

Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI-19)

Macao, China
10-16 August 2019

Volume 1 of 9

Editors:

Sarit Kraus

ISBN: 978-1-5108-9474-7

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© (2019) by International Joint Conference on Artificial Intelligence (IJCAI)
All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

Printed with permission by Curran Associates, Inc. (2020)

For permission requests, please contact International Joint Conference on Artificial Intelligence (IJCAI) at the address below.

International Joint Conference on Artificial Intelligence (IJCAI)
Information Sciences Institute/USC 4676
Admiralty Way, Marina del Rey
CA 90292
USA

Phone: +49 (761) 203-8221

Fax: +49 (761) 203-8222

nebel@ijcai.org

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

VOLUME 1

AGENT-BASED AND MULTI-AGENT SYSTEMS

Flexible Representative Democracy: An Introduction with Binary Issues	3
<i>Ben Abramowitz, Nicholas Mattei</i>	
Portioning Using Ordinal Preferences: Fairness and Efficiency	11
<i>Stéphane Airiau, Haris Aziz, Ioannis Caragiannis, Justin Kruger, Jérôme Lang, Dominik Peters</i>	
An Efficient Algorithm for Skeptical Preferred Acceptance in Dynamic Argumentation Frameworks	18
<i>Gianvincenzo Alfano, Sergio Greco, Francesco Parisi</i>	
Strategic Signaling for Selling Information Goods	25
<i>Shani Alkoby, David Sarne, Igal Milchtaich</i>	
Probabilistic Strategy Logic	32
<i>Benjamin Aminof, Marta Kwiatkowska, Bastien Maubert, Aniello Murano, Sasha Rubin</i>	
Multi-Agent Pathfinding with Continuous Time	39
<i>Anton Andreychuk, Konstantin Yakovlev, Dor Atzmon, Roni Stern</i>	
Weighted Maxmin Fair Share Allocation of Indivisible Chores	46
<i>Haris Aziz, Hau Chan, Bo Li</i>	
Fair Allocation of Indivisible Goods and Chores	53
<i>Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, Toby Walsh</i>	
Strategyproof and Approximately Maxmin Fair Share Allocation of Chores	60
<i>Haris Aziz, Bo Li, Xiaowei Wu</i>	
Stable and Envy-free Partitions in Hedonic Games	67
<i>Nathanaël Barrot, Makoto Yokoo</i>	
How Hard Is the Manipulative Design of Scoring Systems?	74
<i>Dorothea Baumeister, Tobias Högrefe</i>	
The Price of Fairness for Indivisible Goods	81
<i>Xiaohui Bei, Xinhang Lu, Pasin Manurangsi, Warut Suksompong</i>	
Strategy Logic with Simple Goals: Tractable Reasoning about Strategies	88
<i>Francesco Belardinelli, Wojciech Jamroga, Damian Kurpiewski, Vadim Malvone, Aniello Murano</i>	
Fairness Towards Groups of Agents in the Allocation of Indivisible Items	95
<i>Nawal Benabbou, Mithun Chakraborty, Edith Elkind, Yair Zick</i>	
Optimality and Nash Stability in Additive Separable Generalized Group Activity Selection Problems	102
<i>Vittorio Bilò, Angelo Fanelli, Michele Flammini, Gianpiero Monaco, Luca Moscardelli</i>	
An Experimental View on Committees Providing Justified Representation	109
<i>Robert Bredereck, Piotr Faliszewski, Andrzej Kaczmarczyk, Rolf Niedermeier</i>	
A Contribution to the Critique of Liquid Democracy	116
<i>Ioannis Caragiannis, Evi Micha</i>	
Be a Leader or Become a Follower: The Strategy to Commit to with Multiple Leaders	123
<i>Matteo Castiglioni, Alberto Marchesi, Nicola Gatti</i>	
On the Problem of Assigning PhD Grants	130
<i>Katarína Cechlářová, Laurent Gourvès, Julien Lesca</i>	
Maximin-Aware Allocations of Indivisible Goods	137
<i>Hau Chan, Jing Chen, Bo Li, Xiaowei Wu</i>	
Reachability and Coverage Planning for Connected Agents	144
<i>Tristan Charrier, Arthur Queffelec, Ocan Sankur, François Schwarzentruber</i>	
Approximately Maximizing the Broker's Profit in a Two-sided Market	151
<i>Jing Chen, Bo Li, Yingkai Li</i>	
Election with Bribe-Effect Uncertainty: A Dichotomy Result	158
<i>Lin Chen, Lei Xu, Shouhuai Xu, Zhimin Gao, Weidong Shi</i>	
Dispatching Through Pricing: Modeling Ride-Sharing and Designing Dynamic Prices	165
<i>Mengjing Chen, Weiran Shen, Pingzhong Tang, Song Zuo</i>	
ATSIS: Achieving the Ad hoc Teamwork by Sub-task Inference and Selection	172
<i>Shuo Chen, Ewa Andrejczuk, Athirai A. Irissappane, Jie Zhang</i>	
Network Formation under Random Attack and Probabilistic Spread	180
<i>Yu Chen, Shahin Jabbari, Michael Kearns, Sanjeev Khanna, Jamie Morgenstern</i>	

Cap-and-Trade Emissions Regulation: A Strategic Analysis	187
<i>Frank Cheng, Yagil Engel, Michael P. Wellman</i>	
A Value-based Trust Assessment Model for Multi-agent Systems	194
<i>Kinzang Chhogyal, Abhaya Nayak, Aditya Ghose, Hoa K. Dam</i>	
Exploiting Social Influence to Control Elections Based on Scoring Rules	201
<i>Federico Corò, Emilio Cruciani, Gianlorenzo D'Angelo, Stefano Ponziani</i>	
Civic Crowdfunding for Agents with Negative Valuations and Agents with Asymmetric Beliefs	208
<i>Sankarshan Damle, Moin Hussain Moti, Praphul Chandra, Sujit Gujar</i>	
Anytime Heuristic for Weighted Matching Through Altruism-Inspired Behavior	215
<i>Panayiotis Danassis, Aris Filos-Ratsikas, Boi Faltings</i>	
AsymDPOP: Complete Inference for Asymmetric Distributed Constraint Optimization Problems	223
<i>Yanchen Deng, Ziyu Chen, Dingding Chen, Wenxin Zhang, Xingqiong Jiang</i>	
Preferred Deals in General Environments	231
<i>Yuan Deng, Sébastien Lahaie, Vahab Mirrokni</i>	
A Parameterized Perspective on Protecting Elections	238
<i>Palash Dey, Neeldhara Misra, Swaprava Nath, Garima Shakya</i>	
Spotting Collective Behaviour of Online Frauds in Customer Reviews	245
<i>Sarthika Dhawan, Siva Charan Reddy Gangireddy, Shiv Kumar, Tanmoy Chakraborty</i>	
Equilibrium Characterization for Data Acquisition Games	252
<i>Jinshuo Dong, Hadi Elzayn, Shahin Jabbari, Michael Kearns, Zachary Schutzman</i>	
Protecting Elections by Recounting Ballots	259
<i>Edith Elkind, Jiarui Gan, Svetlana Obraztsova, Zinovi Rabinovich, Alexandros A. Voudouris</i>	
Schelling Games on Graphs	266
<i>Edith Elkind, Jiarui Gan, Ayumi Igarashi, Warut Suksompong, Alexandros A. Voudouris</i>	
Reallocating Multiple Facilities on the Line	273
<i>Dimitris Fotakis, Loukas Kavouras, Panagiotis Kostopanagiotis, Philip Lazos, Stratis Skoulakis, Nikos Zarifis</i>	
Equitable Allocations of Indivisible Goods	280
<i>Rupert Freeman, Sujoy Sikdar, Rohit Vaish, Lirong Xia</i>	
Average-case Analysis of the Assignment Problem with Independent Preferences	287
<i>Yansong Gao, Jie Zhang</i>	
Improving Nash Social Welfare Approximations	294
<i>Jugal Garg, Peter McGlaughlin</i>	
On the Efficiency and Equilibria of Rich Ads	301
<i>MohammadAmin Ghiasi, MohammadTaghi Hajiaghayi, Sébastien Lahaie, Hadi Yami</i>	
Identifying vulnerabilities in trust and reputation systems	308
<i>Taha D. Günes, Long Tran-Thanh, Timothy J. Norman</i>	
An Asymptotically Optimal VCG Redistribution Mechanism for the Public Project Problem	315
<i>Mingyu Guo</i>	
On Succinct Encodings for the Tournament Fixing Problem	322
<i>Sushmita Gupta, Saket Saurabh, Ramanujan Sridharan, Meirav Zehavi</i>	
On Computational Tractability for Rational Verification	329
<i>Julian Gutierrez, Muhammad Najib, Giuseppe Perelli, Michael Wooldridge</i>	
Swarm Engineering Through Quantitative Measurement of Swarm Robotic Principles in a 10,000 Robot Swarm	336
<i>John Harwell, Maria Gini</i>	
Achieving a Fairer Future by Changing the Past	343
<i>Jiafan He, Ariel D. Procaccia, Alexandros Psomas, David Zeng</i>	
Compact Representation of Value Function in Partially Observable Stochastic Games	350
<i>Karel Horák, Branislav Bošanský, Christopher Kiekintveld, Charles Kamhoua</i>	
Explicitly Coordinated Policy Iteration	357
<i>Yujing Hu, Yingfeng Chen, Changjie Fan, Jianye Hao</i>	
Robustness against Agent Failure in Hedonic Games	364
<i>Ayumi Igarashi, Kazunori Ota, Yuko Sakurai, Makoto Yokoo</i>	
The Interplay of Emotions and Norms in Multiagent Systems	371
<i>Anup K. Kalia, Nirav Ajmeri, Kevin S. Chan, Jin-Hee Cho, Sibel Adali, Munindar P. Singh</i>	
An Ordinal Banzhaf Index for Social Ranking	378
<i>Hossein Khani, Stefano Moretti, Meltem Öztürk</i>	
Multigoal Committee Selection	385
<i>Maciej Kocot, Anna Kolonko, Edith Elkind, Piotr Faliszewski, Nimrod Talmon</i>	
Neural Networks for Predicting Human Interactions in Repeated Games	392
<i>Yoav Kolumbus, Gali Noti</i>	

Almost Envy-Freeness in Group Resource Allocation	400
<i>Maria Kyropoulou, Warut Suksompong, Alexandros A. Voudouris</i>	
A Quantitative Analysis of Multi-Winner Rules	407
<i>Martin Lackner, Piotr Skowron</i>	
Correlating Preferences and Attributes: Nearly Single-Crossing Profiles	414
<i>Foram Lakhani, Dominik Peters, Edith Elkind</i>	
Automated Negotiation with Gaussian Process-based Utility Models	421
<i>Haralambie Leahu, Michael Kaisers, Tim Baarslag</i>	
Temporal Information Design in Contests	428
<i>Priel Levy, David Sarne, Yonatan Aumann</i>	
Diffusion and Auction on Graphs	435
<i>Bin Li, Dong Hao, Dengji Zhao, Makoto Yokoo</i>	
Improved Heuristics for Multi-Agent Path Finding with Conflict-Based Search	442
<i>Jiaoyang Li, Ariel Felner, Eli Boyarski, Hang Ma, Sven Koenig</i>	
Integrating Decision Sharing with Prediction in Decentralized Planning for Multi-Agent Coordination under Uncertainty	450
<i>Minglong Li, Wenjing Yang, Zhongxuan Cai, Shaowu Yang, Ji Wang</i>	
Value Function Transfer for Deep Multi-Agent Reinforcement Learning Based on N-Step Returns	457
<i>Yong Liu, Yujing Hu, Yang Gao, Yingfeng Chen, Changjie Fan</i>	
Computing Approximate Equilibria in Sequential Adversarial Games by Exploitability Descent	464
<i>Edward Lockhart, Marc Lanctot, Julien Pérolat, Jean-Baptiste Lespiau, Dustin Morrill, Finbarr Timbers, Karl Tuyls</i>	
Computational Aspects of Equilibria in Discrete Preference Games	471
<i>Phani Raj Lolakapuri, Umang Bhaskar, Ramasuri Narayanam, Gyana R Parija, Pankaj S Dayama</i>	
Multi-Robot Planning Under Uncertain Travel Times and Safety Constraints	478
<i>Masoumeh Mansouri, Bruno Lacerda, Nick Hawes, Federico Pecora</i>	
Leadership in Congestion Games: Multiple User Classes and Non-Singleton Actions	485
<i>Alberto Marchesi, Matteo Castiglioni, Nicola Gatti</i>	
Graphical One-Sided Markets	492
<i>Sagar Massand, Sunil Simon</i>	
Reachability Games in Dynamic Epistemic Logic	499
<i>Bastien Maubert, Sophie Pinchinat, François Schwarzentruber</i>	
FaRM: Fair Reward Mechanism for Information Aggregation in Spontaneous Localized Settings	506
<i>Moin Hussain Moti, Dimitris Chatzopoulos, Pan Hui, Sujit Gujar</i>	
Learning Swarm Behaviors using Grammatical Evolution and Behavior Trees	513
<i>Aadesh Neupane, Michael Goodrich</i>	
A Probabilistic Logic for Resource-Bounded Multi-Agent Systems	521
<i>Hoang Nga Nguyen, Abdur Rakib</i>	
Imitative Attacker Deception in Stackelberg Security Games	528
<i>Thanh Nguyen, Haifeng Xu</i>	
Priority Inheritance with Backtracking for Iterative Multi-agent Path Finding	535
<i>Keisuke Okumura, Manao Machida, Xavier Défago, Yasumasa Tamura</i>	
Approval-Based Elections and Distortion of Voting Rules	543
<i>Grzegorz Pierczynski, Piotr Skowron</i>	
Ad Hoc Teamwork With Behavior Switching Agents	550
<i>Manish Ravula, Shani Alkoby, Peter Stone</i>	
Ridesharing with Driver Location Preferences	557
<i>Duncan Rheingans-Yoo, Scott Duke Kominers, Hongyao Ma, David C. Parkes</i>	
Multi-Population Congestion Games With Incomplete Information	565
<i>Charlotte Roman, Paolo Turrini</i>	
Sybil-Resilient Reality-Aware Social Choice	572
<i>Gal Shahaf, Ehud Shapiro, Nimrod Talmon</i>	
Preferences Single-Peaked on a Tree: Sampling and Tree Recognition	580
<i>Jakub Sliwinski, Edith Elkind</i>	
Model-Free Model Reconciliation	587
<i>Sarath Sreedharan, Alberto Olmo Hernandez, Aditya Prasad Mishra, Subbarao Kambhampati</i>	
Aggregating Incomplete Pairwise Preferences by Weight	595
<i>Zoi Terzopoulou, Ulle Endriss</i>	
A Regularized Opponent Model with Maximum Entropy Objective	602
<i>Zheng Tian, Ying Wen, Zhichen Gong, Faiz Punakkath, Shihao Zou, Jun Wang</i>	
Exploring the Task Cooperation in Multi-goal Visual Navigation	609
<i>Yuechen Wu, Zhenhuan Rao, Wei Zhang, Shijian Lu, Weizhi Lu, Zheng-Jun Zha</i>	

On Strategyproof Conference Peer Review	616
<i>Yichong Xu, Han Zhao, Xiaofei Shi, Nihar B. Shah</i>	
Towards Efficient Detection and Optimal Response against Sophisticated Opponents	623
<i>Tianpei Yang, Jianye Hao, Zhaopeng Meng, Chongjie Zhang, Yan Zheng, Ze Zheng</i>	
Large-Scale Home Energy Management Using Entropy-Based Collective Multiagent Deep Reinforcement Learning Framework	630
<i>Yaodong Yang, Jianye Hao, Yan Zheng, Chao Yu</i>	
Complexity of Manipulating and Controlling Approval-Based Multiwinner Voting	637
<i>Yongjie Yang</i>	
On the Tree Representations of Dichotomous Preferences	644
<i>Yongjie Yang</i>	
The Price of Governance: A Middle Ground Solution to Coordination in Organizational Control	651
<i>Chao Yu, Guozhen Tan</i>	
Decentralized Optimization with Edge Sampling	658
<i>Chi Zhang, Qianxiao Li, Peilin Zhao</i>	
Explore Truthful Incentives for Tasks with Heterogenous Levels of Difficulty in the Sharing Economy	665
<i>Pengzhan Zhou, Xin Wei, Cong Wang, Yuanyuan Yang</i>	

COMPUTER VISION

CoSegNet: Image Co-segmentation using a Conditional Siamese Convolutional Network	673
<i>Sayan Banerjee, Avik Hati, Subhasis Chaudhuri, Rajbabu Velmurugan</i>	
Multi-Margin based Decorrelation Learning for Heterogeneous Face Recognition	680
<i>Bing Cao, Nannan Wang, Xinbo Gao, Jie Li, Zhifeng Li</i>	
Generalized Zero-Shot Vehicle Detection in Remote Sensing Imagery via Coarse-to-Fine Framework	687
<i>Hong Chen, Yongtan Luo, Liujuan Cao, Baochang Zhang, Guodong Guo, Cheng Wang, Jonathan Li, Rongrong Ji</i>	
Structure-Aware Residual Pyramid Network for Monocular Depth Estimation	694
<i>Xiaotian Chen, Xuejin Chen, Zheng-Jun Zha</i>	
A Deep Bi-directional Attention Network for Human Motion Recovery	701
<i>Qiongjie Cui, Huaijiang Sun, Yupeng Li, Yue Kong</i>	
On Retrospecting Human Dynamics with Attention	708
<i>Minjing Dong, Chang Xu</i>	
Learning to Draw Text in Natural Images with Conditional Adversarial Networks	715
<i>Shancheng Fang, Hongtao Xie, Jianjun Chen, Jianlong Tan, Yongdong Zhang</i>	
Beyond Product Quantization: Deep Progressive Quantization for Image Retrieval	723
<i>Lianli Gao, Xiaosu Zhu, Jingkuan Song, Zhou Zhao, Heng Tao Shen</i>	

VOLUME 2

ANODE: Unconditionally Accurate Memory-Efficient Gradients for Neural ODEs	730
<i>Amir Gholaminejad, Kurt Keutzer, George Biros</i>	
Asynchronous Stochastic Frank-Wolfe Algorithms for Non-Convex Optimization	737
<i>Bin Gu, Wenhan Xian, Heng Huang</i>	
Dense Temporal Convolution Network for Sign Language Translation	744
<i>Dan Guo, Shuo Wang, Qi Tian, Meng Wang</i>	
Connectionist Temporal Modeling of Video and Language: a Joint Model for Translation and Sign Labeling	751
<i>Dan Guo, Shengeng Tang, Meng Wang</i>	
3DViewGraph: Learning Global Features for 3D Shapes from A Graph of Unordered Views with Attention	758
<i>Zhizhong Han, Xiyang Wang, Chi Man Vong, Yu-Shen Liu, Matthias Zwicker, C. L. Philip Chen</i>	
Parts4Feature: Learning 3D Global Features from Generally Semantic Parts in Multiple Views	766
<i>Zhizhong Han, Xinhai Liu, Yu-Shen Liu, Matthias Zwicker</i>	
MAT-Net: Medial Axis Transform Network for 3D Object Recognition	774
<i>Jianwei Hu, Bin Wang, Lihui Qian, Yiling Pan, Xiaohu Guo, Lingjie Liu, Wenping Wang</i>	
Dynamic Feature Fusion for Semantic Edge Detection	782
<i>Yuan Hu, Yunpeng Chen, Xiang Li, Jiashi Feng</i>	
Multi-Level Visual-Semantic Alignments with Relation-Wise Dual Attention Network for Image and Text Matching	789
<i>Zhibin Hu, Yongsheng Luo, Jiong Lin, Yan Yan, Jian Chen</i>	

