito.it

10:30	DM211	Generating Realistic Tabular Data with Large Language Model	R
-------	-------	---	---

		Dang Nguyen, Sunil Gupta, Kien Do, Thin Nguyen, and Svetha Venkatesh	
10:50	DM245	HyperTime: A Dynamic Hypergraph Approach for Time Series Classification	R
		Raneen Younis and Zahra Ahmadi	
11:10	DM301	Improving Time Series Encoding with Noise-Aware Self-Supervised Learning and an Efficient Encoder	R
		Duy Nguyen Anh, Trang Tran, Hieu Pham Huy, Le Nguyen Phi, and Lam Nguyen Minh	
11:30	DM295	QUCE: The Minimisation and Quantification of Path-Based Uncertainty for Generative Counterfactual Explanations	S
		Jamie Duell, Monika Seisenberger, Hsuan Fu, and Xiuyi Fan	

Session A3-1 Mining from heterogeneous data sources

Room CS 9, 10:30-11:40

Session Chair: Yue He, Tsinghua University, heyuethu@mail.tsinghua.edu.cn

10:30	DM216	Graph Community Augmentation with GMM-based Modeling in Latent Space	R
		Shintaro Fukushima and Kenji Yamanishi	
10:50	DM233	Solving Combinatorial Optimization Problem over Graph through QUBO Transformation and Deep Reinforcement Learning	R
		Tianle Pu, Chao Chen, Li Zeng, Shixuan Liu, Rui Sun, and Changjun	

		Fan	
11:10	DM384	Exploratory Combinatorial Optimization Problem Solving via Gauge Transformation	S
		Tianle Pu, Changjun Fan, Mutian Shen, Yizhou Lu, Li Zeng, Zohar Nussinov, Chao Chen, and Zhong Liu	
11:25	DM259	2DXformer: Dual Transformers for Wind Power Forecasting with Dual Exogenous Variables	S
		Yajuan Zhang, Jiahai Jiang, Yule Yan, liang Yang, and ping zhang	

Session A5-1 Data mining for modelling, visualization, personalization, and recommendation
Room Hive, 10:30-11:40

Session Chair: Di Wu, Sun Yat-Sen University, China, wudi27@mail.sysu.edu.cn

10:30	DM410	Contrastive Learning for Adapting Language Model to Sequential Recommendation	R
		Fei-Yao Liang, Wu-Dong Xi, Xing-Xing Xing, Wei Wan, Chang-Dong Wang, Min Chen, and Mohsen Guizani	
10:50	DM419	Cross-Store Next-Basket Recommendation	R
		Liangchen Ma, Ya Li, Zifeng Mai, Feiyao Liang, Chang-Dong Wang, Min Chen, and Mohsen Guizani	
11:10	DM517	DiffFaiRec: Generative Fair Recommender with Conditional Diffusion Model	S

		Zhenhao Jiang and Jicong Fan	
11:25	DM663	A Parameter Update Balancing Algorithm for Multi-task Ranking Models in Recommendation Systems	S
		Jun Yuan, Guohao Cai, and Zhenhua Dong	

Session A4-1 Data mining systems and platforms

Room CS 5, 13:30-15:00

Session Chair: Juan Garcia, Universidad de Guayaquil, juan.garciap1@ug.edu.ec

13:30	DM366	Designing an attack-defense game: how to increase the robustness of financial transaction models via a competition	R
		Alexey Zaytsev, Alex Natekin, Evgeni Vorsin, Valerii Smirnov, Georgii Smirnov, Oleg Sidorshin, Alexander Senin, Alexander Dudin, Maria Kovaleva, and Dmitry Berestnev	
13:50	DM409	Scaling Disk Failure Prediction via Multi-Source Stream Mining	R
		Shujie Han, Zirui Ou, Qun Huang, and Patrick P. C. Lee	
14:10	DM455	APOLLO: Differential Private Online Multi-Sensor Data Prediction with Certified Performance	R
		Honghui Xu, Wei Li, Shaoen Wu, Liang Zhao, and Zhipeng Cai	
14:30	DM559	FGLBA: Enabling Highly-Effective and Stealthy Backdoor Attack on Federated Graph Learning	S
		Qing Lu, Miao Hu, Di Wu, Yipeng Zhou, Mohsen Guizani, and Quan Z. Sheng	
14:45	DM583	Enhancing Entity Alignment on Probabilistic Knowledge Graphs	S
		Yunfei Li, Lu Chen, Chengfei Liu, Rui Zhou, and Jianxin Li	

Session A5-2 Data mining for modelling, visualization, personalization, and recommendation

Room CS 7, 13:30-15:00

Session Chair: Yejing Wang, City University of Hong Kong, yejing.wang@my.cityu.edu.hk