Learning Unsupervised Visual Grounding Through Semantic Self-Supervision	796
<i>Syed Ashar Javed, Shreyas Saxena, Vineet Gandhi</i>	
Supervised Set-to-Set Hashing in Visual Recognition	803
<i>I-Hong Jhuo</i>	
Generative Image Inpainting with Submanifold Alignment	811
<i>Ang Li, Jianzhong Qi, Rui Zhang, Xingjun Ma, Kotagiri Ramamohanarao</i>	
Detecting Robust Co-Saliency with Recurrent Co-Attention Neural Network	818
<i>Bo Li, Zhengxing Sun, Lv Tang, Yunhan Sun, Jinlong Shi</i>	
Variation Generalized Feature Learning via Intra-view Variation Adaptation	826
<i>Jiawei Li, Mang Ye, Andy Jinhua Ma, Pong C Yuen</i>	
Pedestrian Attribute Recognition by Joint Visual-semantic Reasoning and Knowledge Distillation	833
<i>Qiaozhe Li, Xin Zhao, Ran He, Kaiqi Huang</i>	
Rethinking Loss Design for Large-scale 3D Shape Retrieval	840
<i>Zhaoqun Li, Cheng Xu, Biao Leng</i>	
Attribute-Aware Convolutional Neural Networks for Facial Beauty Prediction	847
<i>Luojun Lin, Lingyu Liang, Lianwen Jin, Weijie Chen</i>	
Rectified Binary Convolutional Networks for Enhancing the Performance of 1-bit DCNNs	854
<i>Chunlei Liu, Wenrui Ding, Xin Xia, Yuan Hu, Baochang Zhang, Jianzhuang Liu, Bohan Zhuang, Guodong Guo</i>	
Nuclei Segmentation via a Deep Panoptic Model with Semantic Feature Fusion	861
<i>Dongnan Liu, Donghao Zhang, Yang Song, Chaoyi Zhang, Fan Zhang, Lauren O'Donnell, Weidong Cai</i>	
Densely Connected Attention Flow for Visual Question Answering	869
<i>Fei Liu, Jing Liu, Zhiwei Fang, Richang Hong, Hanqing Lu</i>	
Unsupervised Learning of Scene Flow Estimation Fusing with Local Rigidity	876
<i>Liang Liu, Guangyao Zhai, Wenlong Ye, Yong Liu</i>	
Resolution-invariant Person Re-Identification	883
<i>Shunan Mao, Shiliang Zhang, Ming Yang</i>	
Low Shot Box Correction for Weakly Supervised Object Detection	890
<i>Tianxiang Pan, Bin Wang, Guiguang Ding, Jungong Han, Junhai Yong</i>	
DBDNet: Learning Bi-directional Dynamics for Early Action Prediction	897
<i>Guoliang Pang, Xionghui Wang, Jian-Fang Hu, Qing Zhang, Wei-Shi Zheng</i>	
Deep Light-field-driven Saliency Detection from a Single View	904
<i>Yongri Piao, Zhengkun Rong, Miao Zhang, Xiao Li, Huchuan Lu</i>	
Deep Recurrent Quantization for Generating Sequential Binary Codes	912
<i>Jingkuan Song, Xiaosu Zhu, Lianli Gao, Xin-Shun Xu, Wu Liu, Heng Tao Shen</i>	
Talking Face Generation by Conditional Recurrent Adversarial Network	919
<i>Yang Song, Jingwen Zhu, Dawei Li, Andy Wang, Hairong Qi</i>	
Hallucinating Optical Flow Features for Video Classification	926
<i>Yongyi Tang, Lin Ma, Lianqiang Zhou</i>	
Color-Sensitive Person Re-Identification	933
<i>Guan'an Wang, Yang Yang, Jian Cheng, Jinqiao Wang, Zengguang Hou</i>	
Convolutional Auto-encoding of Sentence Topics for Image Paragraph Generation	940
<i>Jing Wang, Yingwei Pan, Ting Yao, Jinhui Tang, Tao Mei</i>	
DSRN: A Deep Scale Relationship Network for Scene Text Detection	947
<i>Yuxin Wang, Hongtao Xie, Zilong Fu, Yongdong Zhang</i>	
Transferable Adversarial Attacks for Image and Video Object Detection	954
<i>Xingxing Wei, Siyuan Liang, Ning Chen, Xiaochun Cao</i>	
Video Interactive Captioning with Human Prompts	961
<i>Aming Wu, Yahong Han, Yi Yang</i>	
Mutually Reinforced Spatio-Temporal Convolutional Tube for Human Action Recognition	968
<i>Haoze Wu, Jiawei Liu, Zheng-Jun Zha, Zhenzhong Chen, Xiaoyan Sun</i>	
Densely Supervised Hierarchical Policy-Value Network for Image Paragraph Generation	975
<i>Siyang Wu, Zheng-Jun Zha, Zilei Wang, Houqiang Li, Feng Wu</i>	
Graph Convolutional Network Hashing for Cross-Modal Retrieval	982
<i>Ruiqing Xu, Chao Li, Junchi Yan, Cheng Deng, Xianglong Liu</i>	
MSR: Multi-Scale Shape Regression for Scene Text Detection	989
<i>Chuhui Xue, Shijian Lu, Wei Zhang</i>	
Dynamically Visual Disambiguation of Keyword-based Image Search	996
<i>Yazhou Yao, Zeren Sun, Fumin Shen, Li Liu, Limin Wang, Fan Zhu, Lizhong Ding, Gangshan Wu, Ling Shao</i>	
High Performance Gesture Recognition via Effective and Efficient Temporal Modeling	1003
<i>Yang Yi, Feng Ni, Yuexin Ma, Xinge Zhu, Yuankai Qi, Riming Qiu, Shijie Zhao, Feng Li, Yongtao Wang</i>	

Capturing Spatial and Temporal Patterns for Facial Landmark Tracking through Adversarial Learning	1010
<i>Shi Yin, Shangfei Wang, Guozhu Peng, Xiaoping Chen, Bowen Pan</i>	
Pose-preserving Cross Spectral Face Hallucination	1018
<i>Junchi Yu, Jie Cao, Yi Li, Xiaofei Jia, Ran He</i>	
Generative Visual Dialogue System via Weighted Likelihood Estimation	1025
<i>Heming Zhang, Shalini Ghosh, Larry Heck, Stephen Walsh, Junting Zhang, Jie Zhang, C.-C. Jay Kuo</i>	
Binarized Neural Networks for Resource-Efficient Hashing with Minimizing Quantization Loss	1032
<i>Feng Zheng, Cheng Deng, Heng Huang</i>	
LRDNN: Local-refining based Deep Neural Network for Person Re-Identification with Attribute Discerning	1041
<i>Qinqin Zhou, Bineng Zhong, Xiangyuan Lan, Gan Sun, Yulun Zhang, Mengran Gou</i>	
Face Photo-Sketch Synthesis via Knowledge Transfer	1048
<i>Mingrui Zhu, Nannan Wang, Xinbo Gao, Jie Li, Zhifeng Li</i>	

CONSTRAINTS AND SAT

Athantor: High-Level Local Search Over Abstract Constraint Specifications in Essence	1056
<i>Saad Attieh, Nguyen Dang, Christopher Jefferson, Ian Miguel, Peter Nightingale</i>	
Constraint Programming for Mining Borders of Frequent Itemsets	1064
<i>Mohamed-Bachir Belaid, Christian Bessiere, Nadjib Lazaar</i>	
How to Tame Your Anticipatory Algorithm	1071
<i>Allegra De Filippo, Michele Lombardi, Michela Milano</i>	
Predict+Optimise with Ranking Objectives: Exhaustively Learning Linear Functions	1078
<i>Emir Demirovic, Peter J. Stuckey, James Bailey, Jeffrey Chan, Christopher Leckie, Kotagiri Ramamohanarao, Tias Guns</i>	
Privacy-Preserving Obfuscation of Critical Infrastructure Networks	1086
<i>Ferdinando Fioretto, Terrence W.K. Mak, Pascal Van Hentenryck</i>	
Solving the Satisfiability Problem of Modal Logic S5 Guided by Graph Coloring	1093
<i>Pei Huang, Minghao Liu, Ping Wang, Wenhui Zhang, Feifei Ma, Jian Zhang</i>	
DoubleLex Revisited and Beyond	1101
<i>Xuming Huang, Jimmy Lee</i>	
Model-Based Diagnosis with Multiple Observations	1108
<i>Alexey Ignatiev, Antonio Morgado, Georg Weissenbacher, Joao Marques-Silva</i>	
Enumerating Potential Maximal Cliques via SAT and ASP	1116
<i>Tuukka Korhonen, Jeremias Berg, Matti Järvisalo</i>	
Entropy-Penalized Semidefinite Programming	1123
<i>Mikhail Krechetov, Jakub Marecek, Yury Maximov, Martin Takac</i>	
Acquiring Integer Programs from Data	1130
<i>Mohit Kumar, Stefano Teso, Luc De Raedt</i>	
Stochastic Constraint Propagation for Mining Probabilistic Networks	1137
<i>Anna Louise D. Latour, Behrouz Babaki, Siegfried Nijssen</i>	
Optimizing Constraint Solving via Dynamic Programming	1146
<i>Shu Lin, Na Meng, Wenxin Li</i>	
Constraint-Based Scheduling with Complex Setup Operations: An Iterative Two-Layer Approach	1155
<i>Adriana Pacheco, Cédric Pralet, Stéphanie Roussel</i>	
Phase Transition Behavior of Cardinality and XOR Constraints	1162
<i>Yash Pote, Saurabh Joshi, Kuldeep S. Meel</i>	
GANAK: A Scalable Probabilistic Exact Model Counter	1169
<i>Shubham Sharma, Subhajit Roy, Mate Soos, Kuldeep S. Meel</i>	
Unifying Search-based and Compilation-based Approaches to Multi-agent Path Finding through Satisfiability Modulo Theories	1177
<i>Pavel Surynek</i>	
Integrating Pseudo-Boolean Constraint Reasoning in Multi-Objective Evolutionary Algorithms	1184
<i>Miguel Terra-Neves, Inês Lynce, Vasco Manquinho</i>	
Resolution and Domination: An Improved Exact MaxSAT Algorithm	1191
<i>Chao Xu, Wenjun Li, Yongjie Yang, Jianer Chen, Jianxin Wang</i>	

HEURISTIC SEARCH AND GAME PLAYING

Path Planning with CPD Heuristics	1199
<i>Massimo Bono, Alfonso E. Gerevini, Daniel D. Harabor, Peter J. Stuckey</i>	
A*+IDA*: A Simple Hybrid Search Algorithm	1206
<i>Zhaoxing Bu, Richard E. Korf</i>	
Deanonymizing Social Networks Using Structural Information	1213
<i>Ioannis Caragiannis, Evanthia Tsitsoka</i>	
Conditions for Avoiding Node Re-expansions in Bounded Suboptimal Search	1220
<i>Jingwei Chen, Nathan R. Sturtevant</i>	
An Efficient Evolutionary Algorithm for Minimum Cost Submodular Cover	1227
<i>Victoria G. Crawford</i>	
An Evolution Strategy with Progressive Episode Lengths for Playing Games	1234
<i>Lior Fuks, Noor Awad, Frank Hutter, Marius Lindauer</i>	
Regarding Jump Point Search and Subgoal Graphs	1241
<i>Daniel D. Harabor, Tansel Uras, Peter J. Stuckey, Sven Koenig</i>	
Iterative Budgeted Exponential Search	1249
<i>Malte Helmert, Tor Lattimore, Levi H. S. Lelis, Laurent Orseau, Nathan R. Sturtevant</i>	
Direction-Optimizing Breadth-First Search with External Memory Storage	1258
<i>Shuli Hu, Nathan R. Sturtevant</i>	
DeltaDou: Expert-level Doudizhu AI through Self-play	1265
<i>Qiqi Jiang, Kuangzheng Li, Boyao Du, Hao Chen, Hai Fang</i>	
Graph Mining Meets Crowdsourcing: Extracting Experts for Answer Aggregation	1272
<i>Yasushi Kawase, Yuko Kuroki, Atsushi Miyauchi</i>	
Depth-First Memory-Limited AND/OR Search and Unsolvability in Cyclic Search Spaces	1280
<i>Akihiro Kishimoto, Adi Botea, Radu Marinescu</i>	
Branch-and-Cut-and-Price for Multi-Agent Pathfinding	1289
<i>Edward Lam, Pierre Le Bodic, Daniel D. Harabor, Peter J. Stuckey</i>	
Local Search with Efficient Automatic Configuration for Minimum Vertex Cover	1297
<i>Chuan Luo, Holger H. Hoos, Shaowei Cai, Qingwei Lin, Hongyu Zhang, Dongmei Zhang</i>	
Learning Deep Decentralized Policy Network by Collective Rewards for Real-Time Combat Game	1305
<i>Peixi Peng, Junliang Xing, Lili Cao, Lisen Mu, Chang Huang</i>	
Heuristic Search for Homology Localization Problem and Its Application in Cardiac Trabeculae Reconstruction	1312
<i>Xudong Zhang, Pengxiang Wu, Changhe Yuan, Yusu Wang, Dimitris Metaxas, Chao Chen</i>	
Non-smooth Optimization over Stiefel Manifolds with Applications to Dimensionality Reduction and Graph Clustering	1319
<i>Fariba Zohrizadeh, Mohsen Kheirandishfard, Farhad Kamangar, Ramtin Madani</i>	

HUMANS AND AI

Explaining Reinforcement Learning to Mere Mortals: An Empirical Study	1328
<i>Andrew Anderson, Jonathan Dodge, Amrita Sadarangani, Zoe Juozapaitis, Evan Newman, Jed Irvine, Souti Chattopadhyay, Alan Fern, Margaret Burnett</i>	
Balancing Explicability and Explanations in Human-Aware Planning	1335
<i>Tathagata Chakraborti, Sarath Sreedharan, Subbarao Kambhampati</i>	
Multi-agent Attentional Activity Recognition	1344
<i>Kaixuan Chen, Lina Yao, Dalin Zhang, Bin Guo, Zhiwen Yu</i>	
Deep Adversarial Social Recommendation	1351
<i>Wenqi Fan, Tyler Derr, Yao Ma, Jianping Wang, Jiliang Tang, Qing Li</i>	
A Semantics-based Model for Predicting Children's Vocabulary	1358
<i>Ishaan Grover, Hae Won Park, Cynthia Breazeal</i>	
STCA: Spatio-Temporal Credit Assignment with Delayed Feedback in Deep Spiking Neural Networks	1366
<i>Pengjie Gu, Rong Xiao, Gang Pan, Huajin Tang</i>	
Dynamic Item Block and Prediction Enhancing Block for Sequential Recommendation	1373
<i>Guibing Guo, Shichang Ouyang, Xiaodong He, Fajie Yuan, Xiaohua Liu</i>	
Discrete Trust-aware Matrix Factorization for Fast Recommendation	1380
<i>Guibing Guo, Enneng Yang, Li Shen, Xiaochun Yang, Xiaodong He</i>	
Decoding EEG by Visual-guided Deep Neural Networks	1387
<i>Zhicheng Jiao, Haoxuan You, Fan Yang, Xin Li, Han Zhang, Dinggang Shen</i>	

MISC: Mixed Strategies Crowdsourcing	1394
<i>Ching Yun Ko, Rui Lin, Shu Li, Ngai Wong</i>	
Exploring Computational User Models for Agent Policy Summarization	1401
<i>Isaac Lage, Daphna Lifschitz, Finale Doshi-Velez, Ofra Amir</i>	
Minimizing Time-to-Rank: A Learning and Recommendation Approach	1408
<i>Haoming Li, Sujoy Sikdar, Rohit Vaish, Junming Wang, Lirong Xia, Chaonan Ye</i>	
DeepFlow: Detecting Optimal User Experience From Physiological Data Using Deep Neural Networks	1415
<i>Marco Maier, Daniel Elsner, Chadly Marouane, Meike Zehnle, Christoph Fuchs</i>	
Why Can't You Do That HAL? Explaining Unsolvability of Planning Tasks	1422
<i>Sarath Sreedharan, Siddharth Srivastava, David Smith, Subbarao Kambhampati</i>	
Personalized Multimedia Item and Key Frame Recommendation	1431
<i>Le Wu, Lei Chen, Yonghui Yang, Richang Hong, Yong Ge, Xing Xie, Meng Wang</i>	
Counterfactual Fairness: Unidentification, Bound and Algorithm	1438
<i>Yongkai Wu, Lu Zhang, Xintao Wu</i>	
Fast and Accurate Classification with a Multi-Spike Learning Algorithm for Spiking Neurons	1445
<i>Rong Xiao, Qiang Yu, Rui Yan, Huajin Tang</i>	
Achieving Causal Fairness through Generative Adversarial Networks	1452
<i>Depeng Xu, Yongkai Wu, Shuhan Yuan, Lu Zhang, Xintao Wu</i>	
DeepAPF: Deep Attentive Probabilistic Factorization for Multi-site Video Recommendation	1459
<i>Huan Yan, Xiangning Chen, Chen Gao, Yong Li, Depeng Jin</i>	

VOLUME 3

An Input-aware Factorization Machine for Sparse Prediction	1466
<i>Yantao Yu, Zhen Wang, Bo Yuan</i>	
Multiple Noisy Label Distribution Propagation for Crowdsourcing	1473
<i>Hao Zhang, Liangxiao Jiang, Wenqiang Xu</i>	
FAHT: An Adaptive Fairness-aware Decision Tree Classifier	1480
<i>Wenbin Zhang, Eirini Ntoutsi</i>	

KNOWLEDGE REPRESENTATION AND REASONING

ASP-based Discovery of Semi-Markovian Causal Models under Weaker Assumptions	1488
<i>Zhalama, Jiji Zhang, Frederick Eberhardt, Wolfgang Mayer, Mark Junjie Li</i>	
On the Integration of CP-nets in ASPRIN	1495
<i>Mario Alviano, Javier Romero, Torsten Schaub</i>	
Compilation of Logical Arguments	1502
<i>Leila Amgoud, Dragan Doder</i>	
Observations on Darwiche and Pearl's Approach for Iterated Belief Revision	1509
<i>Theofanis Aravanis, Pavlos Peppas, Mary-Anne Williams</i>	
Do You Need Infinite Time?	1516
<i>Alessandro Artale, Andrea Mazzullo, Ana Ozaki</i>	
Stratified Evidence Logics	1523
<i>Philippe Balbiani, David Fernández-Duque, Andreas Herzig, Emiliano Lorini</i>	
Worst-Case Optimal Querying of Very Expressive Description Logics with Path Expressions and Succinct Counting	1530
<i>Bartosz Bednarczyk, Sebastian Rudolph</i>	
Comparing Options with Argument Schemes Powered by Cancellation	1537
<i>Khaled Belahcene, Christophe Labreuche, Nicolas Maudet, Vincent Mousseau, Wassila Ouerdane</i>	
Possibilistic Games with Incomplete Information	1544
<i>Nahla Ben Amor, Helene Fargier, Régis Sabbadin, Meriem Trabelsi</i>	
Reasoning about Disclosure in Data Integration in the Presence of Source Constraints	1551
<i>Michael Benedikt, Pierre Bourhis, Louis Jachiet, Michaël Thomazo</i>	
Mixed-World Reasoning with Existential Rules under Active-Domain Semantics	1558
<i>Meghyn Bienvenu, Pierre Bourhis</i>	
Guarantees for Sound Abstractions for Generalized Planning	1566
<i>Blai Bonet, Raquel Fuentetaja, Yolanda E-Martín, Daniel Borrajo</i>	
Ontology Approximation in Horn Description Logics	1574
<i>Anneke Böcher, Carsten Lutz, Frank Wolter</i>	

Oblivious and Semi-Oblivious Boundedness for Existential Rules	1581
<i>Pierre Bourhis, Michel Leclère, Marie-Laure Mugnier, Sophie Tison, Federico Ulliana, Lily Gallois</i>	
Reasoning about Quality and Fuzziness of Strategic Behaviours	1588
<i>Patricia Bouyer, Orna Kupferman, Nicolas Markey, Bastien Maubert, Aniello Murano, Giuseppe Perelli</i>	
The Complexity of Model Checking Knowledge and Time	1595
<i>Laura Bozzelli, Bastien Maubert, Aniello Murano</i>	
Planning for LTL_f / LDL_f Goals in Non-Markovian Fully Observable Nondeterministic Domains	1602
<i>Ronen I. Brafman, Giuseppe De Giacomo</i>	
Causal Discovery with Cascade Nonlinear Additive Noise Model	1609
<i>Ruichu Cai, Jie Qiao, Kun Zhang, Zhenjie Zhang, Zhifeng Hao</i>	
Enriching Ontology-based Data Access with Provenance	1616
<i>Diego Calvanese, Davide Lanti, Ana Ozaki, Rafael Penaloza, Guohui Xiao</i>	
Chasing Sets: How to Use Existential Rules for Expressive Reasoning	1624
<i>David Carral, Irina Dragoste, Markus Krötzsch, Christian Lewe</i>	
Simple Conditionals with Constrained Right Weakening	1632
<i>Giovanni Casini, Thomas Meyer, Ivan Varzinczak</i>	
Explanations for Query Answers under Existential Rules	1639
<i>Ismail Ilkan Ceylan, Thomas Lukasiewicz, Enrico Malizia, Andrius Vaicenavicius</i>	
Semantic Characterization of Data Services through Ontologies	1647
<i>Gianluca Cima, Maurizio Lenzerini, Antonella Poggi</i>	
Measuring the Likelihood of Numerical Constraints	1654
<i>Marco Console, Matthias Hofer, Leonid Libkin</i>	
From Statistical Transportability to Estimating the Effect of Stochastic Interventions	1661
<i>Juan D. Correa, Elias Bareinboim</i>	
Answer Set Programming for Judgment Aggregation	1668
<i>Ronald de Haan, Marija Slavkovic</i>	
An ASP Approach to Generate Minimal Countermodels in Intuitionistic Propositional Logic	1675
<i>Camillo Fiorentini</i>	
Learning Description Logic Concepts: When can Positive and Negative Examples be Separated?	1682
<i>Maurice Funk, Jean Christoph Jung, Carsten Lutz, Hadrien Pulcini, Frank Wolter</i>	
Aggressive Driving Saves More Time? Multi-task Learning for Customized Travel Time Estimation	1689
<i>Ruipeng Gao, Xiaoyu Guo, Fuyong Sun, Lin Dai, Jiayan Zhu, Chenxi Hu, Haibo Li</i>	
Approximating Integer Solution Counting via Space Quantification for Linear Constraints	1697
<i>Cunjing Ge, Feifei Ma, Xutong Ma, Fan Zhang, Pei Huang, Jian Zhang</i>	
Best Answers over Incomplete Data : Complexity and First-Order Rewritings	1704
<i>Amélie Gheerbrant, Cristina Sirangelo</i>	
On Division Versus Saturation in Pseudo-Boolean Solving	1711
<i>Stephan Gocht, Jakob Nordström, Amir Yehudayoff</i>	
On Finite and Unrestricted Query Entailment beyond SQ with Number Restrictions on Transitive Roles	1719
<i>Tomasz Gogacz, Víctor Gutiérrez-Basulto, Yuzmín Ibáñez-García, Jean Christoph Jung, Filip Murlak</i>	
Belief Revision Operators with Varying Attitudes Towards Initial Beliefs	1726
<i>Adrian Haret, Stefan Woltran</i>	
Some Things are Easier for the Dumb and the Bright Ones (Beware the Average!)	1734
<i>Wojciech Jamroga, Michal Knapik</i>	
Converging on Common Knowledge	1741
<i>Dominik Klein, Rasmus Kræmmer Rendsvig</i>	
Rational Inference Relations from Maximal Consistent Subsets Selection	1749
<i>Sébastien Konieczny, Pierre Marquis, Srdjan Vesic</i>	
How to Handle Missing Values in Multi-Criteria Decision Aiding?	1756
<i>Christophe Labreuche, Sébastien Destercke</i>	
A Tractable, Expressive, and Eventually Complete First-Order Logic of Limited Belief	1764
<i>Gerhard Lakemeyer, Hector J. Levesque</i>	
Travel Time Estimation without Road Networks: An Urban Morphological Layout Representation Approach	1772
<i>Wuwei Lan, Yanyan Xu, Bin Zhao</i>	
Augmenting Transfer Learning with Semantic Reasoning	1779
<i>Freddy Lécué, Jiaoyan Chen, Jeff Z. Pan, Huajun Chen</i>	
Revisiting Controlled Query Evaluation in Description Logics	1786
<i>Domenico Lembo, Riccardo Rosati, Domenico Fabio Savo</i>	
Unit Selection Based on Counterfactual Logic	1793
<i>Ang Li, Judea Pearl</i>	

Story Ending Prediction by Transferable BERT	1800
<i>Zhongyang Li, Xiao Ding, Ting Liu</i>	
Geo-ALM: POI Recommendation by Fusing Geographical Information and Adversarial Learning Mechanism	1807
<i>Wei Liu, Zhi-Jie Wang, Bin Yao, Jian Yin</i>	
Automatic Verification of FSA Strategies via Counterexample-Guided Local Search for Invariants	1814
<i>Kailun Luo, Yongmei Liu</i>	
BiOWA for Preference Aggregation with Bipolar Scales: Application to Fair Optimization in Combinatorial Domains	1822
<i>Hugo Martin, Patrice Perny</i>	
Satisfaction and Implication of Integrity Constraints in Ontology-based Data Access	1829
<i>Charalampos Nikolaou, Bernardo Cuenca Grau, Egor V. Kostylev, Mark Kaminski, Ian Horrocks</i>	
Monitoring of a Dynamic System Based on Autoencoders	1836
<i>Aomar Osmani, Massinissa Hamidi, Salah Bouhouche</i>	
Boosting for Comparison-Based Learning	1844
<i>Michael Perrot, Ulrike von Luxburg</i>	
Data Complexity and Rewritability of Ontology-Mediated Queries in Metric Temporal Logic under the Event-Based Semantics	1851
<i>Vladislav Ryzhikov, Przemyslaw Andrzej Walega, Michael Zakharyashev</i>	
Belief Update without Compactness in Non-finitary Languages	1858
<i>Jandson S Ribeiro, Abhaya Nayak, Renata Wassermann</i>	
What Has Been Said? Identifying the Change Formula in a Belief Revision Scenario	1865
<i>Nicolas Schwind, Katsumi Inoue, Sébastien Konieczny, Jean-Marie Lagniez, Pierre Marquis</i>	
Estimating Causal Effects of Tone in Online Debates	1872
<i>Dhanya Sridhar, Lise Getoor</i>	
Out of Sight But Not Out of Mind: An Answer Set Programming Based Online Abduction Framework for Visual Sensemaking in Autonomous Driving	1879
<i>Jakob Suchan, Mehul Bhatt, Srikrishna Varadarajan</i>	
DatalogMTL: Computational Complexity and Expressive Power	1886
<i>Przemyslaw A. Walega, Bernardo Cuenca Grau, Mark Kaminski, Egor V. Kostylev</i>	
Cross-City Transfer Learning for Deep Spatio-Temporal Prediction	1893
<i>Leye Wang, Xu Geng, Xiaojuan Ma, Feng Liu, Qiang Yang</i>	
A Modal Characterization Theorem for a Probabilistic Fuzzy Description Logic	1900
<i>Paul Wild, Lutz Schröder, Dirk Pattinson, Barbara König</i>	
Graph WaveNet for Deep Spatial-Temporal Graph Modeling	1907
<i>Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Chengqi Zhang</i>	
Profit-driven Task Assignment in Spatial Crowdsourcing	1914
<i>Jinfu Xia, Yan Zhao, Guanfeng Liu, Jiajie Xu, Min Zhang, Kai Zheng</i>	
Boosting Causal Embeddings via Potential Verb-Mediated Causal Patterns	1921
<i>Zhipeng Xie, Feiteng Mu</i>	
Graph Convolutional Networks using Heat Kernel for Semi-supervised Learning	1928
<i>Bingbing Xu, Huawei Shen, Qi Cao, Keting Cen, Xueqi Cheng</i>	
TransMS: Knowledge Graph Embedding for Complex Relations by Multidirectional Semantics	1935
<i>Shihui Yang, Jidong Tian, Honglun Zhang, Junchi Yan, Hao He, Yaohui Jin</i>	
Neighborhood-Aware Attentional Representation for Multilingual Knowledge Graphs	1943
<i>Qiannan Zhu, Xiaofei Zhou, Jia Wu, Jianlong Tan, Li Guo</i>	

MACHINE LEARNING A-L

The Expected-Length Model of Options	1951
<i>David Abel, John Winder, Marie desJardins, Michael Littman</i>	
Human-in-the-loop Active Covariance Learning for Improving Prediction in Small Data Sets	1959
<i>Homayun Afrabandpey, Tomi Peltola, Samuel Kaski</i>	
Inter-node Hellinger Distance based Decision Tree	1967
<i>Pritom Saha Akash, Md. Eusha Kadir, Amin Ahsan Ali, Mohammad Shoyaib</i>	
Unobserved Is Not Equal to Non-existent: Using Gaussian Processes to Infer Immediate Rewards Across Contexts	1974
<i>Hamoon Azizsoltani, Yeo Jin Kim, Markel Sanz Ausin, Tiffany Barnes, Min Chi</i>	
STG2Seq: Spatial-Temporal Graph to Sequence Model for Multi-step Passenger Demand Forecasting	1981
<i>Lei Bai, Lina Yao, Salil S. Kanhere, Xianzhi Wang, Quan Z. Sheng</i>	

Unsupervised Inductive Graph-Level Representation Learning via Graph-Graph Proximity	1988
<i>Yunsheng Bai, Hao Ding, Yang Qiao, Agustin Marinovic, Ken Gu, Ting Chen, Yizhou Sun, Wei Wang</i>	
Conditional GAN with Discriminative Filter Generation for Text-to-Video Synthesis	1995
<i>Yogesh Balaji, Martin Renqiang Min, Bing Bai, Rama Chellappa, Hans Peter Graf</i>	
An Actor-Critic-Attention Mechanism for Deep Reinforcement Learning in Multi-view Environments	2002
<i>Elaheh Barati, Xuewen Chen</i>	
Motion Invariance in Visual Environments	2009
<i>Alessandro Betti, Marco Gori, Stefano Melacci</i>	
Optimal Exploitation of Clustering and History Information in Multi-armed Bandit	2016
<i>Djallel Bouneffouf, Srinivasan Parthasarathy, Horst Samulowitz, Martin Wistuba</i>	
Incremental Elicitation of Rank-Dependent Aggregation Functions based on Bayesian Linear Regression	2023
<i>Nadjet Bourdache, Patrice Perny, Olivier Spanjaard</i>	
A Gradient-Based Split Criterion for Highly Accurate and Transparent Model Trees	2030
<i>Klaus Broelemann, Gjergji Kasneci</i>	
Matrix Completion in the Unit Hypercube via Structured Matrix Factorization	2038
<i>Emanuele Bugliarello, Swayambhoo Jain, Vineeth Rakesh</i>	
Active Learning within Constrained Environments through Imitation of an Expert Questioner	2045
<i>Kalesha Bullard, Yannick Schroecker, Sonia Chernova</i>	
Multi-View Active Learning for Video Recommendation	2053
<i>Jia-Jia Cai, Jun Tang, Qing-Guo Chen, Yao Hu, Xiaobo Wang, Sheng-Jun Huang</i>	
Learning Disentangled Semantic Representation for Domain Adaptation	2060
<i>Ruichu Cai, Zijian Li, Pengfei Wei, Jie Qiao, Kun Zhang, Zhifeng Hao</i>	
Tree Sampling Divergence: An Information-Theoretic Metric for Hierarchical Graph Clustering	2067
<i>Bertrand Charpentier, Thomas Bonald</i>	
Fake Tables: Using GANs to Generate Functional Dependency Preserving Tables with Bounded Real Data	2074
<i>Haipeng Chen, Sushil Jajodia, Jing Liu, Noseong Park, Vadim Sokolov, V. S. Subrahmanian</i>	
Theoretical Investigation of Generalization Bound for Residual Networks	2081
<i>Hao Chen, Zhanfeng Mo, Zhouwang Yang, Xiao Wang</i>	
Learning Semantic Annotations for Tabular Data	2088
<i>Jiaoyan Chen, Ernesto Jimenez-Ruiz, Ian Horrocks, Charles Sutton</i>	
Matching User with Item Set: Collaborative Bundle Recommendation with Deep Attention Network	2095
<i>Liang Chen, Yang Liu, Xiangnan He, Lianli Gao, Zibin Zheng</i>	
Cooperative Pruning in Cross-Domain Deep Neural Network Compression	2102
<i>Shangyu Chen, Wenya Wang, Sinno Jialin Pan</i>	
Extensible Cross-Modal Hashing	2109
<i>Tian-yi Chen, Lan Zhang, Shi-cong Zhang, Zi-long Li, Bai-chuan Huang</i>	
Semi-supervised User Profiling with Heterogeneous Graph Attention Networks	2116
<i>Weijian Chen, Yulong Gu, Zhaochun Ren, Xiangnan He, Hongtao Xie, Tong Guo, Dawei Yin, Yongdong Zhang</i>	
ActiveHNE: Active Heterogeneous Network Embedding	2123
<i>Xia Chen, Guoxian Yu, Jun Wang, Carlotta Domeniconi, Zhao Li, Xiangliang Zhang</i>	
A Restart-based Rank-1 Evolution Strategy for Reinforcement Learning	2130
<i>Zefeng Chen, Yuren Zhou, Xiao-yu He, Siyu Jiang</i>	
Co-Attentive Multi-Task Learning for Explainable Recommendation	2137
<i>Zhongxia Chen, Xiting Wang, Xing Xie, Tong Wu, Guoqing Bu, Yining Wang, Enhong Chen</i>	
Variational Graph Embedding and Clustering with Laplacian Eigenmaps	2144
<i>Zitai Chen, Chuan Chen, Zong Zhang, Zibin Zheng, Qingsong Zou</i>	
Deep Active Learning for Anchor User Prediction	2151
<i>Anfeng Cheng, Chuan Zhou, Hong Yang, Jia Wu, Lei Li, Jianlong Tan, Li Guo</i>	
Success Prediction on Crowdfunding with Multimodal Deep Learning	2158
<i>Chaoran Cheng, Fei Tan, Xiurui Hou, Zhi Wei</i>	
Approximate Optimal Transport for Continuous Densities with Copulas	2165
<i>Jinjin Chi, Jihong Ouyang, Ximing Li, Yang Wang, Meng Wang</i>	
Ornstein Auto-Encoders	2172
<i>Youngwon Choi, Joong-Ho Won</i>	
A Strongly Asymptotically Optimal Agent in General Environments	2179
<i>Michael K. Cohen, Elliot Catt, Marcus Hutter</i>	
Extrapolating Paths with Graph Neural Networks	2187
<i>Jean-Baptiste Cordonnier, Andreas Loukas</i>	

VOLUME 4

Recommending Links to Maximize the Influence in Social Networks	2195
<i>Federico Corò, Gianlorenzo D'Angelo, Yllka Velaj</i>	
Three-Player Wasserstein GAN via Amortised Duality	2202
<i>Nhan Dam, Quan Hoang, Trung Le, Tu Dinh Nguyen, Hung Bui, Dinh Phung</i>	
Learn Smart with Less: Building Better Online Decision Trees with Fewer Training Examples	2209
<i>Ariyam Das, Jin Wang, Sahil M. Gandhi, Jae Lee, Wei Wang, Carlo Zaniolo</i>	
IRC-GAN: Introspective Recurrent Convolutional GAN for Text-to-video Generation	2216
<i>Kangle Deng, Tianyi Fei, Xin Huang, Yuxin Peng</i>	
Marginal Posterior Sampling for Slate Bandits	2223
<i>Maria Dimakopoulou, Nikos Vlassis, Tony Jebara</i>	
Reinforced Negative Sampling for Recommendation with Exposure Data	2230
<i>Jingtao Ding, Yuhan Quan, Xiangnan He, Yong Li, Depeng Jin</i>	
Group Reconstruction and Max-Pooling Residual Capsule Network	2237
<i>Xinpeng Ding, Nannan Wang, Xinbo Gao, Jie Li, Xiaoyu Wang</i>	
Crafting Efficient Neural Graph of Large Entropy	2244
<i>Minjing Dong, Hanting Chen, Yunhe Wang, Chang Xu</i>	
Joint Link Prediction and Network Alignment via Cross-graph Embedding	2251
<i>Xingbo Du, Junchi Yan, Hongyuan Zha</i>	
Fast Algorithm for K-Truss Discovery on Public-Private Graphs	2258
<i>Soroush Ebadian, Xin Huang</i>	
Mindful Active Learning	2265
<i>Zhila Esna Ashari, Hassan Ghasemzadeh</i>	
iDev: Enhancing Social Coding Security by Cross-platform User Identification Between GitHub and Stack Overflow	2272
<i>Yujie Fan, Yiming Zhang, Shifu Hou, Lingwei Chen, Yanfang Ye, Chuan Shi, Liang Zhao, Shouhuai Xu</i>	
Hybrid Actor-Critic Reinforcement Learning in Parameterized Action Space	2279
<i>Zhou Fan, Rui Su, Weinan Zhang, Yong Yu</i>	
GSTNet: Global Spatial-Temporal Network for Traffic Flow Prediction	2286
<i>Shen Fang, Qi Zhang, Gaofeng Meng, Shiming Xiang, Chunhong Pan</i>	
Partial Label Learning by Semantic Difference Maximization	2294
<i>Lei Feng, Bo An</i>	
Deep Session Interest Network for Click-Through Rate Prediction	2301
<i>Yufei Feng, Fuyu Lv, Weichen Shen, Menghan Wang, Fei Sun, Yu Zhu, Keping Yang</i>	
Curriculum Learning for Cumulative Return Maximization	2308
<i>Francesco Foglino, Christiano Coletto Christakou, Ricardo Luna Gutierrez, Matteo Leonetti</i>	
Advocacy Learning: Learning through Competition and Class-Conditional Representations	2315
<i>Ian Fox, Jenna Wiens</i>	
Neurons Merging Layer: Towards Progressive Redundancy Reduction for Deep Supervised Hashing	2322
<i>Chaoyou Fu, Liangchen Song, Xiang Wu, Guoli Wang, Ran He</i>	
Deep Multi-Agent Reinforcement Learning with Discrete-Continuous Hybrid Action Spaces	2329
<i>Haotian Fu, Hongyao Tang, Jianye Hao, Zihan Lei, Yingfeng Chen, Changjie Fan</i>	
Automatic Successive Reinforcement Learning with Multiple Auxiliary Rewards	2336
<i>Zhao-Yang Fu, De-Chuan Zhan, Xin-Chun Li, Yi-Xing Lu</i>	
RecoNet: An Interpretable Neural Architecture for Recommender Systems	2343
<i>Francesco Fusco, Michalis Vlachos, Vasileios Vasileiadis, Kathrin Wardatzky, Johannes Schneider</i>	
Reward Learning for Efficient Reinforcement Learning in Extractive Document Summarisation	2350
<i>Yang Gao, Christian M. Meyer, Mohsen Mesgar, Iryna Gurevych</i>	
Fully Distributed Bayesian Optimization with Stochastic Policies	2357
<i>Javier Garcia-Barcos, Ruben Martinez-Cantin</i>	
Scalable Semi-Supervised SVM via Triply Stochastic Gradients	2364
<i>Xiang Geng, Bin Gu, Xiang Li, Wanli Shi, Guansheng Zheng, Heng Huang</i>	
Perception-Aware Point-Based Value Iteration for Partially Observable Markov Decision Processes	2371
<i>Mahsa Ghasemi, Ufuk Topcu</i>	
Efficient Regularization Parameter Selection for Latent Variable Graphical Models via Bi-Level Optimization	2378
<i>Joachim Giesen, Frank Nussbaum, Christopher Schneider</i>	
Using Natural Language for Reward Shaping in Reinforcement Learning	2385
<i>Prasoon Goyal, Scott Niekum, Raymond J. Mooney</i>	
Sketched Iterative Algorithms for Structured Generalized Linear Models	2392
<i>Qilong Gu, Arindam Banerjee</i>	

SPINE: Structural Identity Preserved Inductive Network Embedding	2399
<i>Junliang Guo, Linli Xu, Jingchang Liu</i>	
Discriminative Sample Generation for Deep Imbalanced Learning	2406
<i>Ting Guo, Xingquan Zhu, Yang Wang, Fang Chen</i>	
Affine Equivariant Autoencoder	2413
<i>Xifeng Guo, En Zhu, Xinwang Liu, Jianping Yin</i>	
AdaLinUCB: Opportunistic Learning for Contextual Bandits	2420
<i>Xueying Guo, Xiaoxiao Wang, Xin Liu</i>	
Zero-shot Learning with Many Classes by High-rank Deep Embedding Networks	2428
<i>Yuchen Guo, Guiguang Ding, Jungong Han, Hang Shao, Xin Lou, Qionghai Dai</i>	
Landmark Selection for Zero-shot Learning	2435
<i>Yuchen Guo, Guiguang Ding, Jungong Han, Chenggang Yan, Jiyong Zhang, Qionghai Dai</i>	
MineRL: A Large-Scale Dataset of Minecraft Demonstrations	2442
<i>William H. Guss, Brandon Houghton, Nicholay Topin, Phillip Wang, Cayden Codel, Manuela Veloso, Ruslan Salakhutdinov</i>	
Confirmatory Bayesian Online Change Point Detection in the Covariance Structure of Gaussian Processes	2449
<i>Jiyeon Han, Kyowoon Lee, Anh Tong, Jaesik Choi</i>	
Attribute Aware Pooling for Pedestrian Attribute Recognition	2456
<i>Kai Han, Yunhe Wang, Han Shu, Chuanjian Liu, Chunjing Xu, Chang Xu</i>	
Network Embedding under Partial Monitoring for Evolving Networks	2463
<i>Yu Han, Jie Tang, Qian Chen</i>	
Deep Active Learning with Adaptive Acquisition	2470
<i>Manuel Haussmann, Fred Hamprecht, Melih Kandemir</i>	
One Network for Multi-Domains: Domain Adaptive Hashing with Intersectant Generative Adversarial Networks	2477
<i>Tao He, Yuan-Fang Li, Lianli Gao, Dongxiang Zhang, Jingkuan Song</i>	
Deliberation Learning for Image-to-Image Translation	2484
<i>Tianyu He, Yingce Xia, Jianxin Lin, Xu Tan, Di He, Tao Qin, Zhibo Chen</i>	
Online Learning from Capricious Data Streams: A Generative Approach	2491
<i>Yi He, Baijun Wu, Di Wu, Ege Beyazit, Sheng Chen, Xindong Wu</i>	
Learning Topic Models by Neighborhood Aggregation	2498
<i>Ryohei Hisano</i>	
Group-based Learning of Disentangled Representations with Generalizability for Novel Contents	2506
<i>Haruo Hosoya</i>	
Robust Learning from Noisy Side-information by Semidefinite Programming	2514
<i>En-Liang Hu, Quanming Yao</i>	
Hybrid Item-Item Recommendation via Semi-Parametric Embedding	2521
<i>Peng Hu, Rong Du, Yao Hu, Nan Li</i>	
Cascaded Algorithm-Selection and Hyper-Parameter Optimization with Extreme-Region Upper Confidence Bound Bandit	2528
<i>Yi-Qi Hu, Yang Yu, Jun-Da Liao</i>	
Deep Metric Learning: The Generalization Analysis and an Adaptive Algorithm	2535
<i>Mengdi Huai, Hongfei Xue, Chenglin Miao, Liuyi Yao, Lu Su, Changyou Chen, Aidong Zhang</i>	
Privacy-aware Synthesizing for Crowdsourced Data	2542
<i>Mengdi Huai, Di Wang, Chenglin Miao, Jinhui Xu, Aidong Zhang</i>	
Zeroth-Order Stochastic Alternating Direction Method of Multipliers for Nonconvex Nonsmooth Optimization	2549
<i>Feihu Huang, Shangqian Gao, Songcan Chen, Heng Huang</i>	
Nostalgic Adam: Weighting More of the Past Gradients When Designing the Adaptive Learning Rate	2556
<i>Haiwen Huang, Chang Wang, Bin Dong</i>	
Multi-view Spectral Clustering Network	2563
<i>Zhenyu Huang, Joey Tianyi Zhou, Xi Peng, Changqing Zhang, Hongyuan Zhu, Jiancheng Lv</i>	
Conditions on Features for Temporal Difference-Like Methods to Converge	2570
<i>Marcus Hutter, Samuel Yang-Zhao, Sultan Javed Majeed</i>	
Entangled Kernels	2578
<i>Riikka Huusari, Hachem Kadri</i>	
Efficient Protocol for Collaborative Dictionary Learning in Decentralized Networks	2585
<i>Tsuyoshi Idé, Rudy Raymond, Dzong T. Phan</i>	
SlateQ: A Tractable Decomposition for Reinforcement Learning with Recommendation Sets	2592
<i>Eugene Ie, Vihan Jain, Jing Wang, Sanmit Narvekar, Ritesh Agarwal, Rui Wu, Heng-Tze Cheng, Tushar Chandra, Craig Boutilier</i>	