13:30	DM277	Transitivity-Encoded Graph Attention Networks for Complementary Item Recommendations	R
		Jin Shang, Yang Jiao, Chenghuan Guo, Minghao Sun, Yan Gao, Jia Liu, Michinari Momma, Itetsu Taru, and Yi Sun	
13:50	DM288	SR-PredictAO: Session-based Recommendation with High-Capability Predictor Add-On	R
		Ruida WANG, Raymond Chi-Wing Wong, and Weile TAN	
14:10	DM331	Enhancing Embeddings Quality with Stacked Gate for Click-Through Rate Prediction	R
		Caihong Mu, Yunfei Fang, Jialiang Zhou, and Yi Liu	
14:30	DM241	Hi-Gen: Generative Retrieval For Large-Scale Personalized E-commerce Search	S
		YanjingWu Wu, Yinfu Feng, Jian Wang, Wenji Zhou, Yunan Ye, Rong Xiao, and Jun Xiao	
14:45	DM343	Exploitation or Exploration Next? User Behavior Decoupling and Emerging Intent Modeling for Next-Item Recommendation	S
		Nengjun Zhu, Lingdan Sun, Xiangfeng Luo, Jian Cao, Qi Zhang, and Xinjiang Lu	

Session A6-1 Applications of data mining

Room CS 9, 13:30-15:00

Session Chair: Meikang Qiu, Augusta University, qiumeikang@ieee.org

13:30	DM254	Towards Efficient Ridesharing via Order-Vehicle Pre-Matching Using Attention Mechanism	R
		Zhidan Liu, Jinye Lin, Zhiyu Xia, Chao Chen, and Kaishun Wu	

13:50	DM270	DFDG: Data-Free Dual-Generator Adversarial Distillation for One-Shot Federated Learning	R
		Kangyang Luo, Shuai Wang, Yexuan Fu, Renrong Shao, Xiang Li, Yunshi Lan, Ming Gao, and Jinlong Shu	
14:10	DM359	Debunking Fake News in Online Social Networks without Text Analysis	R
		Xing Su, Jian Yang, Jia Wu, and Zitai Qiu	
14:30	DM266	Goal-guided Generative Prompt Injection Attack on Large Language Models	S
		Chong Zhang, Mingyu Jin, Qinkai Yu, Chengzhi Liu, Haochen Xue, and Xiaobo Jin	
14:45	DM385	SplitSEE: A Splittable Self-supervised Framework for Single-channel EEG Representation Learning	S
		Rikuto Kotoge, Zheng Chen, Tasuku Kimura, Yasuko Matsubara, Takufumi Yanagisawa, Haruhiko Kishima, and Yasushi Sakurai	

Session A3-2 Mining from heterogeneous data sources

Room Hive, 13:30-15:00

Session Chair: Djellel Difallah, NYU Abu Dhabi, djellel@nyu.edu

13:30	DM388	ELiCiT: Effective and Lightweight Lossy Compression of Tensors	R
		Jihoon Ko, Taehyung Kwon, Jinhong Jung, and Kijung Shin	
13:50	DM393	LISA: Learning-Integrated Space Partitioning Framework for Traffic Accident Forecasting on Heterogeneous Spatiotemporal Data	R
		Bang An, Xun Zhou, Amin Khezerlou, Nick Street, Jinping Guan, and Jun Luo	

14:10	DM430	Emotional Synchronization for Audio-Driven Talking-Head Generation	R
		Zhao Zhang, Yan Luo, Zhichao Zuo, Richang Hong, Yi Yang, and Meng Wang	
14:30	DM605	SemiFDA: Domain Adaptation in Semi-Supervised Federated Learning	S
		Michele Craighero, Giorgio Rossi, Beatrice Rossi, Diego Carrera, Diego Stucchi, Pasqualina Fragneto, and Giacomo Boracchi	
14:45	DM649	Controllable Visit Trajectory Generation with Spatiotemporal Constraints	S
		Haowen Lin, John Krumm, Cyrus Shahabi, and Li Xiong	

Session A1-2 Foundations, algorithms, models, and theory of data mining

Room CS 5, 15:30-17:00

Session Chair: Yüewen Sun, MBZUAI, yüewen.sun@mbzuai.ac.ae

15:30	DM378	Probabilistic Matrix Factorization-based Three-stage Label Completion for Crowdsourcing	R
		Boyi Yang, Liangxiao Jiang, and Wenjun Zhang	
15:50	DM413	HomoMGC: Homophily-enhanced Adaptive Graph Refinement for Multi-view Graph Clustering	R
		Man-Sheng Chen, Xiao-Sha Cai, Chang-Dong Wang, Dong Huang, Min Chen, and Mohsen Guizani	
16:10	DM442	GADIN: Generative Adversarial Denoise Imputation Network for Incomplete Data	R
		Dong Li, Zhicong Liu, Mingfeng Hu, Baoyan Song, and Xiaohuan	

		Shan	
16:30	DM325	Generalized Sparse Additive Model with Unknown Link Function	S
		Peipei Yuan, Xinge You, Hong Chen, Xuelin Zhang, and Qinmu Peng	
16:45	DM462	Towards Expressive Graph Representations for Graph Neural Networks	S
		Chengsheng Mao, Liang Yao, and Yuan Luo	

Session A2-2 Deep learning and statistical methods for data mining

Room CS 7, 15:30-17:00

Session Chair: Evgenii Tsymbalov, Amazon, etsymbalov@gmail.com

15:30	DM323	Graph Contrastive Learning with Adversarial Structure Refinement (GCL-ASR)	R
		Jiangwen Chen, Kou Guang, Qiyang Li, and Tan Hao	
15:50	DM412	GQ*: Towards Generalizable Deep Q-Learning for Steiner Tree in Graphs	R
		Wei Huang, Hanchen Wang, Dong Wen, Xuefeng Chen, Wenjie zhang, and Ying Zhang	
16:10	DM315	Hierarchical Explanations for Text Classification Models: Fast and Effective	R
		Zhenyu Nie, Zheng Xiao, Huizhang Luo, Xuan Liu, and Anthony Theodore Chronopoulos	
16:30	DM449	Channel-Attentive Graph Neural Networks	S