Accelerating Extreme Classification via Adaptive Feature Agglomeration	2600
<i>Ankit Jalan, Purushottam Kar</i>	
Assumed Density Filtering Q-learning	2607
<i>Heejin Jeong, Clark Zhang, George J. Pappas, Daniel D. Lee</i>	
Learning to Learn Gradient Aggregation by Gradient Descent	2614
<i>Jinlong Ji, Xuhui Chen, Qianlong Wang, Lixing Yu, Pan Li</i>	
Patent Citation Dynamics Modeling via Multi-Attention Recurrent Networks	2621
<i>Taoran Ji, Zhiqian Chen, Nathan Self, Kaiqun Fu, Chang-Tien Lu, Naren Ramakrishnan</i>	
Recurrent Generative Networks for Multi-Resolution Satellite Data: An Application in Cropland Monitoring	2628
<i>Xiaowei Jia, Mengdie Wang, Ankush Khandelwal, Anuj Karpatne, Vipin Kumar</i>	
Dynamic Hypergraph Neural Networks	2635
<i>Jianwen Jiang, Yuxuan Wei, Yifan Feng, Jingxuan Cao, Yue Gao</i>	
Convolutional Gaussian Embeddings for Personalized Recommendation with Uncertainty	2642
<i>Junyang Jiang, Deqing Yang, Yanghua Xiao, Chenlu Shen</i>	
Robust Low-Tubal-Rank Tensor Completion via Convex Optimization	2649
<i>Qiang Jiang, Michael Ng</i>	
CensNet: Convolution with Edge-Node Switching in Graph Neural Networks	2656
<i>Xiaodong Jiang, Pengsheng Ji, Sheng Li</i>	
Network-Specific Variational Auto-Encoder for Embedding in Attribute Networks	2663
<i>Di Jin, Bingyi Li, Pengfei Jiao, Dongxiao He, Weixiong Zhang</i>	
Hypergraph Induced Convolutional Manifold Networks	2670
<i>Taisong Jin, Liujuan Cao, Baochang Zhang, Xiaoshuai Sun, Cheng Deng, Rongrong Ji</i>	
Submodular Batch Selection for Training Deep Neural Networks	2677
<i>K J Joseph, Vamshi Teja R, Krishnakant Singh, Vineeth N Balasubramanian</i>	
Obstacle Tower: A Generalization Challenge in Vision, Control, and Planning	2684
<i>Arthur Juliani, Ahmed Khalifa, Vincent-Pierre Berges, Jonathan Harper, Ervin Teng, Hunter Henry, Adam Crespi, Julian Togelius, Danny Lange</i>	
Interactive Teaching Algorithms for Inverse Reinforcement Learning	2692
<i>Parameswaran Kamalaruban, Rati Devidze, Volkan Cevher, Adish Singla</i>	
Multiple Partitions Aligned Clustering	2701
<i>Zhao Kang, Zipeng Guo, Shudong Huang, Siying Wang, Wenyu Chen, Yuanzhang Su, Zenglin Xu</i>	
Twin-Systems to Explain Artificial Neural Networks using Case-Based Reasoning: Comparative Tests of Feature-Weighting Methods in ANN-CBR Twins for XAI	2708
<i>Eoin M. Kenny, Mark T. Keane</i>	
What to Expect of Classifiers? Reasoning about Logistic Regression with Missing Features	2716
<i>Pasha Khosravi, Yitao Liang, YooJung Choi, Guy Van den Broeck</i>	
Outlier Detection for Time Series with Recurrent Autoencoder Ensembles	2725
<i>Tung Kieu, Bin Yang, Chenjuan Guo, Christian S. Jensen</i>	
DeepMellow: Removing the Need for a Target Network in Deep Q-Learning	2733
<i>Seungchan Kim, Kavosh Asadi, Michael Littman, George Konidaris</i>	
Sequential and Diverse Recommendation with Long Tail	2740
<i>Yejin Kim, Kwangseob Kim, Chanyoung Park, Hwanjo Yu</i>	
Single-Channel Signal Separation and Deconvolution with Generative Adversarial Networks	2747
<i>Qiuqiang Kong, Yong Xu, Philip J. B. Jackson, Wenwu Wang, Mark D. Plumbley</i>	
Autoregressive Policies for Continuous Control Deep Reinforcement Learning	2754
<i>Dmytro Korenkevych, A. Rupam Mahmood, Gautham Vasan, James Bergstra</i>	
Adaptive Ensemble Active Learning for Drifting Data Stream Mining	2763
<i>Bartosz Krawczyk, Alberto Cano</i>	
Learning Sound Events from Webly Labeled Data	2772
<i>Anurag Kumar, Ankit Shah, Alexander Hauptmann, Bhiksha Raj</i>	
Harnessing the Vulnerability of Latent Layers in Adversarially Trained Models	2779
<i>Nupur Kumari, Mayank Singh, Abhishek Sinha, Harshitha Machiraju, Balaji Krishnamurthy, Vineeth N Balasubramanian</i>	
Perturbed-History Exploration in Stochastic Multi-Armed Bandits	2786
<i>Branislav Kveton, Csaba Szepesvári, Mohammad Ghavamzadeh, Craig Boutilier</i>	
Meta Reinforcement Learning with Task Embedding and Shared Policy	2794
<i>Lin Lan, Zhenguo Li, Xiaohong Guan, Pinghui Wang</i>	
The Dangers of Post-hoc Interpretability: Unjustified Counterfactual Explanations	2801
<i>Thibault Laugel, Marie-Jeanne Lesot, Christophe Marsala, Xavier Renard, Marcin Detryniecki</i>	
Correlation-Sensitive Next-Basket Recommendation	2808
<i>Duc-Trong Le, Hady W. Lauw, Yuan Fang</i>	

Learning Multiple Maps from Conditional Ordinal Triplets	2815
<i>Dung D. Le, Hady W. Lauw</i>	
Learning Generative Adversarial Networks from Multiple Data Sources	2823
<i>Trung Le, Quan Hoang, Hung Vu, Tu Dinh Nguyen, Hung Bui, Dinh Phung</i>	
Action Space Learning for Heterogeneous User Behavior Prediction	2830
<i>Dongha Lee, Chanyoung Park, Hyunjun Ju, Junyoung Hwang, Hwanjo Yu</i>	
Learning Shared Knowledge for Deep Lifelong Learning using Deconvolutional Networks	2837
<i>Seungwon Lee, James Stokes, Eric Eaton</i>	
Similarity Preserving Representation Learning for Time Series Clustering	2845
<i>Qi Lei, Jinfeng Yi, Roman Vaculin, Lingfei Wu, Inderjit S. Dhillon</i>	
Differentially Private Optimal Transport: Application to Domain Adaptation	2852
<i>Nam LeTien, Amaury Habrard, Marc Sebban</i>	
Cascading Non-Stationary Bandits: Online Learning to Rank in the Non-Stationary Cascade Model	2859
<i>Chang Li, Maarten de Rijke</i>	
A Review-Driven Neural Model for Sequential Recommendation	2866
<i>Chenliang Li, Xichuan Niu, Xiangyang Luo, Zhenzhong Chen, Cong Quan</i>	
Hierarchical Representation Learning for Bipartite Graphs	2873
<i>Chong Li, Kunyang Jia, Dan Shen, C.J. Richard Shi, Hongxia Yang</i>	
Multi-Class Learning using Unlabeled Samples: Theory and Algorithm	2880
<i>Jian Li, Yong Liu, Rong Yin, Weiping Wang</i>	
Approximate Manifold Regularization: Scalable Algorithm and Generalization Analysis	2887
<i>Jian Li, Yong Liu, Rong Yin, Weiping Wang</i>	
Dense Transformer Networks for Brain Electron Microscopy Image Segmentation	2894
<i>Jun Li, Yongjun Chen, Lei Cai, Ian Davidson, Shuiwang Ji</i>	
Learning Interpretable Deep State Space Model for Probabilistic Time Series Forecasting	2901
<i>Longyuan Li, Junchi Yan, Xiaokang Yang, Yaohui Jin</i>	
Improving the Robustness of Deep Neural Networks via Adversarial Training with Triplet Loss	2909
<i>Pengcheng Li, Jinfeng Yi, Bowen Zhou, Lijun Zhang</i>	
Flexible Multi-View Representation Learning for Subspace Clustering	2916
<i>Ruihuang Li, Changqing Zhang, Qinghua Hu, Pengfei Zhu, Zheng Wang</i>	

VOLUME 5

Improved Algorithm on Online Clustering of Bandits	2923
<i>Shuai Li, Wei Chen, Shuai Li, Kwong-Sak Leung</i>	
Linear Time Complexity Time Series Clustering with Symbolic Pattern Forest	2930
<i>Xiaosheng Li, Jessica Lin, Liang Zhao</i>	
Learning Network Embedding with Community Structural Information	2937
<i>Yu Li, Ying Wang, Tingting Zhang, Jiawei Zhang, Yi Chang</i>	
ARMIN: Towards a More Efficient and Light-weight Recurrent Memory Network	2944
<i>Zhangheng Li, Jia-Xing Zhong, Jingjia Huang, Tao Zhang, Thomas Li, Ge Li</i>	
Deep Adversarial Multi-view Clustering Network	2952
<i>Zhaoyang Li, Qianqian Wang, Zhiqiang Tao, Quanxue Gao, Zhaohua Yang</i>	
GCN-LASE: Towards Adequately Incorporating Link Attributes in Graph Convolutional Networks	2959
<i>Ziyao Li, Liang Zhang, Guojie Song</i>	
Learning K-way D-dimensional Discrete Embedding for Hierarchical Data Visualization and Retrieval	2966
<i>Xiaoyuan Liang, Martin Renqiang Min, Hongyu Guo, Gulling Wang</i>	
Worst-Case Discriminative Feature Selection	2973
<i>Shuangli Liao, Quanxue Gao, Feiping Nie, Yang Liu, Xiangdong Zhang</i>	
Image-to-Image Translation with Multi-Path Consistency Regularization	2980
<i>Jianxin Lin, Yingce Xia, Yijun Wang, Tao Qin, Zhibo Chen</i>	
Balanced Clustering: A Uniform Model and Fast Algorithm	2987
<i>Weibo Lin, Zhu He, Mingyu Xiao</i>	
Feature Prioritization and Regularization Improve Standard Accuracy and Adversarial Robustness	2994
<i>Chihuang Liu, Joseph JaJa</i>	
Learning Instance-wise Sparsity for Accelerating Deep Models	3001
<i>Chuanjian Liu, Yunhe Wang, Kai Han, Chunjing Xu, Chang Xu</i>	
Learning Robust Distance Metric with Side Information via Ratio Minimization of Orthogonally Constrained L21-Norm Distances	3008
<i>Kai Liu, Lodewijk Brand, Hua Wang, Feiping Nie</i>	

Prototype Propagation Networks (PPN) for Weakly-supervised Few-shot Learning on Category Graph	3015
<i>Lu Liu, Tianyi Zhou, Guodong Long, Jing Jiang, Lina Yao, Chengqi Zhang</i>	
Margin Learning Embedded Prediction for Video Anomaly Detection with A Few Anomalies	3023
<i>Wen Liu, Weixin Luo, Zhengxin Li, Peilin Zhao, Shenghua Gao</i>	
Supervised Short-Length Hashing	3031
<i>Xingbo Liu, Xiushan Nie, Quan Zhou, Xiaoming Xi, Lei Zhu, Yilong Yin</i>	
Graph and Autoencoder Based Feature Extraction for Zero-shot Learning	3038
<i>Yang Liu, Deyan Xie, Quanxue Gao, Jungong Han, Shujian Wang, Xinbo Gao</i>	
Accelerated Incremental Gradient Descent using Momentum Acceleration with Scaling Factor	3045
<i>Yuanyuan Liu, Fanhua Shang, Licheng Jiao</i>	
Omnidirectional Scene Text Detection with Sequential-free Box Discretization	3052
<i>Yuliang Liu, Sheng Zhang, Lianwen Jin, Lele Xie, Yaqiang Wu, Zhepeng Wang</i>	
Hi-Fi Ark: Deep User Representation via High-Fidelity Archive Network	3059
<i>Zheng Liu, Yu Xing, Fangzhao Wu, Mingxiao An, Xing Xie</i>	
Learning Low-precision Neural Networks without Straight-Through Estimator (STE)	3066
<i>Zhi-Gang Liu, Matthew Mattina</i>	
Parametric Manifold Learning of Gaussian Mixture Models	3073
<i>Ziquan Liu, Lei Yu, Janet H. Hsiao, Antoni B. Chan</i>	
Multi-Objective Generalized Linear Bandits	3080
<i>Shiyin Lu, Guanghui Wang, Yao Hu, Lijun Zhang</i>	
Knowledge Amalgamation from Heterogeneous Networks by Common Feature Learning	3087
<i>Sihui Luo, Xinchao Wang, Gongfan Fang, Yao Hu, Dapeng Tao, Mingli Song</i>	
E²GAN: End-to-End Generative Adversarial Network for Multivariate Time Series Imputation	3094
<i>Yonghong Luo, Ying Zhang, Xiangrui Cai, Xiaojie Yuan</i>	
Weakly Supervised Multi-Label Learning via Label Enhancement	3101
<i>JiaQi Lv, Ning Xu, RenYi Zheng, Xin Geng</i>	

MACHINE LEARNING M-Z

AttnSense: Multi-level Attention Mechanism For Multimodal Human Activity Recognition	3109
<i>HaoJie Ma, Wenzhong Li, Xiao Zhang, Songcheng Gao, Sanglu Lu</i>	
Monte Carlo Tree Search for Policy Optimization	3116
<i>Xiaobai Ma, Katherine Driggs-Campbell, Zongzhang Zhang, Mykel J. Kochenderfer</i>	
Coarse-to-Fine Image Inpainting via Region-wise Convolutions and Non-Local Correlation	3123
<i>Yuqing Ma, Xianglong Liu, Shihao Bai, Lei Wang, Dailan He, Aishan Liu</i>	
On Principled Entropy Exploration in Policy Optimization	3130
<i>Jincheng Mei, Chenjun Xiao, Ruitong Huang, Dale Schuurmans, Martin Müller</i>	
Anytime Bottom-Up Rule Learning for Knowledge Graph Completion	3137
<i>Christian Meilicke, Melisachew Wudage Chekol, Daniel Ruffinelli, Heiner Stuckenschmidt</i>	
Unsupervised Hierarchical Temporal Abstraction by Simultaneously Learning Expectations and Representations	3144
<i>Katherine Metcalfe, David Leake</i>	
Meta-Learning for Low-resource Natural Language Generation in Task-oriented Dialogue Systems	3151
<i>Fei Mi, Minlie Huang, Jiyong Zhang, Boi Faltings</i>	
Robust Flexible Feature Selection via Exclusive L21 Regularization	3158
<i>Di Ming, Chris Ding</i>	
Advantage Amplification in Slowly Evolving Latent-State Environments	3165
<i>Martin Mladenov, Ofer Meshi, Jayden Ooi, Dale Schuurmans, Craig Boutilier</i>	
Deep Variational Koopman Models: Inferring Koopman Observations for Uncertainty-Aware Dynamics Modeling and Control	3173
<i>Jeremy Morton, Freddie D. Witherden, Mykel J. Kochenderfer</i>	
DyAt Nets: Dynamic Attention Networks for State Forecasting in Cyber-Physical Systems	3180
<i>Nikhil Muralidhar, Sathappan Muthiah, Naren Ramakrishnan</i>	
Outlier-Robust Multi-Aspect Streaming Tensor Completion and Factorization	3187
<i>Mehrnaz Najafi, Lifang He, Philip S. Yu</i>	
Incremental Learning of Planning Actions in Model-Based Reinforcement Learning	3195
<i>Jun Hao Alvin Ng, Ronald P. A. Petrick</i>	
Group LASSO with Asymmetric Structure Estimation for Multi-Task Learning	3202
<i>Saullo H. G. Oliveira, André R. Gonçalves, Fernando J. Von Zuben</i>	

Hill Climbing on Value Estimates for Search-control in Dyna	3209
<i>Yangchen Pan, Hengshuai Yao, Amir-massoud Farahmand, Martha White</i>	
Indirect Trust is Simple to Establish	3216
<i>Elham Parhizkar, Mohammad Hossein Nikravan, Sandra Zilles</i>	
Exploiting Interaction Links for Node Classification with Deep Graph Neural Networks	3223
<i>Hogun Park, Jennifer Neville</i>	
Improving Cross-lingual Entity Alignment via Optimal Transport	3231
<i>Shichao Pei, Lu Yu, Xiangliang Zhang</i>	
Fine-grained Event Categorization with Heterogeneous Graph Convolutional Networks	3238
<i>Hao Peng, Jianxin Li, Qiran Gong, Yangqiu Song, Yuanxin Ning, Kunfeng Lai, Philip S. Yu</i>	
A Practical Semi-Parametric Contextual Bandit	3246
<i>Yi Peng, Miao Xie, Jiahao Liu, Xuying Meng, Nan Li, Cheng Yang, Tao Yao, Rong Jin</i>	
Graph Space Embedding	3253
<i>João Pereira, Albert K. Groen, Erik S. G. Stroes, Evgeni Levin</i>	
An Atari Model Zoo for Analyzing, Visualizing, and Comparing Deep Reinforcement Learning Agents	3260
<i>Felipe Petroski Such, Vashisht Madhavan, Rosanne Liu, Rui Wang, Pablo Samuel Castro, Yulun Li, Jiale Zhi, Ludwig Schubert, Marc G. Bellemare, Jeff Clune, Joel Lehman</i>	
Improving representation learning in autoencoders via multidimensional interpolation and dual regularizations	3268
<i>Sheng Qian, Guanyue Li, Wen-Ming Cao, Cheng Liu, Si Wu, Hau San Wong</i>	
Scalable Bayesian Non-linear Matrix Completion	3275
<i>Xiangju Qin, Paul Blomstedt, Samuel Kaski</i>	
Noise-Resilient Similarity Preserving Network Embedding for Social Networks	3282
<i>Zhenyu Qiu, Wenbin Hu, Jia Wu, ZhongZheng Tang, Xiaohua Jia</i>	
Fairwalk: Towards Fair Graph Embedding	3289
<i>Tahleen Rahman, Bartłomiej Surma, Michael Backes, Yang Zhang</i>	
Automated Machine Learning with Monte-Carlo Tree Search	3296
<i>Herilalaina Rakotoarison, Marc Schoenauer, Michèle Sebag</i>	
Successor Options: An Option Discovery Framework for Reinforcement Learning	3304
<i>Rahul Ramesh, Manan Tomar, Balaraman Ravindran</i>	
Unifying the Stochastic and the Adversarial Bandits with Knapsack	3311
<i>Anshuka Rangi, Massimo Franceschetti, Long Tran-Thanh</i>	
Label distribution learning with label-specific features	3318
<i>Tingting Ren, Xiuyi Jia, Weiwei Li, Lei Chen, Zechao Li</i>	
Label Distribution Learning with Label Correlations via Low-Rank Approximation	3325
<i>Tingting Ren, Xiuyi Jia, Weiwei Li, Shu Zhao</i>	
Closed-Loop Memory GAN for Continual Learning	3332
<i>Amanda Rios, Laurent Itti</i>	
Complementary Learning for Overcoming Catastrophic Forgetting Using Experience Replay	3339
<i>Mohammad Rostami, Soheil Kolouri, Praveen K. Pilly</i>	
Discovering Regularities from Traditional Chinese Medicine Prescriptions via Bipartite Embedding Model	3346
<i>Chunyang Ruan, Jiangang Ma, Ye Wang, Yanchun Zhang, Yun Yang</i>	
A Degeneracy Framework for Scalable Graph Autoencoders	3353
<i>Guillaume Salha, Romain Hennequin, Viet Anh Tran, Michalis Vazirgiannis</i>	
Deterministic Routing between Layout Abstractions for Multi-Scale Classification of Visually Rich Documents	3360
<i>Ritesh Sarkhel, Arnab Nandi</i>	
SynthNet: Learning to Synthesize Music End-to-End	3367
<i>Florin Schimbinschi, Christian Walder, Sarah M. Erfani, James Bailey</i>	
Weakly Supervised Multi-task Learning for Semantic Parsing	3375
<i>Bo Shao, Yeyun Gong, Junwei Bao, Jianshu Ji, Guihong Cao, Xiaola Lin, Nan Duan</i>	
Community Detection and Link Prediction via Cluster-driven Low-rank Matrix Completion	3382
<i>Junming Shao, Zhong Zhang, Zhongjing Yu, Jun Wang, Yi Zhao, Qinli Yang</i>	
On the Effectiveness of Low Frequency Perturbations	3389
<i>Yash Sharma, Gavin Weiguang Ding, Marcus A. Brubaker</i>	
A Part Power Set Model for Scale-Free Person Retrieval	3397
<i>Yunhang Shen, Rongrong Ji, Xiaopeng Hong, Feng Zheng, Xiaowei Guo, Yongjian Wu, Feiyue Huang</i>	
Rapid Performance Gain through Active Model Reuse	3404
<i>Feng Shi, Yu-Feng Li</i>	

A Convergence Analysis of Distributed SGD with Communication-Efficient Gradient Sparsification	3411
<i>Shaohuai Shi, Kaiyong Zhao, Qiang Wang, Zhenheng Tang, Xiaowen Chu</i>	
Quadruply Stochastic Gradients for Large Scale Nonlinear Semi-Supervised AUC Optimization	3418
<i>Wanli Shi, Bin Gu, Xiang Li, Xiang Geng, Heng Huang</i>	
Soft Policy Gradient Method for Maximum Entropy Deep Reinforcement Learning	3425
<i>Wenjie Shi, Shiji Song, Cheng Wu</i>	
Gradient Boosting with Piece-Wise Linear Regression Trees	3432
<i>Yu Shi, Jian Li, Zhize Li</i>	
The Pupil Has Become the Master: Teacher-Student Model-Based Word Embedding Distillation with Ensemble Learning	3439
<i>Bonggun Shin, Hao Yang, Jinho D. Choi</i>	
A Principled Approach for Learning Task Similarity in Multitask Learning	3446
<i>Changjian Shui, Mahdieh Abbasi, Louis-Émile Robitaille, Boyu Wang, Christian Gagné</i>	
Structure Learning for Safe Policy Improvement	3453
<i>Thiago D. Simão, Matthijs T. J. Spaan</i>	
Play and Prune: Adaptive Filter Pruning for Deep Model Compression	3460
<i>Pravendra Singh, Vinay Kumar Verma, Piyush Rai, Vinay P. Namboodiri</i>	
Solving Continual Combinatorial Selection via Deep Reinforcement Learning	3467
<i>Hyungseok Song, Hyeryung Jang, Hai H. Tran, Se-eun Yoon, Kyunghwan Son, Donggyu Yun, Hyoju Chung, Yung Yi</i>	
Playing FPS Games With Environment-Aware Hierarchical Reinforcement Learning	3475
<i>Shihong Song, Jiayi Weng, Hang Su, Dong Yan, Haosheng Zou, Jun Zhu</i>	
Parallel Wasserstein Generative Adversarial Nets with Multiple Discriminators	3483
<i>Yuxin Su, Shenglin Zhao, Xixian Chen, Irwin King, Michael Lyu</i>	
Finding Statistically Significant Interactions between Continuous Features	3490
<i>Mahito Sugiyama, Karsten Borgwardt</i>	
Fast and Robust Multi-View Multi-Task Learning via Group Sparsity	3499
<i>Lu Sun, Canh Hao Nguyen, Hiroshi Mamitsuka</i>	
Multiplicative Sparse Feature Decomposition for Efficient Multi-View Multi-Task Learning	3506
<i>Lu Sun, Canh Hao Nguyen, Hiroshi Mamitsuka</i>	
Adversarial Imitation Learning from Incomplete Demonstrations	3513
<i>Mingfei Sun, Xiaojuan Ma</i>	
Heavy-ball Algorithms Always Escape Saddle Points	3520
<i>Tao Sun, Dongsheng Li, Zhe Quan, Hao Jiang, Shengguo Li, Yong Dou</i>	
MEGAN: A Generative Adversarial Network for Multi-View Network Embedding	3527
<i>Yiwei Sun, Suhang Wang, Tsung-Yu Hsieh, Xianfeng Tang, Vasant Honavar</i>	
Metric Learning on Healthcare Data with Incomplete Modalities	3534
<i>Qiuling Suo, Weida Zhong, Fenglong Ma, Ye Yuan, Jing Gao, Aidong Zhang</i>	
HMLasso: Lasso with High Missing Rate	3541
<i>Masaaki Takada, Hironori Fujisawa, Takeichiro Nishikawa</i>	
Deeply-learned Hybrid Representations for Facial Age Estimation	3548
<i>Zichang Tan, Yang Yang, Jun Wan, Guodong Guo, Stan Z. Li</i>	
AugBoost: Gradient Boosting Enhanced with Step-Wise Feature Augmentation	3555
<i>Philip Tammor, Lior Rokach</i>	
Adversarial Graph Embedding for Ensemble Clustering	3562
<i>Zhiqiang Tao, Hongfu Liu, Jun Li, Zhaowen Wang, Yun Fu</i>	
Hierarchical Inter-Attention Network for Document Classification with Multi-Task Learning	3569
<i>Bing Tian, Yong Zhang, Jin Wang, Chunxiao Xing</i>	
Image Captioning with Compositional Neural Module Networks	3576
<i>Junjiao Tian, Jean Oh</i>	
Imitation Learning from Video by Leveraging Proprioception	3585
<i>Faraz Torabi, Garrett Warnell, Peter Stone</i>	
Exchangeability and Kernel Invariance in Trained MLPs	3592
<i>Russell Tsuchida, Fred Roosta, Marcus Gallagher</i>	
Deeper Connections between Neural Networks and Gaussian Processes Speed-up Active Learning	3599
<i>Evgenii Tsybalov, Sergei Makarychev, Alexander Shapeev, Maxim Panov</i>	
Object Detection based Deep Unsupervised Hashing	3606
<i>Rong-Cheng Tu, Xian-Ling Mao, Bo-Si Feng, Shu-ying Yu</i>	
Ensemble-based Ultrahigh-dimensional Variable Screening	3613
<i>Wei Tu, Dong Yang, Linglong Kong, Menglu Che, Qian Shi, Guodong Li, Guangjian Tian</i>	
Learning to Interpret Satellite Images using Wikipedia	3620
<i>Burak Uzkent, Evan Sheehan, Chenlin Meng, Zhongyi Tang, Marshall Burke, David Lobell, Stefano Ermon</i>	

DeepCU: Integrating both Common and Unique Latent Information for Multimodal Sentiment Analysis	3627
<i>Sunny Verma, Chen Wang, Liming Zhu, Wei Liu</i>	
Interpolation Consistency Training for Semi-supervised Learning	3635
<i>Vikas Verma, Alex Lamb, Juho Kannala, Yoshua Bengio, David Lopez-Paz</i>	
Sharing Experience in Multitask Reinforcement Learning	3642
<i>Tung-Long Vuong, Do-Van Nguyen, Tai-Long Nguyen, Cong-Minh Bui, Hai-Dang Kieu, Viet-Cuong Ta, Quoc-Long Tran, Thanh-Ha Le</i>	

VOLUME 6

Planning with Expectation Models	3649
<i>Yi Wan, Muhammad Zaheer, Adam White, Martha White, Richard S. Sutton</i>	
Recurrent Existence Determination Through Policy Optimization	3656
<i>Baoxiang Wang</i>	
Boundary Perception Guidance: A Scribble-Supervised Semantic Segmentation Approach	3663
<i>Bin Wang, Guojun Qi, Sheng Tang, Tianzhu Zhang, Yunchao Wei, Linghui Li, Yongdong Zhang</i>	
Attributed Graph Clustering: A Deep Attentional Embedding Approach	3670
<i>Chun Wang, Shirui Pan, Ruiqi Hu, Guodong Long, Jing Jiang, Chengqi Zhang</i>	
Spectral Perturbation Meets Incomplete Multi-view Data	3677
<i>Hao Wang, Linlin Zong, Bing Liu, Yan Yang, Wei Zhou</i>	
Measuring Structural Similarities in Finite MDPs	3684
<i>Hao Wang, Shaokang Dong, Ling Shao</i>	
Discriminative and Correlative Partial Multi-Label Learning	3691
<i>Haobo Wang, Weiwei Liu, Yang Zhao, Chen Zhang, Tianlei Hu, Gang Chen</i>	
DMRAN:A Hierarchical Fine-Grained Attention-Based Network for Recommendation	3698
<i>Huizhao Wang, Guanfeng Liu, An Liu, Zhixu Li, Kai Zheng</i>	
CLVSA: A Convolutional LSTM Based Variational Sequence-to-Sequence Model with Attention for Predicting Trends of Financial Markets	3705
<i>Jia Wang, Tong Sun, Benyuan Liu, Yu Cao, Hongwei Zhu</i>	
Classification with Label Distribution Learning	3712
<i>Jing Wang, Xin Geng</i>	
Attributed Subspace Clustering	3719
<i>Jing Wang, Linchuan Xu, Feng Tian, Atsushi Suzuki, Changqing Zhang, Kenji Yamanishi</i>	
Deep Cascade Generation on Point Sets	3726
<i>Kaiqi Wang, Ke Chen, Kui Jia</i>	
Discrete Binary Coding based Label Distribution Learning	3733
<i>Ke Wang, Xin Geng</i>	
Differentially Private Iterative Gradient Hard Thresholding for Sparse Learning	3740
<i>Lingxiao Wang, Quanquan Gu</i>	
MUSICAL: Multi-Scale Image Contextual Attention Learning for Inpainting	3748
<i>Ning Wang, Jingyuan Li, Lefei Zhang, Bo Du</i>	
Partial Label Learning with Unlabeled Data	3755
<i>Qian-Wei Wang, Yu-Feng Li, Zhi-Hua Zhou</i>	
Heterogeneous Graph Matching Networks for Unknown Malware Detection	3762
<i>Shen Wang, Zhengzhang Chen, Xiao Yu, Ding Li, Jingchao Ni, Lu-An Tang, Jiaping Gui, Zhichun Li, Haifeng Chen, Philip S. Yu</i>	
Modeling Multi-Purpose Sessions for Next-Item Recommendations via Mixture-Channel Purpose Routing Networks	3771
<i>Shoujin Wang, Liang Hu, Yan Wang, Quan Z. Sheng, Mehmet Orgun, Longbing Cao</i>	
Multi-view Clustering via Late Fusion Alignment Maximization	3778
<i>Siwei Wang, Xinwang Liu, En Zhu, Chang Tang, Jiyuan Liu, Jingtao Hu, Jingyuan Xia, Jianping Yin</i>	
COP: Customized Deep Model Compression via Regularized Correlation-Based Filter-Level Pruning	3785
<i>Wenxiao Wang, Cong Fu, Jishun Guo, Deng Cai, Xiaofei He</i>	
Position Focused Attention Network for Image-Text Matching	3792
<i>Yaxiong Wang, Hao Yang, Xueming Qian, Lin Ma, Jing Lu, Biao Li, Xin Fan</i>	
Tag2Gauss: Learning Tag Representations via Gaussian Distribution in Tagged Networks	3799
<i>Yun Wang, Lun Du, Guojie Song, Xiaojun Ma, Lichen Jin, Wei Lin, Fei Sun</i>	
Weak Supervision Enhanced Generative Network for Question Generation	3806
<i>Yutong Wang, Jiyuan Zheng, Qijiong Liu, Zhou Zhao, Jun Xiao, Yueting Zhuang</i>	

Unified Embedding Model over Heterogeneous Information Network for Personalized Recommendation	3813
<i>Zekai Wang, Hongzhi Liu, Yingpeng Du, Zhonghai Wu, Xing Zhang</i>	
Interactive Reinforcement Learning with Dynamic Reuse of Prior Knowledge from Human and Agent Demonstrations	3820
<i>Zhaodong Wang, Matthew E. Taylor</i>	
Hierarchical Diffusion Attention Network	3828
<i>Zhitao Wang, Wenjie Li</i>	
Learning Multi-Objective Rewards and User Utility Function in Contextual Bandits for Personalized Ranking	3835
<i>Nirandika Wanigasekara, Yuxuan Liang, Siong Thye Goh, Ye Liu, Joseph Jay Williams, David S. Rosenblum</i>	
Learning for Tail Label Data: A Label-Specific Feature Approach	3842
<i>Tong Wei, Wei-Wei Tu, Yu-Feng Li</i>	
Bayesian Uncertainty Matching for Unsupervised Domain Adaptation	3849
<i>Jun Wen, Nenggan Zheng, Junsong Yuan, Zhefeng Gong, Changyou Chen</i>	
RobustTrend: A Huber Loss with a Combined First and Second Order Difference Regularization for Time Series Trend Filtering	3856
<i>Qingsong Wen, Jingkun Gao, Xiaomin Song, Liang Sun, Jian Tan</i>	
Neural News Recommendation with Attentive Multi-View Learning	3863
<i>Chuhan Wu, Fangzhao Wu, Mingxiao An, Jianqiang Huang, Yongfeng Huang, Xing Xie</i>	
PD-GAN: Adversarial Learning for Personalized Diversity-Promoting Recommendation	3870
<i>Qiong Wu, Yong Liu, Chunyan Miao, Binqiang Zhao, Yin Zhao, Lu Guan</i>	
Feature Evolution Based Multi-Task Learning for Collaborative Filtering with Social Trust	3877
<i>Qitian Wu, Lei Jiang, Xiaofeng Gao, Xiaochun Yang, Guihai Chen</i>	
Multi-View Multi-Label Learning with View-Specific Information Extraction	3884
<i>Xuan Wu, Qing-Guo Chen, Yao Hu, Dengbao Wang, Xiaodong Chang, Xiaobo Wang, Min-Ling Zhang</i>	
Trend-Aware Tensor Factorization for Job Skill Demand Analysis	3891
<i>Xunxian Wu, Tong Xu, Hengshu Zhu, Le Zhang, Enhong Chen, Hui Xiong</i>	
Graph Convolutional Networks on User Mobility Heterogeneous Graphs for Social Relationship Inference	3898
<i>Yongji Wu, Defu Lian, Shuwei Jin, Enhong Chen</i>	
BPAM: Recommendation Based on BP Neural Network with Attention Mechanism	3905
<i>Wu-Dong Xi, Ling Huang, Chang-Dong Wang, Yin-Yu Zheng, Jianhuang Lai</i>	
Incremental Few-Shot Learning for Pedestrian Attribute Recognition	3912
<i>Liyuy Xiang, Xiaoming Jin, Guiguang Ding, Jungong Han, Leida Li</i>	
Reparameterizable Subset Sampling via Continuous Relaxations	3919
<i>Sang Michael Xie, Stefano Ermon</i>	
CFM: Convolutional Factorization Machines for Context-Aware Recommendation	3926
<i>Xin Xin, Bo Chen, Xiangnan He, Dong Wang, Yue Ding, Joemon Jose</i>	
Adversarial Incomplete Multi-view Clustering	3933
<i>Cai Xu, Ziyu Guan, Wei Zhao, Hongchang Wu, Yunfei Niu, Beilei Ling</i>	
Graph Contextualized Self-Attention Network for Session-based Recommendation	3940
<i>Chengfeng Xu, Pengpeng Zhao, Yanchi Liu, Victor S. Sheng, Jiajie Xu, Fuzhen Zhuang, Junhua Fang, Xiaofang Zhou</i>	
Spatio-Temporal Attentive RNN for Node Classification in Temporal Attributed Graphs	3947
<i>Dongkuan Xu, Wei Cheng, Dongsheng Luo, Xiao Liu, Xiang Zhang</i>	
Learning a Generative Model for Fusing Infrared and Visible Images via Conditional Generative Adversarial Network with Dual Discriminators	3954
<i>Han Xu, Pengwei Liang, Wei Yu, Junjun Jiang, Jiayi Ma</i>	
Topology Attack and Defense for Graph Neural Networks: An Optimization Perspective	3961
<i>Kaidi Xu, Hongge Chen, Sijia Liu, Pin-Yu Chen, Tsui-Wei Weng, Mingyi Hong, Xue Lin</i>	
MR-GNN: Multi-Resolution and Dual Graph Neural Network for Predicting Structured Entity Interactions	3968
<i>Nuo Xu, Pinghui Wang, Long Chen, Jing Tao, Junzhou Zhao</i>	
Commit Message Generation for Source Code Changes	3975
<i>Shengbin Xu, Yuan Yao, Feng Xu, Tianxiao Gu, Hanghang Tong, Jian Lu</i>	
Latent Semantics Encoding for Label Distribution Learning	3982
<i>Suping Xu, Lin Shang, Furao Shen</i>	
Learning Image-Specific Attributes by Hyperbolic Neighborhood Graph Propagation	3989
<i>Xiaofeng Xu, Ivor W. Tsang, Xiaofeng Cao, Ruiheng Zhang, Chuancai Liu</i>	
Zero-shot Metric Learning	3996
<i>Xinyi Xu, Huanhuan Cao, Yanhua Yang, Erkun Yang, Cheng Deng</i>	

On the Convergence of (Stochastic) Gradient Descent with Extrapolation for Non-Convex Minimization	4003
<i>Yi Xu, Zhuoning Yuan, Sen Yang, Rong Jin, Tianbao Yang</i>	
Transfer of Temporal Logic Formulas in Reinforcement Learning	4010
<i>Zhe Xu, Ufuk Topcu</i>	
Deep Spectral Kernel Learning	4019
<i>Hui Xue, Zheng-Fan Wu, Wei-Xiang Sun</i>	
Deep Correlated Predictive Subspace Learning for Incomplete Multi-View Semi-Supervised Classification	4026
<i>Zhe Xue, Junping Du, Dawei Du, Wenqi Ren, Siwei Lyu</i>	
Multi-scale Information Diffusion Prediction with Reinforced Recurrent Networks	4033
<i>Cheng Yang, Jian Tang, Maosong Sun, Ganqu Cui, Zhiyuan Liu</i>	
Learning Strictly Orthogonal p-Order Nonnegative Laplacian Embedding via Smoothed Iterative Reweighted Method	4040
<i>Haoxuan Yang, Kai Liu, Hua Wang, Feiping Nie</i>	
Low-Bit Quantization for Attributed Network Representation Learning	4047
<i>Hong Yang, Shirui Pan, Ling Chen, Chuan Zhou, Peng Zhang</i>	
Topology Optimization based Graph Convolutional Network	4054
<i>Liang Yang, Zesheng Kang, Xiaochun Cao, Di Jin, Bo Yang, Yuanfang Guo</i>	
Dual Self-Paced Graph Convolutional Network: Towards Reducing Attribute Distortions Induced by Topology	4062
<i>Liang Yang, Zhiyang Chen, Junhua Gu, Yuanfang Guo</i>	
Masked Graph Convolutional Network	4070
<i>Liang Yang, Fan Wu, Yingkui Wang, Junhua Gu, Yuanfang Guo</i>	
Deep Multi-Task Learning with Adversarial-and-Cooperative Nets	4078
<i>Pei Yang, Qi Tan, Jieping Ye, Hanghang Tong, Jingrui He</i>	
Legal Judgment Prediction via Multi-Perspective Bi-Feedback Network	4085
<i>Wenmian Yang, Weijia Jia, Xiaojie Zhou, Yutao Luo</i>	
Comprehensive Semi-Supervised Multi-Modal Learning	4092
<i>Yang Yang, Ke-Tao Wang, De-Chuan Zhan, Hui Xiong, Yuan Jiang</i>	
SPAGAN: Shortest Path Graph Attention Network	4099
<i>Yiding Yang, Xinchao Wang, Mingli Song, Junsong Yuan, Dacheng Tao</i>	
On the Estimation of Treatment Effect with Text Covariates	4106
<i>Liuyi Yao, Sheng Li, Yaliang Li, Hongfei Xue, Jing Gao, Aidong Zhang</i>	
Privacy-Preserving Stacking with Application to Cross-organizational Diabetes Prediction	4114
<i>Quanming Yao, Xiawei Guo, James Kwok, Weiwei Tu, Yuqiang Chen, Wenyuan Dai, Qiang Yang</i>	
Multi-View Multiple Clustering	4121
<i>Shixin Yao, Guoxian Yu, Jun Wang, Carlotta Domeniconi, Xiangliang Zhang</i>	
Amalgamating Filtered Knowledge: Learning Task-customized Student from Multi-task Teachers	4128
<i>Jingwen Ye, Xinchao Wang, Yixin Ji, Kairi Ou, Mingli Song</i>	
A Vectorized Relational Graph Convolutional Network for Multi-Relational Network Alignment	4135
<i>Rui Ye, Xin Li, Yujie Fang, Hongyu Zang, Mingzhong Wang</i>	
Distributed Collaborative Feature Selection Based on Intermediate Representation	4142
<i>Xiucui Ye, Hongmin Li, Akira Imakura, Tetsuya Sakurai</i>	
Out-of-sample Node Representation Learning for Heterogeneous Graph in Real-time Android Malware Detection	4150
<i>Yanfang Ye, Shifu Hou, Lingwei Chen, Jingwei Lei, Wenqiang Wan, Jiabin Wang, Qi Xiong, Fudong Shao</i>	
Neural Network based Continuous Conditional Random Field for Fine-grained Crime Prediction	4157
<i>Fei Yi, Zhiwen Yu, Fuzhen Zhuang, Bin Guo</i>	
BN-invariant Sharpness Regularizes the Training Model to Better Generalization	4164
<i>Mingyang Yi, Huishuai Zhang, Wei Chen, Zhi-Ming Ma, Tie-Yan Liu</i>	
Geometric Understanding for Unsupervised Subspace Learning	4171
<i>Shihui Ying, Lipeng Cai, Changzhou He, Yaxin Peng</i>	
Belief Propagation Network for Hard Inductive Semi-Supervised Learning	4178
<i>Jaemin Yoo, Hyunsik Jeon, U Kang</i>	
Metatrace Actor-Critic: Online Step-Size Tuning by Meta-gradient Descent for Reinforcement Learning Control	4185
<i>Kenny Young, Baoxiang Wang, Matthew E. Taylor</i>	
Semi-supervised Three-dimensional Reconstruction Framework with GAN	4192
<i>Chong Yu</i>	
Interpreting and Evaluating Neural Network Robustness	4199
<i>Fuxun Yu, Zhuwei Qin, Chenchen Liu, Liang Zhao, Yanzhi Wang, Xiang Chen</i>	