		Tuğrul Hasan Karabulut and İnci M. Baytaş	
16:45	DM580	Cascading Multimodal Feature Enhanced Contrast Learning for Music Recommendation	S
		Qimeng Yang, Shijia Wang, Da Guo, Dongjin Yu, Qiang Xiao, Dongjing Wang, and Chuanjiang Luo	

Session A3-3 Mining from heterogeneous data sources

Room CS 9, 15:30-17:00

Session Chair: Guangyi Chen, MBZUAI, Guangyi.Chen@mbzuai.ac.ae

15:30	DM327	Adaptive Loss-aware Modulation for Multimedia Retrieval	R
		Jian Zhu, Yu Cui, Zeyi Sun, Yuyang Dai, Xi Wang, Lei Liu, Cheng Luo, and Li-Rong Dai	
15:50	DM337	Towards Cross-domain Few-shot Graph Anomaly Detection	R
		Jiazhen Chen, Sichao Fu, Zhibin Zhang, Zheng Ma, Mingbin Feng, Tony Wirjanto, and Qinmu Peng	
16:10	DM383	Informative Subgraphs Aware Masked Auto-Encoder in Dynamic Graphs	R
		Pengfei Jiao, Xinxun Zhang, Mengzhou Gao, and Tianpeng Li	
16:30	DM320	A Momentum Contrastive Learning Framework for Query-POI Matching	S
		Yuting Qiang, Jianbin Zheng, Lixia Wu, Haomin Wen, Junhong Lou, and Minhui Deng	

16:45	DM371	Multi-modal Sarcasm Detection via Dual Synergetic Perception Graph Convolutional Networks	S
		Xingjie Zhuang and Zhixin Li	

Session A6-2 Applications of data mining

Room Hive, 15:30-17:00

Session Chair: Xun Zhou, Harbin Institute of Technology, Shenzhenzhouxun2023@hit.edu.cn

15:30	DM690	Dual Cross-Stage Partial Learning for Enhanced Object Detection in Dehazed Images	R
		Jinbiao Zhao, Zhao Zhang, Jiahuan Ren, Haijun Zhang, Zhongqiu Zhao, and Meng Wang	
15:50	DM697	Resource2Box: Learning To Rank Resources in Distributed Search Using Box Embedding	R
		Ulugbek Ergashev, Geon Lee, Kijung Shin, Eduard Dragut, and Weiyi Meng	
16:10	DM709	ChronoCTI: Mining Knowledge Graph of Temporal Relations among Cyberattack Actions	R
		Md Rayhanur Rahman, Brandon Wroblewski, Quinn Matthews, Brantley Morgan, Timothy Menzies, and Laurie Williams	
16:30	DM749	Addressing Delayed Feedback in Conversion Rate Prediction: A Domain Adaptation Approach	S
		Leisheng Yu, Yanxiao Cai, Lucas Chen, Minxing Zhang, Wei-Yen Day, Li Li, Rui Chen, Soo-Hyun Choi, and Xia Hu	
16:45	DM753	Hypergraph-Enhanced Contrastively Regularized Transformer for Multi-Behavior E-commerce Product Recommendation	S
		Shuiying Liao and P. Y. Mok	

Day 3: December 11, 2024

Session A1-3 Foundations, algorithms, models, and theory of data mining

Room CS 5, 09:00-10:00

Session Chair: Mengyue Yang, Bristol University, mengyue.yang.20@ucl.ac.uk

09:00	DM363	Scalable Order-Preserving Pattern Mining	R
		Ling Li, Wiktor Zuba, Grigorios Loukides, Solon Pissis, and Maria Matsangidou	
09:20	DM546	Efficiently Manipulating Structural Graph Clustering Under Jaccard Similarity	R
		Chuanyu Zong, Rui Fang, Meng-xiang Wang, Tao Qiu, and Anzhen Zhang	
09:40	DM617	IIFE: Interaction Information Based Automated Feature Engineering	S
		Tom Overman, Diego Klabjan, and Jean Utke	

Session A3-4 Mining from heterogeneous data sources

Room CS 7, 09:00-10:00

Session Chair: Haoxuan Li, Peking University, hxli@stu.pku.edu.cn

09:00	DM322	Adaptive Graph Neural Networks for Cold-start Multimedia Recommendation	R
		Zhen Li, Jibin Wang, Zhuo Chen, Kun Wu, Yuanzhen Wei, and Hai Huang	
09:20	DM482	EEiF: Efficient Isolated Forest with e Branches for Anomaly Detection	R
		Yifan Zhang, Haolong Xiang, Xuyun Zhang, Xiaolong Xu, Wei Fan, Qin	

		Zhang, and Lianyong Qi	
09:40	DM223	MetaSTC: A Meta Spatio-Temporal Learning Paradigm for Traffic Flow Prediction	S
		Kexin Xu, Zhemeng Yu, Yucen Gao, Songjian Zhang, Jun Fang, Xiaofeng Gao, and Guihai Chen	