VAEGAN: A Collaborative Filtering Framework based on Adversarial Variational Autoencoders	4206
<i>Xianwen Yu, Xiaoning Zhang, Yang Cao, Min Xia</i>	
Adaptive User Modeling with Long and Short-Term Preferences for Personalized Recommendation	4213
<i>Zeping Yu, Jianxun Lian, Ahmad Mahmoody, Gongshen Liu, Xing Xie</i>	
Progressive Transfer for Person Re-identification	4220
<i>Zhengxu Yu, Zhongming Jin, Long Wei, Jishun Guo, Jianqiang Huang, Deng Cai, Xiaofei He, Xian-Sheng Hua</i>	
DARec: Deep Domain Adaptation for Cross-Domain Recommendation via Transferring Rating Patterns	4227
<i>Feng Yuan, Lina Yao, Boualem Benatallah</i>	
KCNN: Kernel-wise Quantization to Remarkably Decrease Multiplications in Convolutional Neural Network	4234
<i>Linghua Zeng, Zhangcheng Wang, Xinmei Tian</i>	
Experience Replay Optimization	4243
<i>Daochen Zha, Kwei-Herng Lai, Kaixiong Zhou, Xia Hu</i>	
Positive and Unlabeled Learning with Label Disambiguation	4250
<i>Chuang Zhang, Dexin Ren, Tongliang Liu, Jian Yang, Chen Gong</i>	
Generalized Majorization-Minimization for Non-Convex Optimization	4257
<i>Hu Zhang, Pan Zhou, Yi Yang, Jiashi Feng</i>	
STAR-GCN: Stacked and Reconstructed Graph Convolutional Networks for Recommender Systems	4264
<i>Jiani Zhang, Xingjian Shi, Shenglin Zhao, Irwin King</i>	
Light-Weight Hybrid Convolutional Network for Liver Tumor Segmentation	4271
<i>Jianpeng Zhang, Yutong Xie, Pingping Zhang, Hao Chen, Yong Xia, Chunhua Shen</i>	
ProNE: Fast and Scalable Network Representation Learning	4278
<i>Jie Zhang, Yuxiao Dong, Yan Wang, Jie Tang, Ming Ding</i>	
Towards Robust ResNet: A Small Step but a Giant Leap	4285
<i>Jingfeng Zhang, Bo Han, Laura Wynter, Bryan Kian Hsiang Low, Mohan Kankanhalli</i>	
High Dimensional Bayesian Optimization via Supervised Dimension Reduction	4292
<i>Miao Zhang, Huiqi Li, Steven Su</i>	
Efficient Non-parametric Bayesian Hawkes Processes	4299
<i>Rui Zhang, Christian Walder, Marian-Andrei Rizoiiu, Lexing Xie</i>	
Inferring Substitutable Products with Deep Network Embedding	4306
<i>Shijie Zhang, Hongzhi Yin, Qinyong Wang, Tong Chen, Hongxu Chen, Quoc Viet Hung Nguyen</i>	
Quaternion Collaborative Filtering for Recommendation	4313
<i>Shuai Zhang, Lina Yao, Lucas Vinh Tran, Aston Zhang, Yi Tay</i>	
Feature-level Deeper Self-Attention Network for Sequential Recommendation	4320
<i>Tingting Zhang, Pengpeng Zhao, Yanchi Liu, Victor S. Sheng, Jiajie Xu, Deqing Wang, Guanfeng Liu, Xiaofang Zhou</i>	
Attributed Graph Clustering via Adaptive Graph Convolution	4327
<i>Xiaotong Zhang, Han Liu, Qimai Li, Xiao-Ming Wu</i>	
InteractionNN: A Neural Network for Learning Hidden Features in Sparse Prediction	4334
<i>Xiaowang Zhang, Qiang Gao, Zhiyong Feng</i>	
Multi-Group Encoder-Decoder Networks to Fuse Heterogeneous Data for Next-Day Air Quality Prediction	4341
<i>Yawen Zhang, Qin Lv, Duanfeng Gao, Si Shen, Robert Dick, Michael Hammigan, Qi Liu</i>	
Taming the Noisy Gradient: Train Deep Neural Networks with Small Batch Sizes	4348
<i>Yikai Zhang, Hui Qu, Chao Chen, Dimitris Metaxas</i>	
Accelerated Inference Framework of Sparse Neural Network Based on Nested Bitmask Structure	4355
<i>Yipeng Zhang, Bo Du, Lefei Zhang, Rongchun Li, Yong Dou</i>	
DANE: Domain Adaptive Network Embedding	4362
<i>Yizhou Zhang, Guojie Song, Lun Du, Shuwen Yang, Yilun Jin</i>	
ATTAIN: Attention-based Time-Aware LSTM Networks for Disease Progression Modeling	4369
<i>Yuan Zhang, Xi Yang, Julie Ivy, Min Chi</i>	

VOLUME 7

Scalable Block-Diagonal Locality-Constrained Projective Dictionary Learning	4376
<i>Zhao Zhang, Weiming Jiang, Zheng Zhang, Sheng Li, Guangcan Liu, Jie Qin</i>	
Open-Ended Long-Form Video Question Answering via Hierarchical Convolutional Self-Attention Networks	4383
<i>Zhu Zhang, Zhou Zhao, Zhijie Lin, Jingkuan Song, Xiaofei He</i>	

Localizing Unseen Activities in Video via Image Query	4390
<i>Zhu Zhang, Zhou Zhao, Zhijie Lin, Jingkuan Song, Deng Cai</i>	
Multi-Prototype Networks for Unconstrained Set-based Face Recognition	4397
<i>Jian Zhao, Jianshu Li, Xiaoguang Tu, Fang Zhao, Yuan Xin, Junliang Xing, Hengzhu Liu, Shuicheng Yan, Jiashi Feng</i>	
GAN-EM: GAN Based EM Learning Framework	4404
<i>Wentian Zhao, Shaojie Wang, Zhihui Xie, Jing Shi, Chenliang Xu</i>	
Large Scale Evolving Graphs with Burst Detection	4412
<i>Yifeng Zhao, Xiangwei Wang, Hongxia Yang, Le Song, Jie Tang</i>	
AddGraph: Anomaly Detection in Dynamic Graph Using Attention-based Temporal GCN	4419
<i>Li Zheng, Zhenpeng Li, Jian Li, Zhao Li, Jun Gao</i>	
Metadata-driven Task Relation Discovery for Multi-task Learning	4426
<i>Zimu Zheng, Yuqi Wang, Quanyu Dai, Huadi Zheng, Dan Wang</i>	
BeatGAN: Anomalous Rhythm Detection using Adversarially Generated Time Series	4433
<i>Bin Zhou, Shenghua Liu, Bryan Hooi, Xueqi Cheng, Jing Ye</i>	
Latent Distribution Preserving Deep Subspace Clustering	4440
<i>Lei Zhou, Xiao Bai, Dong Wang, Xianglong Liu, Jun Zhou, Edwin Hancock</i>	
Reinforcement Learning Experience Reuse with Policy Residual Representation	4447
<i>Wenji Zhou, Yang Yu, Yingfeng Chen, Kai Guan, Tangjie Lv, Changjie Fan, Zhi-Hua Zhou</i>	
Collaborative Metric Learning with Memory Network for Multi-Relational Recommender Systems	4454
<i>Xiao Zhou, Danyang Liu, Jianxun Lian, Xing Xie</i>	
One-Shot Texture Retrieval with Global Context Metric	4461
<i>Kai Zhu, Wei Zhai, Zheng-Jun Zha, Yang Cao</i>	
HDI-Forest: Highest Density Interval Regression Forest	4468
<i>Lin Zhu, Jiaying Lu, Yihong Chen</i>	
Prediction of Mild Cognitive Impairment Conversion Using Auxiliary Information	4475
<i>Xiaofeng Zhu</i>	
Simultaneous Representation Learning and Clustering for Incomplete Multi-view Data	4482
<i>Wenzhang Zhuge, Chenping Hou, Xinwang Liu, Hong Tao, Dongyun Yi</i>	
Persistence Bag-of-Words for Topological Data Analysis	4489
<i>Bartosz Zielinski, Michal Lipinski, Mateusz Juda, Matthias Zeppelzauer, Pawel Dlotko</i>	
Exploiting the Sign of the Advantage Function to Learn Deterministic Policies in Continuous Domains	4496
<i>Matthieu Zimmer, Paul Weng</i>	

MACHINE LEARNING APPLICATIONS

Predicting the Visual Focus of Attention in Multi-Person Discussion Videos	4504
<i>Chongyang Bai, Srijan Kumar, Jure Leskovec, Miriam Metzger, Jay F. Nunamaker, V. S. Subrahmanian</i>	
A Quantum-inspired Classical Algorithm for Separable Non-negative Matrix Factorization	4511
<i>Zhihui Chen, Yinan Li, Xiaoming Sun, Pei Yuan, Jialin Zhang</i>	
MLRDA: A Multi-Task Semi-Supervised Learning Framework for Drug-Drug Interaction Prediction	4518
<i>Xu Chu, Yang Lin, Yasha Wang, Leye Wang, Jiangtao Wang, Jingyue Gao</i>	
Combining ADMM and the Augmented Lagrangian Method for Efficiently Handling Many Constraints	4525
<i>Joachim Giesen, Soeren Laue</i>	
Playing Card-Based RTS Games with Deep Reinforcement Learning	4540
<i>Tianyu Liu, Zijie Zheng, Hongchang Li, Kaigui Bian, Lingyang Song</i>	
FSM: A Fast Similarity Measurement for Gene Regulatory Networks via Genes' Influence Power	4547
<i>Zhongzhou Liu, Wenbin Hu</i>	
Pseudo Supervised Matrix Factorization in Discriminative Subspace	4554
<i>Jiaqi Ma, Yipeng Zhang, Lefei Zhang, Bo Du, Dapeng Tao</i>	
Representation Learning-Assisted Click-Through Rate Prediction	4561
<i>Wentao Ouyang, Xiuyu Zhang, Shukui Ren, Chao Qi, Zhaojie Liu, Yanlong Du</i>	
Dynamic Electronic Toll Collection via Multi-Agent Deep Reinforcement Learning with Edge-Based Graph Convolutional Networks	4568
<i>Wei Qiu, Haipeng Chen, Bo An</i>	
FireCast: Leveraging Deep Learning to Predict Wildfire Spread	4575
<i>David Radke, Anna Hessler, Dan Ellsworth</i>	

Faster Distributed Deep Net Training: Computation and Communication Decoupled Stochastic Gradient Descent	4582
<i>Shuheng Shen, Linli Xu, Jingchang Liu, Xianfeng Liang, Yifei Cheng</i>	
Randomized Adversarial Imitation Learning for Autonomous Driving	4590
<i>MyungJae Shin, Joongheon Kim</i>	
Scaling Fine-grained Modularity Clustering for Massive Graphs	4597
<i>Hiroaki Shiokawa, Toshiyuki Amagasa, Hiroyuki Kitagawa</i>	
Node Embedding over Temporal Graphs	4605
<i>Uriel Singer, Ido Guy, Kira Radinsky</i>	
Medical Concept Embedding with Multiple Ontological Representations	4613
<i>Lihong Song, Chin Wang Cheong, Kejing Yin, William K. Cheung, Benjamin C. M. Fung, Jonathan Poon</i>	
Learning Shared Vertex Representation in Heterogeneous Graphs with Convolutional Networks for Recommendation	4620
<i>Yanan Xu, Yanmin Zhu, Yanyan Shen, Jiadi Yu</i>	
Dual-Path in Dual-Path Network for Single Image Dehazing	4627
<i>Aiping Yang, Haixin Wang, Zhong Ji, Yanwei Pang, Ling Shao</i>	
Disparity-preserved Deep Cross-platform Association for Cross-platform Video Recommendation	4635
<i>Shengze Yu, Xin Wang, Wenwu Zhu, Peng Cui, Jingdong Wang</i>	

MULTIDISCIPLINARY TOPICS AND APPLICATIONS

Predicting Dominance in Multi-person Videos	4643
<i>Chongyang Bai, Maksim Bolonkin, Srijan Kumar, Jure Leskovec, Judee Burgoon, Norah Dunbar, V. S. Subrahmanian</i>	
Procedural Generation of Initial States of Sokoban	4651
<i>Dâmaris S. Bento, André G. Pereira, Levi H. S. Lelis</i>	
DeepInspect: A Black-box Trojan Detection and Mitigation Framework for Deep Neural Networks	4658
<i>Huili Chen, Cheng Fu, Jishen Zhao, Farinaz Koushanfar</i>	
VulSniper: Focus Your Attention to Shoot Fine-Grained Vulnerabilities	4665
<i>Xu Duan, Jingzheng Wu, Shouling Ji, Zhiqing Rui, Tianyue Luo, Mutian Yang, Yanjun Wu</i>	
Real-Time Adversarial Attacks	4672
<i>Yuan Gong, Boyang Li, Christian Poellabauer, Yiyu Shi</i>	
Explainable Fashion Recommendation: A Semantic Attribute Region Guided Approach	4681
<i>Min Hou, Le Wu, Enhong Chen, Zhi Li, Vincent W. Zheng, Qi Liu</i>	
Model-Agnostic Adversarial Detection by Random Perturbations	4689
<i>Bo Huang, Yi Wang, Wei Wang</i>	
Musical Composition Style Transfer via Disentangled Timbre Representations	4697
<i>Yun-Ning Hung, I-Tung Chiang, Yi-An Chen, Yi-Hsuan Yang</i>	
Multiple Policy Value Monte Carlo Tree Search	4704
<i>Li-Cheng Lan, Wei Li, Ting-Han Wei, I-Chen Wu</i>	
Robustra: Training Provable Robust Neural Networks over Reference Adversarial Space	4711
<i>Linyi Li, Zexuan Zhong, Bo Li, Tao Xie</i>	
Dilated Convolution with Dilated GRU for Music Source Separation	4718
<i>Jen-Yu Liu, Yi-Hsuan Yang</i>	
Locate-Then-Detect: Real-time Web Attack Detection via Attention-based Deep Neural Networks	4725
<i>Tianlong Liu, Yu Qi, Liang Shi, Jianan Yan</i>	
Data Poisoning against Differentially-Private Learners: Attacks and Defenses	4732
<i>Yuzhe Ma, Xiaojin Zhu, Justin Hsu</i>	
LogAnomaly: Unsupervised Detection of Sequential and Quantitative Anomalies in Unstructured Logs	4739
<i>Weibin Meng, Ying Liu, Yichen Zhu, Shenglin Zhang, Dan Pei, Yuqing Liu, Yihao Chen, Ruizhi Zhang, Shimin Tao, Pei Sun, Rong Zhou</i>	
Decidability of Model Checking Multi-Agent Systems with Regular Expressions against Epistemic HS Specifications	4746
<i>Jakub Michaliszyn, Piotr Witkowski</i>	
Heterogeneous Gaussian Mechanism: Preserving Differential Privacy in Deep Learning with Provable Robustness	4753
<i>NhatHai Phan, Minh N. Vu, Yang Liu, Ruoming Jin, Dejing Dou, Xintao Wu, My T. Thai</i>	
Demystifying the Combination of Dynamic Slicing and Spectrum-based Fault Localization	4760
<i>Sofia Reis, Rui Abreu, Marcelo d'Amorim</i>	

Equally-Guided Discriminative Hashing for Cross-modal Retrieval	4767
<i>Yufeng Shi, Xinge You, Feng Zheng, Shuo Wang, Qinmu Peng</i>	
A Privacy Preserving Collusion Secure DCOP Algorithm	4774
<i>Tamir Tassa, Tal Grinshpoun, Avishai Yanay</i>	
Two-Stage Generative Models of Simulating Training Data at The Voxel Level for Large-Scale Microscopy Bioimage Segmentation	4781
<i>Deli Wang, Ting Zhao, Nenggan Zheng, Zhefeng Gong</i>	
Lower Bound of Locally Differentially Private Sparse Covariance Matrix Estimation	4788
<i>Di Wang, Jinhui Xu</i>	
Principal Component Analysis in the Local Differential Privacy Model	4795
<i>Di Wang, Jinhui Xu</i>	
Binarized Collaborative Filtering with Distilling Graph Convolutional Network	4802
<i>Haoyu Wang, Defu Lian, Yong Ge</i>	
Novel Collaborative Filtering Recommender Friendly to Privacy Protection	4809
<i>Jun Wang, Qiang Tang, Afonso Arriaga, Peter Y. A. Ryan</i>	
Adversarial Examples for Graph Data: Deep Insights into Attack and Defense	4816
<i>Huijun Wu, Chen Wang, Yuriy Tyshetskiy, Andrew Docherty, Kai Lu, Liming Zhu</i>	
FABA: An Algorithm for Fast Aggregation against Byzantine Attacks in Distributed Neural Networks	4824
<i>Qi Xia, Zeyi Tao, Zijiang Hao, Qun Li</i>	
BAYHENN: Combining Bayesian Deep Learning and Homomorphic Encryption for Secure DNN Inference	4831
<i>Peichen Xie, Bingzhe Wu, Guangyu Sun</i>	
Toward Efficient Navigation of Massive-Scale Geo-Textual Streams	4838
<i>Chengcheng Yang, Lisi Chen, Shuo Shang, Fan Zhu, Li Liu, Ling Shao</i>	
Temporal Pyramid Pooling Convolutional Neural Network for Cover Song Identification	4846
<i>Zhesong Yu, Xiaoshuo Xu, Xiaouu Chen, Deshun Yang</i>	
Data Poisoning Attack against Knowledge Graph Embedding	4853
<i>Hengtong Zhang, Tianhang Zheng, Jing Gao, Chenglin Miao, Lu Su, Yaliang Li, Kui Ren</i>	
On Privacy Protection of Latent Dirichlet Allocation Model Training	4860
<i>Fangyuan Zhao, Xuebin Ren, Shusen Yang, Xinyu Yang</i>	
K-Core Maximization: An Edge Addition Approach	4867
<i>Zhongxin Zhou, Fan Zhang, Xuemin Lin, Wenjie Zhang, Chen Chen</i>	
Pivotal Relationship Identification: The K-Truss Minimization Problem	4874
<i>Weijie Zhu, Mengqi Zhang, Chen Chen, Xiaoyang Wang, Fan Zhang, Xuemin Lin</i>	

NATURAL LANGUAGE PROCESSING

Early Discovery of Emerging Entities in Microblogs	4882
<i>Satoshi Akasaki, Naoki Yoshinaga, Masashi Toyoda</i>	
Neural Program Induction for KBQA Without Gold Programs or Query Annotations	4890
<i>Ghulam Ahmed Ansari, Amrita Saha, Vishwajeet Kumar, Mohan Bhamhani, Karthik Sankaranarayanan, Soumen Chakrabarti</i>	
Medical Concept Representation Learning from Multi-source Data	4897
<i>Tian Bai, Brian L. Egleston, Richard Bleicher, Slobodan Vucetic</i>	
Multi-Domain Sentiment Classification Based on Domain-Aware Embedding and Attention	4904
<i>Yitao Cai, Xiaojun Wan</i>	
A Latent Variable Model for Learning Distributional Relation Vectors	4911
<i>Jose Camacho-Collados, Luis Espinosa-Anke, Shoaib Jameel, Steven Schockaert</i>	
Generating Multiple Diverse Responses with Multi-Mapping and Posterior Mapping Selection	4918
<i>Chaotao Chen, Jinhua Peng, Fan Wang, Jun Xu, Hua Wu</i>	
Sentiment-Controllable Chinese Poetry Generation	4925
<i>Huimin Chen, Xiaoyuan Yi, Maosong Sun, Wenhao Li, Cheng Yang, Zhipeng Guo</i>	
From Words to Sentences: A Progressive Learning Approach for Zero-resource Machine Translation with Visual Pivots	4932
<i>Shizhe Chen, Qin Jin, Jianlong Fu</i>	
Learning towards Abstractive Timeline Summarization	4939
<i>Xiuying Chen, Zhangming Chan, Shen Gao, Meng-Hsuan Yu, Dongyan Zhao, Rui Yan</i>	
Coreference Aware Representation Learning for Neural Named Entity Recognition	4946
<i>Zeyu Dai, Hongliang Fei, Ping Li</i>	