Session A5-3 Data mining for modelling, visualization, personalization, and recommendation

Room CS 9, 09:00-10:00

Session Chair: Przemyslaw Kazienko, Wroclaw Tech, kazienko@pwr.edu.pl

09:00	DM438	Early Fire Detection based on Local Morphological Knowledge Matching	R
		Xinzhi Wang, Mengyue Li, Nengjun Zhu, Jiayan Qian, and Zhanyi Zheng	
09:20	DM402	RecCoder: Reformulating Sequential Recommendation as Large Language Model-Based Code Completion	R
		Kai-Huang Lai, Wudong Xi, Xingxing Xing, Wei Wan, Chang-Dong Wang, Min Chen, and Mohsen Guizani	
09:40	DM726	ExoTST: Exogenous-Aware Temporal Sequence Transformer for Time Series Prediction	S
		Kshitij Tayal, Arvind Renganathan, Xiaowei Jia, Vipin Kumar, and Dan Lu	

Session A6-3 Applications of data mining

Room Hive, 09:00-10:00

Session Chair: Maurizio Atzori, University of Cagliari, atzori@unica.it

09:00	DM655	Financial Risk Assessment via Long-term Payment Behavior Sequence Folding	R
		Yiran Qiao, Yateng Tang, Xiang Ao, Qi Yuan, Ziming Liu, Chen Shen, and Xuehao Zheng	
09:20	DM743	Adaptive Process-Guided Learning: An Application in Predicting Lake DO Concentrations	R
		Runlong Yu, Chonghao Qiu, Robert Ladwig, Paul Hanson, Yiqun Xie, Yanhua Li, and Xiaowei Jia	
09:40	DM334	Interdependency Matters: Graph Alignment for Multivariate Time Series Anomaly Detection	S
		Yuanyi Wang, Haifeng Sun, Chengsen Wang, Mengde Zhu, Wei Tang, Jingyu Wang, Qi Qi, Zirui Zhuang, and Jianxin Liao	

Session A5-4 Data mining for modelling, visualization, personalization, and recommendation
Room CS 5, 10:30-11:35

Session Chair: Parham Moradi, RMIT University, parham.moradi@rmit.edu.au

10:30	DM367	Continuous Exact Explanations of Neural Networks	R
		Alice Dethise and Marco Canini	
10:50	DM414	Periodic Prompt on Dynamic Heterogeneous Graph for Next Basket Recommendation	S
		Ru-Bin Li, Man-Sheng Chen, Xin-Yu Ding, Chang-Dong Wang, Sihong Xie, Shuangyin Liu, Min Chen, and Mohsen Guizani	
11:05	DM454	A Condensed Transition Graph Framework for Zero-shot Link Prediction with Large Language Models	S
		Mingchen Li, Chen Ling, rui Zhang, and Liang Zhao	
11:20	DM495	Influence-aware Group Recommendation for Social Media	S

		Propagation	
		Chengkun He, Xiangmin Zhou, Chen Wang, Longbing Cao, Jie Shao, and Zahir Tari	

Session A6-4 Applications of data mining

Room CS 7, 10:30-11:40

Session Chair: Flavio Giobergia, Politecnico di Torino, flavio.giobergia@polito.it

10:30	DM373	Utilitarian Online Learning from Open-World Soft Sensing	R
		Heng Lian, Yu Huang, Xingquan Zhu, and Yi He	
10:50	DM641	CounterFair: Group Counterfactuals for Bias Detection, Mitigation and Subgroup Identification	R
		Alejandro Kuratomi, Zed Lee, Panayiotis Tsaparas, Guilherme Dinis Junior, Evaggelia Pitoura, Tony Lindgren, and Panagiotis Papapetrou	
11:10	DM573	D-Cube : Exploiting Hyper-Features of Diffusion Model for Robust Medical Classification	S
		Minhee Jang, Juheon Son, Thanaporn Viriyasaranon, Junho Kim, and Jang-hwan Choi	
11:25	DM648	Survival Analysis with Multiple Noisy Labels	S
		Donna Tjandra and Jenna Wiens	

Session A1-4 Foundations, algorithms, models, and theory of data mining

Room CS 9, 10:30-11:40

Session Chair: Mubarak Gwaza Abdu-Aguye, MBZUAI, Mubarak.Abdu-Aguye@mbzuai.ac.ae

10:30	DM488	Margin-bounded Confidence Scores for Out-of-Distribution Detection	R
-------	-------	--	---

		Lakpa Tamang, Mohamed Reda Bouadjenek, Richard Dazeley, and Sunil Aryal	
10:50	DM515	Fast and Accurate Triangle Counting in Graph Streams Using Predictions	R
		Cristian Boldrin and Fabio Vandin	
11:10	DM390	Accurate and Fast Estimation of Temporal Motifs using Path Sampling	S
		Yunjie Pan, Omkar Bhalerao, C. Seshadhri, and Nishil Talati	
11:25	DM326	SHADE: Deep Density-based Clustering	S
		Anna Beer, Pascal Weber, Lukas Miklautz, Collin Leiber, Walid Durani, Christian Böhm, and Claudia Plant	

Session A2-3 Deep learning and statistical methods for data mining

Room CHB, 10:30-11:45

Session Chair: Evgenii Tsymbalov, Amazon, etsymbalov@gmail.com

10:30	DM461	Combining Self-Supervision and Privileged Information for Representation Learning from Tabular Data	R
		Haoyu Yang, Gyorgy Simon, Michael Steinbach, Genevieve Melton, and Vipin Kumar	
10:50	DM510	Towards Dynamic University Course Timetabling Problem: An Automated Approach Augmented via Reinforcement Learning	R
		Yanan Xiao, XiangLin Li, Lu Jiang, Pengfei Wang, Kaidi Wang, and Na Luo	

11:10	DM591	HFGNN: Efficient Graph Neural Networks using Hub-Fringe Structures	R
		Pak Lon Ip, Sheng Hui Zhang, Xue Kai Wei, Tsz Nam Chan, and Leong Hou U	
11:30	DM628	Unsupervised Domain Adaptation for Action Recognition via Self-Ensembling and Conditional Embedding Alignment	S
		Indrajeet Ghosh, Garvit Chugh, Abu Zaher Md Faridee, and Nirmalya Roy	

Women forum

Day 3: December 11, 2024	
Women Forum	
10:30- 12:00	
Room: Room Hive	
Co-chairs : Prof. Xiaochun Yang & Prof. Xiaofeng Gao	
	Forum Opening
10:30	Prof. Kun Zhang, Program Committee Co-Chair
	Warm-Up Speech (with Personal Experience Sharing)
10:35	Prof. Xiaochun Yang & Prof. Xiaofeng Gao, Women Forum Co-Chairs
	Intelligent Knowledge Discovery—Explorations in Talent Analytics
10:50	Dr. Ying Sun The Hong Kong University of Science and Technology, Guangzhou, China
	Breaking Barriers in Time Series Analysis
11:10	Dr. Zahra Ahmadi Hannover Medical School, Germany
	Exploring Data Science: A Personal Journey
11:30	Ms. Li Qian Ludwig-Maximilians-Universität München, Germany

11:45	The Research on Machine Learning for Data Management
	Ms. Chaohong Ma Renmin University of China
12:00	Closing Speech
	Prof. Elena Baralis, Program Committee Co-Chair

Day 4: December 12, 2024

Session A3-5 Mining from heterogeneous data sources

Room CS 5, 10:30-11:40

Session Chair: Djellel Difallah, NYU Abu Dhabi, djellel@nyu.edu

10:30	DM436	High-Fidelity Diffusion Editor for Zero-Shot Text-Guided Video Editing	R
		Yan Luo, Zhichao Zuo, Zhao Zhang, Zhongqiu Zhao, Haijun Zhang, and Richang Hong	
10:50	DM475	Align Along Time and Space: A Graph Latent Diffusion Model for Traffic Dynamics Prediction	R
		Yuhang Liu, Yingxue Zhang, Xin Zhang, Yu Yang, Yiqun Xie, Sahar Ghanipoor Machiani, Yanhua Li, and Jun Luo	
11:10	DM483	Futures Quantitative Investment with Heterogeneous Continual Graph Neural Network	S
		Zhizhong Tan, Min Hu, Bin Liu, and Guosheng Yin	
11:25	DM497	Multi-Hyperbolic Space-based Heterogeneous Graph Attention Network	S

		Jongmin Park, Seunghoon Han, Jong-Ryul Lee, and Sungsu Lim
--	--	--

Session A5-5 Data mining for modelling, visualization, personalization, and recommendation
Room CS 7, 10:30-11:25
Session Chair: Shirui Pan, Griffith University, s.pan@griffith.edu.au

10:30	DM747	DISCO: A Hierarchical Disentangled Cognitive Diagnosis Framework for Interpretable Job Recommendation	R
		Xiaoshan Yu, Chuan Qin, Qi Zhang, Chen Zhu, Haiping Ma, Xingyi Zhang, and Hengshu Zhu	
10:50	DM778	Bi-level User Modeling for Deep Recommender Systems	R
		Yejing Wang, Dong Xu, Xiangyu Zhao, Zhiren Mao, Peng Xiang, Ling Yan, Yao Hu, Zijian Zhang, Xuetao Wei, and Qidong Liu	
11:10	DM708	An Explainable Recommender System by Integrating Graph Neural Networks and User Reviews	S
		Sahar Batmani, Parham Moradi, Narges Haidari, and Mahdi Jalili	

Session A6-5 Applications of data mining
Room CS 9, 10:30-11:25
Session Chair: Ling Chen, University of Technology Sydney, ling.chen@uts.edu.au

10:30	DM806	A Learned Approach to Index Algorithm Selection	R
		Chaohong Ma, Xiaohui Yu, Yifan Li, Aishan Maolinyazi, and Xiaofeng Meng	

10:50	DM772	TAN: A Tripartite Alignment Network Enhancing Composed Image Retrieval with Momentum Distillation	R
		Yongquan Wan, Erhe Yang, Cairong Yan, Guobing Zou, and Bofeng Zhang	
11:10	DM604	AnomalyLLM: Few-shot Anomaly Edge Detection for Dynamic Graphs using Large Language Models	S
		Shuo Liu, Yao Di, Lanting Fang, Zhetao Li, Wenbin Li, Kaiyu Feng, Xiaowen Ji, and Jingping Bi	