Learning Assistance from an Adversarial Critic for Multi-Outputs Prediction	4954
<i>Yue Deng, Yilin Shen, Hongxia Jin</i>	
End-to-End Multi-Perspective Matching for Entity Resolution	4961
<i>Cheng Fu, Xianpei Han, Le Sun, Bo Chen, Wei Zhang, Suhui Wu, Hao Kong</i>	
Difficulty Controllable Generation of Reading Comprehension Questions	4968
<i>Yifan Gao, Lidong Bing, Wang Chen, Michael Lyu, Irwin King</i>	
Modeling Source Syntax and Semantics for Neural AMR Parsing	4975
<i>DongLai Ge, Junhui Li, Muhua Zhu, Shoushan Li</i>	
CNN-Based Chinese NER with Lexicon Rethinking	4982
<i>Tao Gui, Ruotian Ma, Qi Zhang, Lujun Zhao, Yu-Gang Jiang, Xuanjing Huang</i>	
Dual Visual Attention Network for Visual Dialog	4989
<i>Dan Guo, Hui Wang, Meng Wang</i>	
AmazonQA: A Review-Based Question Answering Task	4996
<i>Mansi Gupta, Nitish Kulkarni, Raghuv eer Chanda, Anirudha Rayasam, Zachary C. Lipton</i>	
Answering Binary Causal Questions Through Large-Scale Text Mining: An Evaluation Using Cause-Effect Pairs from Human Experts	5003
<i>Oktie Hassanzadeh, Debarun Bhattacharjya, Mark Feblowitz, Kavitha Srinivas, Michael Perrone, Shirin Sohrabi, Michael Katz</i>	
GSN: A Graph-Structured Network for Multi-Party Dialogues	5010
<i>Wenpeng Hu, Zhangming Chan, Bing Liu, Dongyan Zhao, Jinwen Ma, Rui Yan</i>	
Leap-LSTM: Enhancing Long Short-Term Memory for Text Categorization	5017
<i>Ting Huang, Gehui Shen, Zhi-Hong Deng</i>	
Relation Extraction Using Supervision from Topic Knowledge of Relation Labels	5024
<i>Haiyun Jiang, Li Cui, Zhe Xu, Deqing Yang, Jindong Chen, Chenguang Li, Jingping Liu, Jiaqing Liang, Chao Wang, Yanghua Xiao, Wei Wang</i>	
Representation Learning with Weighted Inner Product for Universal Approximation of General Similarities	5031
<i>Geewook Kim, Akifumi Okuno, Kazuki Fukui, Hidetoshi Shimodaira</i>	
Incorporating Structural Information for Better Coreference Resolution	5039
<i>Fang Kong, Fu Jian</i>	
Knowledge Base Question Answering with Topic Units	5046
<i>Yunshi Lan, Shuohang Wang, Jing Jiang</i>	
Adversarial Transfer for Named Entity Boundary Detection with Pointer Networks	5053
<i>Jing Li, Deheng Ye, Shuo Shang</i>	
Towards Discriminative Representation Learning for Speech Emotion Recognition	5060
<i>Runnan Li, Zhiyong Wu, Jia Jia, Yaohua Bu, Sheng Zhao, Helen Meng</i>	
Self-attentive Biaffine Dependency Parsing	5067
<i>Ying Li, Zhenghua Li, Min Zhang, Rui Wang, Sheng Li, Luo Si</i>	
Reading selectively via Binary Input Gated Recurrent Unit	5074
<i>Zhe Li, Peisong Wang, Hanqing Lu, Jian Cheng</i>	
Learning to Select Knowledge for Response Generation in Dialog Systems	5081
<i>Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, Hua Wu</i>	
Deep Mask Memory Network with Semantic Dependency and Context Moment for Aspect Level Sentiment Classification	5088
<i>Peiqin Lin, Meng Yang, Jianhuang Lai</i>	
Exploring and Distilling Cross-Modal Information for Image Captioning	5095
<i>Fenglin Liu, Xuancheng Ren, Yuanxin Liu, Kai Lei, Xu Sun</i>	
Network Embedding with Dual Generation Tasks	5102
<i>Jie Liu, Na Li, Zhicheng He</i>	
Building Personalized Simulator for Interactive Search	5109
<i>Qianlong Liu, Baoliang Cui, Zhongyu Wei, Baolin Peng, Haikuan Huang, Hongbo Deng, Jianye Hao, Xuanjing Huang, Kam-Fai Wong</i>	

VOLUME 8

A Dual Reinforcement Learning Framework for Unsupervised Text Style Transfer	5116
<i>Fuli Luo, Peng Li, Jie Zhou, Pengcheng Yang, Baobao Chang, Xu Sun, Zhifang Sui</i>	
Unsupervised Neural Aspect Extraction with Sememes	5123
<i>Ling Luo, Xiang Ao, Yan Song, Jinyao Li, Xiaopeng Yang, Qing He, Dong Yu</i>	
Randomized Greedy Search for Structured Prediction: Amortized Inference and Learning	5130
<i>Chao Ma, F A Rezaur Rahman Chowdhury, Aryan Deshwal, Md Rakibul Islam, Janardhan Rao Doppa, Dan Roth</i>	

Aspect-Based Sentiment Classification with Attentive Neural Turing Machines	5139
<i>Qianren Mao, Jianxin Li, Senzhang Wang, Yuanning Zhang, Hao Peng, Min He, Lihong Wang</i>	
Learning Task-Specific Representation for Novel Words in Sequence Labeling	5146
<i>Minlong Peng, Qi Zhang, Xiaoyu Xing, Tao Gui, Jinlan Fu, Xuanjing Huang</i>	
Improving Cross-Domain Performance for Relation Extraction via Dependency Prediction and Information Flow Control	5153
<i>Amir Pouran Ben Veyseh, Thien Nguyen, Dejing Dou</i>	
Learn to Select via Hierarchical Gate Mechanism for Aspect-Based Sentiment Analysis	5160
<i>Xiangying Ran, Yuanyuan Pan, Wei Sun, Chongjun Wang</i>	
Aligning Learning Outcomes to Learning Resources: A Lexico-Semantic Spatial Approach	5168
<i>Swarnadeep Saha, Malolan Chetlur, Tejas Indulal Dhamecha, W M Gayathri K Wijayarathna, Red Mendoza, Paul Gagnon, Nabil Zary, Shantanu Godbole</i>	
A Deep Generative Model for Code Switched Text	5175
<i>Bidisha Samanta, Sharmila Reddy, Hussain Jagirdar, Niloy Ganguly, Soumen Chakrabarti</i>	
Knowledge Aware Semantic Concept Expansion for Image-Text Matching	5182
<i>Botian Shi, Lei Ji, Pan Lu, Zhendong Niu, Nan Duan</i>	
Exploiting Persona Information for Diverse Generation of Conversational Responses	5190
<i>Haoyu Song, Wei-Nan Zhang, Yiming Cui, Dong Wang, Ting Liu</i>	
Cold-Start Aware Deep Memory Network for Multi-Entity Aspect-Based Sentiment Analysis	5197
<i>Kaisong Song, Wei Gao, Lujun Zhao, Jun Lin, Changlong Sun, Xiaozhong Liu</i>	
GANs for Semi-Supervised Opinion Spam Detection	5204
<i>Gray Stanton, Athirai A. Irissappane</i>	
PROFET: Predicting the Risk of Firms from Event Transcripts	5211
<i>Christoph Kilian Theil, Samuel Broscheit, Heiner Stuckenschmidt</i>	
Unsupervised Embedding Enhancements of Knowledge Graphs using Textual Associations	5218
<i>Neil Veira, Brian Keng, Kanchana Padmanabhan, Andreas Veneris</i>	
Swell-and-Shrink: Decomposing Image Captioning by Transformation and Summarization	5226
<i>Hanzhang Wang, Hanli Wang, Kaisheng Xu</i>	
T-CVAE: Transformer-Based Conditioned Variational Autoencoder for Story Completion	5233
<i>Tianming Wang, Xiaojun Wan</i>	
Robust Embedding with Multi-Level Structures for Link Prediction	5240
<i>Zihan Wang, Zhaochun Ren, Chunyu He, Peng Zhang, Yue Hu</i>	
Revealing Semantic Structures of Texts: Multi-grained Framework for Automatic Mind-map Generation	5247
<i>Yang Wei, Honglei Guo, Jinmao Wei, Zhong Su</i>	
Correct-and-Memorize: Learning to Translate from Interactive Revisions	5255
<i>Rongxiang Weng, Hao Zhou, Shujian Huang, Lei Li, Yifan Xia, Jiajun Chen</i>	
Modeling Noisy Hierarchical Types in Fine-Grained Entity Typing: A Content-Based Weighting Approach	5264
<i>Junshuang Wu, Richong Zhang, Yongyi Mao, Hongyu Guo, Jinpeng Huai</i>	
Mask and Infill: Applying Masked Language Model for Sentiment Transfer	5271
<i>Xing Wu, Tao Zhang, Liangjun Zang, Jizhong Han, Songlin Hu</i>	
Relation-Aware Entity Alignment for Heterogeneous Knowledge Graphs	5278
<i>Yuting Wu, Xiao Liu, Yansong Feng, Zheng Wang, Rui Yan, Dongyan Zhao</i>	
RTHN: A RNN-Transformer Hierarchical Network for Emotion Cause Extraction	5285
<i>Rui Xia, Mengran Zhang, Zixiang Ding</i>	
Sharing Attention Weights for Fast Transformer	5292
<i>Tong Xiao, Yinqiao Li, Jingbo Zhu, Zhengtao Yu, Tongran Liu</i>	
A Goal-Driven Tree-Structured Neural Model for Math Word Problems	5299
<i>Zhipeng Xie, Shichao Sun</i>	
Dual-View Variational Autoencoders for Semi-Supervised Text Matching	5306
<i>Zhongbin Xie, Shuai Ma</i>	
Earlier Attention? Aspect-Aware LSTM for Aspect-Based Sentiment Analysis	5313
<i>Bowen Xing, Lejian Liao, Dandan Song, Jingang Wang, Fuzheng Zhang, Zhongyuan Wang, Heyan Huang</i>	
Polygon-Net: A General Framework for Jointly Boosting Multiple Unsupervised Neural Machine Translation Models	5320
<i>Chang Xu, Tao Qin, Gang Wang, Tie-Yan Liu</i>	
Neural Collective Entity Linking Based on Recurrent Random Walk Network Learning	5327
<i>Mengge Xue, Weiming Cai, Jinsong Su, Linfeng Song, Yubin Ge, Yubao Liu, Bin Wang</i>	
Robust Audio Adversarial Example for a Physical Attack	5334
<i>Hironu Yakura, Jun Sakuma</i>	

HorNet: A Hierarchical Offshoot Recurrent Network for Improving Person Re-ID via Image Captioning	5342
<i>Shiyang Yan, Jun Xu, Yuai Liu, Lin Xu</i>	
Knowledge-enhanced Hierarchical Attention for Community Question Answering with Multi-task and Adaptive Learning	5349
<i>Min Yang, Lei Chen, Xiaojun Chen, Qingyao Wu, Wei Zhou, Ying Shen</i>	
Knowledgeable Storyteller: A Commonsense-Driven Generative Model for Visual Storytelling	5356
<i>Pengcheng Yang, Fuli Luo, Peng Chen, Lei Li, Zhiyi Yin, Xiaodong He, Xu Sun</i>	
Triplet Enhanced AutoEncoder: Model-free Discriminative Network Embedding	5363
<i>Yao Yang, Haoran Chen, Junming Shao</i>	
Improving Multilingual Sentence Embedding using Bi-directional Dual Encoder with Additive Margin Softmax	5370
<i>Yinfei Yang, Gustavo Hernandez Abrego, Steve Yuan, Mandy Guo, Qinlan Shen, Daniel Cer, Yun-hsuan Sung, Brian Strope, Ray Kurzweil</i>	
Utilizing Non-Parallel Text for Style Transfer by Making Partial Comparisons	5379
<i>Di Yin, Shujian Huang, Xin-Yu Dai, Jiajun Chen</i>	
Graph-based Neural Sentence Ordering	5387
<i>Yongjing Yin, Linfeng Song, Jinsong Su, Jiali Zeng, Chulun Zhou, Jiebo Luo</i>	
Refining Word Representations by Manifold Learning	5394
<i>Chu Yonghe, Hongfei Lin, Liang Yang, Yufeng Diao, Shaowu Zhang, Fan Xiaochao</i>	
Beyond Word Attention: Using Segment Attention in Neural Relation Extraction	5401
<i>Bowen Yu, Zhenyu Zhang, Tingwen Liu, Bin Wang, Sujian Li, Quangang Li</i>	
Adapting BERT for Target-Oriented Multimodal Sentiment Classification	5408
<i>Jianfei Yu, Jing Jiang</i>	
Modeling both Context- and Speaker-Sensitive Dependence for Emotion Detection in Multi-speaker Conversations	5415
<i>Dong Zhang, Liangqing Wu, Changlong Sun, Shoushan Li, Qiaoming Zhu, Guodong Zhou</i>	
Extracting Entities and Events as a Single Task Using a Transition-Based Neural Model	5422
<i>Junchi Zhang, Yanxia Qin, Yue Zhang, Mengchi Liu, Donghong Ji</i>	
Multi-view Knowledge Graph Embedding for Entity Alignment	5429
<i>Qingheng Zhang, Zequn Sun, Wei Hu, Muhao Chen, Lingbing Guo, Yuzhong Qu</i>	
Quantum-Inspired Interactive Networks for Conversational Sentiment Analysis	5436
<i>Yazhou Zhang, Qiuchi Li, Dawei Song, Peng Zhang, Panpan Wang</i>	
A Document-grounded Matching Network for Response Selection in Retrieval-based Chatbots	5443
<i>Xueliang Zhao, Chongyang Tao, Wei Wu, Can Xu, Dongyan Zhao, Rui Yan</i>	
Recurrent Neural Network for Text Classification with Hierarchical Multiscale Dense Connections	5450
<i>Yi Zhao, Yanyan Shen, Junjie Yao</i>	
RLTM: An Efficient Neural IR Framework for Long Documents	5457
<i>Chen Zheng, Yu Sun, Shengxian Wan, Dianhai Yu</i>	
Dynamically Route Hierarchical Structure Representation to Attentive Capsule for Text Classification	5464
<i>Wanshan Zheng, Zibin Zheng, Hai Wan, Chuan Chen</i>	
Sequence Generation: From Both Sides to the Middle	5471
<i>Long Zhou, Jiajun Zhang, Chengqing Zong, Heng Yu</i>	
Getting in Shape: Word Embedding SubSpaces	5478
<i>Tianyuan Zhou, João Sedoc, Jordan Rodu</i>	
A Span-based Joint Model for Opinion Target Extraction and Target Sentiment Classification	5485
<i>Yan Zhou, Longtao Huang, Tao Guo, Jizhong Han, Songlin Hu</i>	

PLANNING AND SCHEDULING

Earliest-Completion Scheduling of Contract Algorithms with End Guarantees	5493
<i>Spyros Angelopoulos, Shendan Jin</i>	
Finding Optimal Solutions in HTN Planning - A SAT-based Approach	5500
<i>Gregor Behnke, Daniel Höller, Susanne Biundo</i>	
Faster Dynamic Controllability Checking in Temporal Networks with Integer Bounds	5509
<i>Nikhil Bhargava, Brian C. Williams</i>	
Regular Decision Processes: A Model for Non-Markovian Domains	5516
<i>Ronen I. Brafman, Giuseppe De Giacomo</i>	
Strong Fully Observable Non-Deterministic Planning with LTL and LTLf Goals	5523
<i>Alberto Camacho, Sheila A. McIlraith</i>	

Counterexample-Guided Strategy Improvement for POMDPs Using Recurrent Neural Networks	5532
<i>Steven Carr, Nils Jansen, Ralf Wimmer, Alexandru Serban, Bernd Becker, Ufuk Topcu</i>	
Influence of State-Variable Constraints on Partially Observable Monte Carlo Planning	5540
<i>Alberto Castellini, Georgios Chalkiadakis, Alessandro Farinelli</i>	
Online Probabilistic Goal Recognition over Nominal Models	5547
<i>Ramon Fraga Pereira, Mor Vered, Felipe Meneguzzi, Miquel Ramírez</i>	
Generalized Potential Heuristics for Classical Planning	5554
<i>Guillem Francès, Augusto B. Corrêa, Cedric Geissmann, Florian Pommerening</i>	
Subgoal-Based Temporal Abstraction in Monte-Carlo Tree Search	5562
<i>Thomas Gabor, Jan Peter, Thomy Phan, Christian Meyer, Claudia Linnhoff-Popien</i>	
Fair Online Allocation of Perishable Goods and its Application to Electric Vehicle Charging	5569
<i>Enrico H. Gerding, Alvaro Perez-Diaz, Haris Aziz, Serge Gaspers, Antonia Marcu, Nicholas Mattei, Toby Walsh</i>	
Dynamic logic of parallel propositional assignments and its applications to planning	5576
<i>Andreas Herzig, Frédéric Maris, Julien Vianey</i>	
Approximability of Constant-horizon Constrained POMDP	5583
<i>Majid Khonji, Ashkan Jasour, Brian Williams</i>	
Bayesian Inference of Linear Temporal Logic Specifications for Contrastive Explanations	5591
<i>Joseph Kim, Christian Muise, Ankit Shah, Shubham Agarwal, Julie Shah</i>	
Partitioning Techniques in LTLf Synthesis	5599
<i>Lucas Martinelli Tabajara, Moshe Y. Vardi</i>	
Adaptive Thompson Sampling Stacks for Memory Bounded Open-Loop Planning	5607
<i>Thomy Phan, Thomas Gabor, Robert Müller, Christoph Roch, Claudia Linnhoff-Popien</i>	
A Novel Distribution-Embedded Neural Network for Sensor-Based Activity Recognition	5614
<i>Hangwei Qian, Sinno Jialin Pan, Bingshui Da, Chunyan Miao</i>	
Pattern Selection for Optimal Classical Planning with Saturated Cost Partitioning	5621
<i>Jendrik Seipp</i>	
Scheduling Jobs with Stochastic Processing Time on Parallel Identical Machines	5628
<i>Richard Stec, Antonin Novak, Premysl Sucha, Zdenek Hanzalek</i>	
On Computational Complexity of Pickup-and-Delivery Problems with Precedence Constraints or Time Windows	5635
<i>Xing Tan, Jimmy Xiangji Huang</i>	
Merge-and-Shrink Task Reformulation for Classical Planning	5644
<i>Álvaro Torralba, Silvan Sievers</i>	
Steady-State Policy Synthesis for Verifiable Control	5653
<i>Alvaro Velasquez</i>	

ROBOTICS

Energy-Efficient Slithering Gait Exploration for a Snake-Like Robot Based on Reinforcement Learning	5663
<i>Zhenshan Bing, Christian Lemke, Zhuangyi Jiang, Kai Huang, Alois Knoll</i>	
The Parameterized Complexity of Motion Planning for Snake-Like Robots	5670
<i>Siddharth Gupta, Guy Sa'ar, Meirav Zehavi</i>	
Unsupervised Learning of Monocular Depth and Ego-Motion using Conditional PatchGANs	5677
<i>Madhu Vankadari, Swagat Kumar, Anima Majumder, Kaushik Das</i>	
Region Deformer Networks for Unsupervised Depth Estimation from Unconstrained Monocular Videos	5685
<i>Haofei Xu, Jianmin Zheng, Jianfei Cai, Juyong Zhang</i>	

UNCERTAINTY IN AI

Statistical Guarantees for the Robustness of Bayesian Neural Networks	5693
<i>Luca Cardelli, Marta Kwiatkowska, Luca Laurenti, Nicola Paoletti, Andrea Patane, Matthew Wicker</i>	
Lifted Message Passing for Hybrid Probabilistic Inference	5701
<i>Yuqiao Chen, Nicholas Ruoizzi, Sriraam Natarajan</i>	
Bayesian Parameter Estimation for Nonlinear Dynamics Using Sensitivity Analysis	5708
<i>Yi Chou, Sriram Sankaranarayanan</i>	
Thompson Sampling on Symmetric Alpha-Stable Bandits	5715
<i>Abhimanyu Dubey, Alex `Sandy' Pentland</i>	

On Constrained Open-World Probabilistic Databases	5722
<i>Tal Friedman, Guy Van den Broeck</i>	
An End-to-End Community Detection Model: Integrating LDA into Markov Random Field via Factor Graph	5730
<i>Dongxiao He, Wenze Song, Di Jin, Zhiyong Feng, Yuxiao Huang</i>	
Exact Bernoulli Scan Statistics using Binary Decision Diagrams	5737
<i>Masakazu Ishihata, Takanori Maehara</i>	
Hyper-parameter Tuning under a Budget Constraint	5744
<i>Zhiyun Lu, Liyu Chen, Chao-Kai Chiang, Fei Sha</i>	
Cutset Bayesian Networks: A New Representation for Learning Rao-Blackwellised Graphical Models	5751
<i>Tahrima Rahman, Shasha Jin, Vibhav Gogate</i>	
Ranked Programming	5758
<i>Tjitze Rienstra</i>	
ISLF: Interest Shift and Latent Factors Combination Model for Session-based Recommendation	5765
<i>Jing Song, Hong Shen, Zijing Ou, Junyi Zhang, Teng Xiao, Shangsong Liang</i>	
DiffChaser: Detecting Disagreements for Deep Neural Networks	5772
<i>Xiaofei Xie, Lei Ma, Haijun Wang, Yuekang Li, Yang Liu, Xiaohong Li</i>	

AI FOR IMPROVING HUMAN WELL-BEING

SparseSense: Human Activity Recognition from Highly Sparse Sensor Data-streams Using Set-based Neural Networks	5780
<i>Alireza Abedin, S. Hamid RezaTofighi, Qinfeng Shi, Damith C. Ranasinghe</i>	
Governance by Glass-Box: Implementing Transparent Moral Bounds for AI Behaviour	5787
<i>Andrea Aler Tubella, Andreas Theodorou, Frank Dignum, Virginia Dignum</i>	
Decision Making for Improving Maritime Traffic Safety Using Constraint Programming	5794
<i>Saumya Bhatnagar, Akshat Kumar, Hoong Chuin Lau</i>	
Evaluating the Interpretability of the Knowledge Compilation Map: Communicating Logical Statements Effectively	5801
<i>Serena Booth, Christian MuiSe, Julie Shah</i>	
AI-powered Posture Training: Application of Machine Learning in Sitting Posture Recognition Using the LifeChair Smart Cushion	5808
<i>Katia Bourahmoune, Toshiyuki Amagasa</i>	
Improving Law Enforcement Daily Deployment Through Machine Learning-Informed Optimization under Uncertainty	5815
<i>Jonathan Chase, Duc Thien Nguyen, Haiyang Sun, Hoong Chuin Lau</i>	
Risk Assessment for Networked-guarantee Loans Using High-order Graph Attention Representation	5822
<i>Dawei Cheng, Yi Tu, Zhenwei Ma, Zhibin Niu, Liqing Zhang</i>	
PI-Bully: Personalized Cyberbullying Detection with Peer Influence	5829
<i>Lu Cheng, Jundong Li, Yasin Silva, Deborah Hall, Huan Liu</i>	
The Price of Local Fairness in Multistage Selection	5836
<i>Vitalii Emelianov, George Arvanitakis, Nicolas Gast, Krishna Gummadi, Patrick Loiseau</i>	
Enhancing Stock Movement Prediction with Adversarial Training	5843
<i>Fuli Feng, Huimin Chen, Xiangnan He, Ji Ding, Maosong Sun, Tat-Seng Chua</i>	