Session A1-5 Foundations, algorithms, models, and theory of data mining

Room Hive, 10:30-11:40

Session Chair: Shaoan Xie, Carnegie Mellon University, shaoan@cmu.edu

10:30	DM667	Scalable Graph Classification via Random Walk Fingerprints	R
		Peiyan Li, Honglian Wang, and Christian Böhm	
10:50	DM717	Warm-Starting Contextual Bandits under Latent Reward Scaling	R
		Bastian Oetomo, R. Malinga Perera, Renata Borovica-Gajic, and Benjamin I. P. Rubinstein	
11:10	DM446	Constructing ϵ -Constrained Sparsified β^s -Complexes using Space Partitioning Trees	S
		Rohit Singh and Philip Wilsey	
11:25	DM394	DynoGraph: Dynamic Graph Construction for Nonlinear Dimensionality Reduction	S
		Li Qian, Claudia Plant, Yalan Qin, Jing Qian, and Christian Böhm	

Session A1-6 Foundations, algorithms, models, and theory of data mining

Room CS 5, 13:30-14:55

Session Chair: Jiuyong Li, University of South Australia, jiuyong.li@unisa.edu.au

13:30	DM776	PROMIPL:A Probabilistic Generative Model for Multi-Instance Partial-Label Learning	R
		Yin-Fang Yang, Wei Tang, and Min-Ling Zhang	
13:50	DM783	A Novel Shadow Variable Catcher for Addressing Selection Bias in Recommendation Systems	R
		Qingfeng Chen, Boquan Wei, Debo Cheng, Jiuyong Li, Lin Liu, and Shichao Zhang	
14:10	DM672	Reducing Unfairness in Distributed Community Detection	S
		Hao Zhang, Malith Jayaweera, Bin Ren, Yanzhi Wang, and Sucheta Soundarajan	
14:25	DM780	An Efficient Graph Autoencoder with Lightweight Desmoothing Decoder and Long-Range Modeling	S
		Jinyong Wen, Tao Zhang, Chunxia Zhang, Shiming Xiang, Chunhong Pan	
14:40	DM798	MoRE-LLM: Mixture of Rule Experts Guided by a Large Language Model	S
		Alexander Koebler, Ingo Thon, and Florian Buettner	

Session A2-5 Deep learning and statistical methods for data mining

Room CS 7, 13:30-15:00

Session Chair: Chuan Zhou, Peking University, zhouchuancn@pku.edu.cn

13:30	DM634	Counterfactual Brain Graph Augmentation Guided Bi-Level Contrastive Learning for Disorder Analysis	R
		Guangwei Dong, Xuexiong Luo, Jing Du, Jia Wu, Shan Xue, Jian Yang, and Amin Beheshti	
13:50	DM734	Feature Map Purification for Enhancing Adversarial Robustness of Deep Timeseries Classifiers	R
		Mubarak Abdu-Aguye, Zaigham Zaheer, and Karthik Nandakumar	
14:10	DM790	EMIT - Event Based Masked Auto Encoding for Irregular Time Series	R
		Hrshikesh Patel, Ruihong Qiu, Adam Irwin, Shazia Sadiq, and Sen Wang	
14:30	DM809	PC3: Enhancing Concurrency in High-Conflict Transactions with Prior Cascading Control	S
		Zhibin Wang, Jiangtao Cui, Xiyue Gao, Hui Zhang, Guiqi Ren, Yixiao Liu, Hui Li, and Kankan Zhao	
14:45	DM795	Handling Non-IID Data in Federated Learning Using Metaheuristic Optimization Techniques	S
		Amin Birashk, Sadaf MD Halim, and Latifur Khan	

Session A3-6 Mining from heterogeneous data sources

Room CS 9, 13:30-15:00

Session Chair: Kijung Shin, KAIST, kijungs@kaist.ac.kr

13:30	DM713	Traffic Pattern Sharing for Federated Traffic Flow Prediction with Personalization	R
		Hang Zhou, Wentao Yu, Sheng Wan, Yongxin Tong, Tianlong Gu, and Chen Gong	

13:50	DM745	TROPICAL: Transformer-based Hypergraph Learning for Camouflaged Fraudsters Detection	R
		Venus Haghighi, Behnaz Soltani, Nasrin Shabani, Jia Wu, Yang Zhang, Lina Yao, Quan Z. Sheng, and Jian Yang	
14:10	DM760	MOSTream: A Modular and Self-Optimizing Data Stream Clustering Algorithm	R
		Zhengru Wang, Xin Wang, and Shuhao Zhang	
14:30	DM467	Weakly-Supervised Graph Classification with Even a Single Key Subgraph Per Class	S
		Lu Zhang, Chenbo Zhang, Jihong Guan, and Shuigeng Zhou	
14:45	DM681	Graph Rhythm Network: Beyond Energy Modeling for Deep Graph Neural Networks	S
		Yufei Jin and Xingquan Zhu	

Session A2-4 Deep learning and statistical methods for data mining

Room Hive, 13:30-15:00

Session Chair: Omkar Bhalerao , University of California, Santa Cruz, obhalera@ucsc.edu

13:30	DM610	A Bayesian Hierarchical Model for Orthogonal Tucker Decomposition with Oblivious Tensor Compression	R
		Matthew Pietrosanu, Bei Jiang, and Linglong Kong	
13:50	DM611	Normalizing self-supervised learning for provably reliable Change Point Detection	R

		Alexandra Bazarova, Evgenia Romanenkova, and Alexey Zaytsev	
14:10	DM729	Enhancing Distribution and Label Consistency for Graph Out-of-Distribution Generalization	S
		Song Wang, Xiaodong Yang, Rashidul Islam, Huiyuan Chen, Minghua Xu, Jundong Li, and Yiwei Cai	
14:25	DM741	CAKD: A Correlation-Aware Knowledge Distillation Framework Based on Decoupling Kullback-Leibler Divergence	S
		Zao Zhang, Huaming Chen, Pei Ning, Nan Yang, and Dong Yuan	
14:40	DM812	Rank Supervised Contrastive Learning for Time Series Classification	S
		Qianying Ren, Dongsheng Luo, and Dongjin Song	

Tutorials

Tutorial 1: Causality and Large Models

Presenters: Haoxuan Li, Chuan Zhou, Mengyue Yang, Mingming Gong, Jun Wang, Xiao-Hua Zhou

Abstract: Our tutorial aims to explore the synergies between causality and large models, also known as “foundation models,” which have demonstrated remarkable capabilities across for helping data mining in healthcare, finance, and education. However, there are increasingly concerns about the trustworthy and interpretability of these complex “black-box” LLMs behind the promising performance in data mining domains. A growing community of researchers is turning towards a more principled framework to address these concerns, better understand the behavior of large models, and improve their reliability and interpretability. Specifically, this tutorial will focus on three directions: causal agents for decision-making, LLMs for causality, and benefiting LLMs with causality. Besides, we introduce some open challenges and potential future directions for this area. We hope this tutorial could stimulate more ideas on this topic and facilitate the development of causality-aware large models.

Duration: One whole Day

Tutorial 2: Hypergraph Neural Networks: An In-Depth and Step-by-Step Guide

Presenters: Sunwoo Kim, Soo Yong Lee, Yue Gao, Alessia Antelmi, Mirko Polato, Kijung Shin

Abstract: Higher-order interactions (HOIs) are ubiquitous in real-world networks. Investigation of deep learning for networks of HOIs, expressed as hypergraphs, has become an important agenda for the data mining and machine learning communities. Thus, hypergraph neural networks (HNNs) have emerged as a powerful tool for representation learning on hypergraphs. Given the emerging trend, we provide a timely tutorial dedicated to HNNs. We cover the (1) inputs, (2) message passing schemes, (3) training strategies, (4) applications (e.g., recommender systems and time series analysis), and (5) open problems of HNNs. This tutorial is intended for researchers and practitioners who are interested in (hyper)graph representation learning and its applications.

Duration: Half-a-day

Tutorial 3: Uncertain Boundaries: Multidisciplinary Approaches to Copyright Issues in Generative AI

Presenters: As generative AI systems become more prevalent in creative fields, concerns about intellectual property rights have grown, particularly regarding the production of content that closely resembles human-created work. Recent controversies, where AI models have generated near-replicas of copyrighted material, underscore the urgency of reviewing the current copyright framework and developing methods to mitigate infringement risks. To this end, this tutorial offers a comprehensive analysis of these copyright challenges, examining them throughout the AI development life cycle and providing developers with actionable strategies. It begins by discussing the foundational goals and considerations for copyright in generative AI, followed by methods for detecting and assessing potential violations in AI outputs. Next, it introduces techniques to safeguard creative works and datasets from unauthorized replication. The tutorial also covers training methods aimed at minimising the risk of AI models reproducing protected content. Finally, it reviews the state of AI copyright regulation and suggests future research pathways to address existing gaps.

Duration: Half-a-day

Keynotes

Keynote 1: Preslav Nakov

Title: Towards Safe, Truly Open, and Factual Large Language Models

Abstract: We will discuss several initiatives towards safe, truly open, and factual large language models (LLMs). First, we will present Do-Not-Answer, a dataset for evaluating the guardrails of LLMs, which is at the core of the safety mechanisms incorporated in Jais, the world's leading open Arabic-centric foundation and instruction-tuned large language model, and Nanda, our recently released open Hindi LLM. Next, we will discuss the LLM360 initiative of MBZUAI's Institute on Foundation Models, aiming at developing fully transparent open-source LLMs. We

will then examine the factuality challenges associated with large language models, and we will present some recent relevant tools for addressing these challenges developed at MBZUAI: (i) OpenFactCheck, a framework for fact-checking LLM output, for building customized fact-checking systems, and for benchmarking LLMs for factuality, (ii) LM-Polygraph, a tool for predicting an LLM's uncertainty in its output using cheap and fast uncertainty quantification techniques, and (iii) LLM-DetectAlve, a tool for machine-generated text detection.

Bio: Preslav Nakov is Professor and Department Chair for NLP at the Mohamed bin Zayed University of Artificial Intelligence. He is part of the core team at MBZUAI's Institute for Foundation Models that developed Jais, the world's best open-source Arabic-centric LLM, Nanda, the world's best Hindi model, and LLM360, the first truly open LLM. Previously, he was Principal Scientist at the Qatar Computing Research Institute, HBKU, where he led the Tanbih mega-project, developed in collaboration with MIT, which aims to limit the impact of "fake news", propaganda and media bias by making users aware of what they are reading, thus promoting media literacy and critical thinking. He received his PhD degree in Computer Science from the University of California at Berkeley, supported by a Fulbright grant. He is Chair-Elect of the European Chapter of the Association for Computational Linguistics (EACL), Secretary of ACL SIGSLAV, and Secretary of the Truth and Trust Online board of trustees. Formerly, he was PC chair of ACL 2022, and President of ACL SIGLEX. He is also member of the editorial board of several journals including Computational Linguistics, TACL, ACM TOIS, IEEE TASL, IEEE TAC, CS&L, NLE, AI Communications, and Frontiers in AI. He authored a Morgan & Claypool book on Semantic Relations between Nominals, two books on computer algorithms, and 250+ research papers. He received a Best Paper Award at ACM WebSci'2022, a Best Long Paper Award at CIKM'2020, a Best Resource Paper Award at EACL'2024, a Best Demo Paper Award (Honorable Mention) at ACL'2020, a Best Task Paper Award (Honorable Mention) at SemEval'2020, a Best Poster Award at SocInfo'2019, and the Young Researcher Award at RANLP'2011. He was also the first to receive the Bulgarian President's John Atanasoff award, named after the inventor of the first automatic electronic digital computer. His research was featured by over 100 news outlets, including Reuters, Forbes, Financial Times, CNN, Boston Globe, Aljazeera, DefenseOne, Business Insider, MIT Technology Review, Science Daily, Popular Science, Fast Company, The Register, WIRED, and Engadget, among others.

Photo: <https://mbzuai.ac.ae/study/faculty/preslav-nakov/>

Keynote 2: Bernhard Schölkopf

Title: Towards causal world models and digital twins

Abstract: Research on understanding and building artificially intelligent systems has moved from symbolic approaches to statistical learning, and is now beginning to study interventional models relying on concepts of causality. Some of the hard open problems of machine learning and AI are intrinsically related to causality, and progress may require advances in our

understanding of how to model and infer causality from data, as well as conceptual progress on what constitutes a causal representation and a causal world model. I will present basic concepts and thoughts, and some applications to astronomy.

Bio: Bernhard Schölkopf's scientific interests are in machine learning and causal inference. He has applied his methods to a number of different fields, ranging from biomedical problems to computational photography and astronomy. Bernhard studied physics and mathematics and earned his Ph.D. in computer science in 1997, becoming a Max Planck director in 2001. He has (co-)received the Berlin-Brandenburg Academy Prize, the Royal Society Milner Award, the Leibniz Award, the BBVA Foundation Frontiers of Knowledge Award, and the ACM AAAI Allen Newell Award. He is Fellow of the CIFAR Program "Learning in Machines and Brains", and a Professor at ETH Zurich. He helped start the MLSS series of Machine Learning Summer Schools. In 2023, he founded the ELLIS Institute Tuebingen, and acts as its scientific director.

Keynote 3: Claudia Plant

Title: Clustering: Balancing Abstraction and Representation

Abstract: How to find a natural grouping of a large real data set? Clustering requires a balance between abstraction and representation. To identify clusters, we need to abstract from superfluous details of individual objects. But we also need a rich representation that emphasizes the key features shared by groups of objects that distinguish them from other groups of objects.

Each clustering algorithm implements a different trade-off between abstraction and representation. Classical K-means implements a high level of abstraction - details are simply averaged out - combined with a very simple representation - all clusters are Gaussians in the original data space. We will see how approaches to subspace and deep clustering support high-dimensional and complex data by allowing richer representations. However, with increasing representational expressiveness comes the need to explicitly enforce abstraction in the objective function to ensure that the resulting method performs clustering and not just representation learning. We will see how current deep clustering methods define and enforce abstraction through centroid-based and density-based clustering losses. Balancing the conflicting goals of abstraction and representation is challenging. Ideas from subspace clustering help by learning one latent space for the information that is relevant to clustering and another latent space to capture all other information in the data.

The talk ends with an outlook on future research in clustering. In my view, future methods will more adaptively balance abstraction and representation to improve performance, energy efficiency and interpretability. By automatically finding the sweet spot between abstraction and representation, the human brain is very good at clustering and other related tasks such as single-shot learning. So, there is still much to be explored.

Bio: Claudia Plant is full professor, leader of the Data Mining and Machine Learning research group at the Faculty of Computer Science University of Vienna, Austria. Her group focuses on new methods for exploratory data mining, e.g., clustering, anomaly detection, graph mining and matrix factorization. Many approaches relate unsupervised learning to data compression, i.e. the better the found patterns compress the data the more information we have learned. Other methods rely on finding statistically independent patterns or multiple non-redundant solutions, relying on deep learning or nature-inspired concepts such as synchronization. Indexing techniques and methods for parallel hardware support exploring massive data. Claudia Plant has co-authored over 150 peer-reviewed publications, among them more than 30 contributions to KDD and ICDM and 4 Best Paper Awards. Papers on scalability aspects appeared at SIGMOD, ICDE, and the results of interdisciplinary projects in leading application-related journals such as Bioinformatics, Cerebral Cortex and Water Research.