VOLUME 9

Safe Contextual Bayesian Optimization for Sustainable Room Temperature PID Control Tuning	5850
<i>Marcello Fiducioso, Sebastian Curi, Benedikt Schumacher, Markus Gwerder, Andreas Krause</i>	
DDL: Deep Dictionary Learning for Predictive Phenotyping	5857
<i>Tianfan Fu, Trong Nghia Hoang, Cao Xiao, Jimeng Sun</i>	
Improving Customer Satisfaction in Bike Sharing Systems through Dynamic Repositioning	5864
<i>Supriyo Ghosh, Jing Yu Koh, Patrick Jaillet</i>	
mdfa: Multi-Differential Fairness Auditor for Black Box Classifiers	5871
<i>Xavier Gitiaux, Huzefa Rangwala</i>	
Counterfactual Regression with Importance Sampling Weights	5880
<i>Negar Hassanpour, Russell Greiner</i>	
MINA: Multilevel Knowledge-Guided Attention for Modeling Electrocardiography Signals	5888
<i>Shenda Hong, Cao Xiao, Tengfei Ma, Hongyan Li, Jimeng Sun</i>	

RDPD: Rich Data Helps Poor Data via Imitation	5895
<i>Shenda Hong, Cao Xiao, Trong Nghia Hoang, Tengfei Ma, Hongyan Li, Jimeng Sun</i>	
Systematic Conservation Planning for Sustainable Land-use Policies: A Constrained Partitioning Approach to Reserve Selection and Design.	5902
<i>Dimitri Justeau-Allaire, Philippe Vismara, Philippe Birnbaum, Xavier Lorca</i>	
Truly Batch Apprenticeship Learning with Deep Successor Features	5909
<i>Donghun Lee, Srivatsan Srinivasan, Finale Doshi-Velez</i>	
Scribble-to-Painting Transformation with Multi-Task Generative Adversarial Networks	5916
<i>Jinning Li, Yexiang Xue</i>	
Diversity-Inducing Policy Gradient: Using Maximum Mean Discrepancy to Find a Set of Diverse Policies	5923
<i>Muhammad Masood, Finale Doshi-Velez</i>	
Kitchenette: Predicting and Ranking Food Ingredient Pairings using Siamese Neural Network	5930
<i>Donghyeon Park, Keonwoo Kim, Yonggyu Park, Jungwoon Shin, Jaewoo Kang</i>	
MNN: Multimodal Attentional Neural Networks for Diagnosis Prediction	5937
<i>Zhi Qiao, Xian Wu, Shen Ge, Wei Fan</i>	
Global Robustness Evaluation of Deep Neural Networks with Provable Guarantees for the Hamming Distance	5944
<i>Wenjie Ruan, Min Wu, Youcheng Sun, Xiaowei Huang, Daniel Kroening, Marta Kwiatkowska</i>	
Pre-training of Graph Augmented Transformers for Medication Recommendation	5953
<i>Junyuan Shang, Tengfei Ma, Cao Xiao, Jimeng Sun</i>	
Three-quarter Sibling Regression for Denoising Observational Data	5960
<i>Shiv Shankar, Daniel Sheldon, Tao Sun, John Pickering, Thomas G. Dietterich</i>	
Daytime Sleepiness Level Prediction Using Respiratory Information	5967
<i>Kazuhiko Shinoda, Masahiko Yoshii, Hayato Yamaguchi, Hirotaka Kaji</i>	
Simultaneous Prediction Intervals for Patient-Specific Survival Curves	5975
<i>Samuel Sokota, Ryan D'Orazio, Khurram Javed, Humza Haider, Russell Greiner</i>	
Controllable Neural Story Plot Generation via Reward Shaping	5982
<i>Pradyumna Tambwekar, Murtaza Dhuliawala, Lara J. Martin, Animesh Mehta, Brent Harrison, Mark O. Riedl</i>	
Bidirectional Active Learning with Gold-Instance-Based Human Training	5989
<i>Feilong Tang</i>	
Group-Fairness in Influence Maximization	5997
<i>Alan Tsang, Bryan Wilder, Eric Rice, Milind Tambe, Yair Zick</i>	
Failure-Scenario Maker for Rule-Based Agent using Multi-agent Adversarial Reinforcement Learning and its Application to Autonomous Driving	6006
<i>Akifumi Wachi</i>	
Protecting Neural Networks with Hierarchical Random Switching: Towards Better Robustness-Accuracy Trade-off for Stochastic Defenses	6013
<i>Xiao Wang, Siyue Wang, Pin-Yu Chen, Yanzhi Wang, Brian Kulis, Xue Lin, Sang Chin</i>	
Who Should Pay the Cost: A Game-theoretic Model for Government Subsidized Investments to Improve National Cybersecurity	6020
<i>Xinrun Wang, Bo An, Hau Chan</i>	
Automatic Grassland Degradation Estimation Using Deep Learning	6028
<i>Xiyu Yan, Yong Jiang, Shuai Chen, Zihao He, Chunmei Li, Shu-Tao Xia, Tao Dai, Shuo Dong, Feng Zheng</i>	
Balanced Ranking with Diversity Constraints	6035
<i>Ke Yang, Vasilis Gkatzelis, Julia Stoyanovich</i>	
A Decomposition Approach for Urban Anomaly Detection Across Spatiotemporal Data	6043
<i>Mingyang Zhang, Tong Li, Hongzhi Shi, Yong Li, Pan Hui</i>	
Learning Interpretable Relational Structures of Hinge-loss Markov Random Fields	6050
<i>Yue Zhang, Arti Ramesh</i>	
K-margin-based Residual-Convolution-Recurrent Neural Network for Atrial Fibrillation Detection	6057
<i>Yuxi Zhou, Shenda Hong, Junyuan Shang, Meng Wu, Qingyun Wang, Hongyan Li, Junqing Xie</i>	

UNDERSTANDING INTELLIGENCE AND HUMAN-LEVEL AI IN THE NEW MACHINE LEARNING ERA

LTL and Beyond: Formal Languages for Reward Function Specification in Reinforcement Learning	6065
<i>Alberto Camacho, Rodrigo Toro Icarte, Toryn Q. Klassen, Richard Valenzano, Sheila A. McIlraith</i>	
Playgol: Learning Programs Through Play	6074
<i>Andrew Cropper</i>	

Learning Relational Representations with Auto-encoding Logic Programs	6081
<i>Sebastijan Dumancic, Tias Guns, Wannas Meert, Hendrik Blockeel</i>	
A Comparative Study of Distributional and Symbolic Paradigms for Relational Learning	6088
<i>Sebastijan Dumancic, Alberto Garcia-Duran, Mathias Niepert</i>	
Learning Hierarchical Symbolic Representations to Support Interactive Task Learning and Knowledge Transfer	6095
<i>James R. Kirk, John E. Laird</i>	
EL Embeddings: Geometric Construction of Models for the Description Logic EL^{++}	6103
<i>Maxat Kulmanov, Wang Liu-Wei, Yuan Yan, Robert Hoehndorf</i>	
How Well Do Machines Perform on IQ tests: a Comparison Study on a Large-Scale Dataset	6110
<i>Yusen Liu, Fangyuan He, Haodi Zhang, Guozheng Rao, Zhiyong Feng, Yi Zhou</i>	
Synthesizing Datalog Programs using Numerical Relaxation	6117
<i>Xujie Si, Mukund Raghothaman, Kihong Heo, Mayur Naik</i>	

BEST SISTER CONFERENCES

A Refined Understanding of Cost-optimal Planning with Polytree Causal Graphs	6126
<i>Christer Bäckström, Peter Jonsson, Sebastian Ordyniak</i>	
Closed-World Semantics for Conjunctive Queries with Negation over ELH-bottom Ontologies	6131
<i>Stefan Borgwardt, Walter Forkel</i>	
Quality Control Attack Schemes in Crowdsourcing	6136
<i>Alessandro Checco, Jo Bates, Gianluca Demartini</i>	
Do We Need Many-valued Logics for Incomplete Information?	6141
<i>Marco Console, Paolo Guagliardo, Leonid Libkin</i>	
Addressing Age-Related Bias in Sentiment Analysis	6146
<i>Mark Díaz, Isaac Johnson, Amanda Lazar, Anne Marie Piper, Darren Gergle</i>	
Sharpness of the Satisfiability Threshold for Non-Uniform Random k-SAT	6151
<i>Tobias Friedrich, Ralf Rothenberger</i>	
A Dual Approach to Verify and Train Deep Networks	6156
<i>Sven Gowal, Krishnamurthy Dvijotham, Robert Stanforth, Timothy Mann, Pushmeet Kohli</i>	
The Provable Virtue of Laziness in Motion Planning	6161
<i>Nika Haghtalab, Simon Mackenzie, Ariel D. Procaccia, Oren Salzman, Siddhartha Srinivasa</i>	
Clause Learning and New Bounds for Graph Coloring	6166
<i>Emmanuel Hebrard, George Katsirelos</i>	
On Guiding Search in HTN Planning with Classical Planning Heuristics	6171
<i>Daniel Höller, Pascal Bercher, Gregor Behnke, Susanne Biundo</i>	
The Power of Context in Networks: Ideal Point Models with Social Interactions	6176
<i>Mohammad T. Irfan, Tucker Gordon</i>	
On Causal Identification under Markov Equivalence	6181
<i>Amin Jaber, Jiji Zhang, Elias Bareinboim</i>	
Meta-Interpretive Learning Using HEX-Programs	6186
<i>Tobias Kaminski, Thomas Eiter, Katsumi Inoue</i>	
A Walkthrough for the Principle of Logit Separation	6191
<i>Gil Keren, Sivan Sabato, Björn Schuller</i>	
Delayed Impact of Fair Machine Learning	6196
<i>Lydia T. Liu, Sarah Dean, Esther Rolf, Max Simchowitz, Moritz Hardt</i>	
Impact of Consuming Suggested Items on the Assessment of Recommendations in User Studies on Recommender Systems	6201
<i>Benedikt Loepp, Tim Donkers, Timm Kleemann, Jürgen Ziegler</i>	
Discovering Reliable Dependencies from Data: Hardness and Improved Algorithms	6206
<i>Panagiotis Mandros, Mario Boley, Jilles Vreeken</i>	
Not All FPRASs are Equal: Demystifying FPRASs for DNF-Counting (Extended Abstract)	6211
<i>Kuldeep S. Meel, Aditya A. Shrotri, Moshe Y. Vardi</i>	
Constraint Games for Stable and Optimal Allocation of Demands in SDN	6216
<i>Anthony Palmieri, Arnaud Lallouet, Luc Pons</i>	
Optimally Efficient Bidirectional Search	6221
<i>Eshed Shaham, Ariel Felner, Nathan R. Sturtevant, Jeffrey S. Rosenschein</i>	
Trust Dynamics and Transfer across Human-Robot Interaction Tasks: Bayesian and Neural Computational Models	6226
<i>Harold Soh, Shu Pan, Min Chen, David Hsu</i>	

Differentiable Physics and Stable Modes for Tool-Use and Manipulation Planning - Extended Abstract	6231
<i>Marc Toussaint, Kelsey R. Allen, Kevin A. Smith, Joshua B. Tenenbaum</i>	
Causal Embeddings for Recommendation: An Extended Abstract	6236
<i>Flavian Vasile, Stephen Bonner</i>	
Taskonomy: Disentangling Task Transfer Learning	6241
<i>Amir Zamir, Alexander Sax, William Shen, Leonidas Guibas, Jitendra Malik, Silvio Savarese</i>	
Adversarial Attacks on Neural Networks for Graph Data	6246
<i>Daniel Zügner, Amir Akbarnejad, Stephan Günnemann</i>	

SURVEY TRACK

Integrating Knowledge and Reasoning in Image Understanding	6252
<i>Somak Aditya, Yezhou Yang, Chitta Baral</i>	
A Replication Study of Semantics in Argumentation	6260
<i>Leila Amgoud</i>	
A Survey on Hierarchical Planning – One Abstract Idea, Many Concrete Realizations	6267
<i>Pascal Bercher, Ron Alford, Daniel Höller</i>	
Counterfactuals in Explainable Artificial Intelligence (XAI): Evidence from Human Reasoning	6276
<i>Ruth M. J. Byrne</i>	
Deep Learning for Video Captioning: A Review	6283
<i>Shaoxiang Chen, Ting Yao, Yu-Gang Jiang</i>	
Learning and Inference for Structured Prediction: A Unifying Perspective	6291
<i>Aryan Deshwal, Janardhan Rao Doppa, Dan Roth</i>	
Automated Essay Scoring: A Survey of the State of the Art	6300
<i>Zixuan Ke, Vincent Ng</i>	
A Survey of Reinforcement Learning Informed by Natural Language	6309
<i>Jelena Luketina, Nantas Nardelli, Gregory Farquhar, Jakob Foerster, Jacob Andreas, Edward Grefenstette, Shimon Whiteson, Tim Rocktäschel</i>	
Social Media-based User Embedding: A Literature Review	6318
<i>Shimei Pan, Tao Ding</i>	
Recent Advances in Imitation Learning from Observation	6325
<i>Faraz Torabi, Garrett Warnell, Peter Stone</i>	
Sequential Recommender Systems: Challenges, Progress and Prospects	6332
<i>Shoujin Wang, Liang Hu, Yan Wang, Longbing Cao, Quan Z. Sheng, Mehmet Orgun</i>	
Leveraging Human Guidance for Deep Reinforcement Learning Tasks	6339
<i>Ruohan Zhang, Faraz Torabi, Lin Guan, Dana H. Ballard, Peter Stone</i>	

JOURNAL TRACK

Learning in the Machine: Random Backpropagation and the Deep Learning Channel (Extended Abstract)	6348
<i>Pierre Baldi, Peter Sadowski, Zhiqin Lu</i>	
Complexity Bounds for the Controllability of Temporal Networks with Conditions, Disjunctions, and Uncertainty (Extended Abstract)	6353
<i>Nikhil Bhargava, Brian C. Williams</i>	
A Core Method for the Weak Completion Semantics with Skeptical Abduction (Extended Abstract)	6358
<i>Emmanuelle-Anna Dietz Saldanha, Steffen Hölldobler, Carroline Dewi Puspa Kencana Ramli, Luis Palacios Medinacelli</i>	
Complexity of Fundamental Problems in Probabilistic Abstract Argumentation: Beyond Independence (Extended Abstract)	6362
<i>Bettina Fazzinga, Sergio Flesca, Filippo Furfaro</i>	
Shielded Base Contraction (Extended Abstract)	6367
<i>Marco Garapa, Eduardo Fermé, Maurício D. L. Reis</i>	
Implicitly Coordinated Multi-Agent Path Finding under Destination Uncertainty: Success Guarantees and Computational Complexity (Extended Abstract)	6372
<i>Bernhard Nebel, Thomas Bolander, Thorsten Engesser, Robert Mattmüller</i>	
Teaching AI Agents Ethical Values Using Reinforcement Learning and Policy Orchestration	6377
<i>Ritesh Noothigattu, Djallel Bouneffouf, Nicholas Mattei, Rachita Chandra, Piyush Madan, Kush R. Varshney, Murray Campbell, Moninder Singh, Francesca Rossi</i>	

On the Responsibility for Undecisiveness in Preferred and Stable Labellings in Abstract Argumentation (Extended Abstract)	6382
<i>Claudia Schulz, Francesca Toni</i>	

EARLY CAREER

The Quest For "Always-On" Autonomous Mobile Robots	6388
<i>Joydeep Biswas</i>	
Integrating Learning with Game Theory for Societal Challenges	6393
<i>Fei Fang</i>	
Multiagent Decision Making and Learning in Urban Environments	6398
<i>Akshat Kumar</i>	
From Data to Knowledge Engineering for Cybersecurity	6403
<i>Gerardo I. Simari</i>	
AI Planning for Enterprise: Putting Theory Into Practice	6408
<i>Shirin Sohrabi</i>	
Domain-Dependent and Domain-Independent Problem Solving Techniques	6411
<i>Roni Stern</i>	
What Does the Evidence Say? Models to Help Make Sense of the Biomedical Literature	6416
<i>Byron C. Wallace</i>	

DOCTORAL CONSORTIUM

Conditional Preference Network with Constraints and Uncertainty	6422
<i>Sultan Ahmed</i>	
Constraint Solving and Optimization Using Evolutionary Techniques	6424
<i>Mahdi Bidar, Malek Mouhoub</i>	
Visionary Security: Using Uncertain Real-Time Information in Signaling Games	6426
<i>Elizabeth Bondi</i>	
AI in Recruiting. Multi-agent Systems Architecture for Ethical and Legal Auditing	6428
<i>María del Carmen Fernández Martínez, Alberto Fernández</i>	
A Similarity Measurement Method Based on Graph Kernel for Disconnected Graphs	6430
<i>Jun Gao, Jianliang Gao</i>	
Finding Justifications by Approximating Core for Large-scale Ontologies	6432
<i>Mengyu Gao, Yuxin Ye, Dantong Ouyang, Bin Wang</i>	
Teaching Robots to Interact with Humans in a Smart Environment	6434
<i>Shivam Goel</i>	
Vision beyond Pixels: Visual Reasoning via Blocksworld Abstractions	6436
<i>Tejas Gokhale</i>	
A Unified Mathematical Approach for Foraging and Construction Systems in a 1,000,000 Robot Swarm	6438
<i>John Harwell</i>	
Can Meta-Interpretive Learning outperform Deep Reinforcement Learning of Evaluable Game strategies?	6440
<i>Céline Hocquette</i>	
Unsupervised Multi-view Learning	6442
<i>Ling Huang</i>	
Intelligent Agent for Assessing and Guiding Rehabilitation Exercises	6444
<i>Min Hun Lee</i>	
Optimizing Interactive Systems with Data-Driven Objectives	6446
<i>Ziming Li</i>	
Split Q Learning: Reinforcement Learning with Two-Stream Rewards	6448
<i>Baihan Lin, Djallel Bouneffouf, Guillermo Cecchi</i>	
Event Prediction in Complex Social Graphs using One-Dimensional Convolutional Neural Network	6450
<i>Bonaventure Molokwu</i>	
Towards Architecture-Agnostic Neural Transfer: a Knowledge-Enhanced Approach	6452
<i>Seán Quinn, Alessandra Mileo</i>	
Technical, Hard and Explainable Question Answering (THE-QA)	6454
<i>Shailaja Sampat</i>	

Preference Elicitation and Explanation in Iterative Planning	6456
<i>Lindsay Sanneman</i>	
Intelligent Querying in Camera Networks for Efficient Target Tracking	6458
<i>Anil Sharma</i>	
Safe and Sample-Efficient Reinforcement Learning Algorithms for Factored Environments	6460
<i>Thiago D. Simão</i>	
Deep Learning with Relational Logic Representations	6462
<i>Gustav Sourek</i>	
Matching with Constraints	6464
<i>Zhaohong Sun</i>	
Cautious Rule-Based Collective Inference	6466
<i>Martin Svatos</i>	
The Design of Human Oversight in Autonomous Weapon Systems	6468
<i>Ilse Verdieen</i>	
A Compliance Checking Framework for DNN Models	6470
<i>Sunny Verma, Chen Wang, Liming Zhu, Wei Liu</i>	
Adversarial Machine Learning with Double Oracle	6472
<i>Kai Wang</i>	
AI at the Margins: Data, Decisions, and Inclusive Social Impact	6474
<i>Bryan Wilder</i>	
Self-Organizing Incremental Neural Networks for Continual Learning	6476
<i>Chayut Wiwatcharakoses, Daniel Berrar</i>	
Evolutionary Learning of Existential Rules	6478
<i>Lianlong Wu</i>	
Entity Alignment for Cross-lingual Knowledge Graph with Graph Convolutional Networks	6480
<i>Fan Xiong, Jianliang Gao</i>	
Personal Context Recognition via Skeptical Learning	6482
<i>Wanyi Zhang</i>	

DEMOS

OpenMarkov, an Open-Source Tool for Probabilistic Graphical Models	6485
<i>Manuel Arias, Jorge Pérez-Martín, Manuel Luque, Francisco J. Díez</i>	
Explainable Deep Neural Networks for Multivariate Time Series Predictions	6488
<i>Roy Assaf, Anika Schumann</i>	
Multi-Agent Path Finding on Ozobots	6491
<i>Roman Barták, Ivan Krasícenko, Jirí Švancara</i>	
Hintikka's World: Scalable Higher-order Knowledge	6494
<i>Tristan Charrier, Sébastien Gamblin, Alexandre Niveau, François Schwarzentruber</i>	
AntProphet: an Intention Mining System behind Alipay's Intelligent Customer Service Bot	6497
<i>Cen Chen, Xiaolu Zhang, Sheng Ju, Chilin Fu, Caizhi Tang, Jun Zhou, Xiaolong Li</i>	
Crowd View: Converting Investors' Opinions into Indicators	6500
<i>Chung-Chi Chen, Hen-Hsen Huang, Hsin-Hsi Chen</i>	
VEST: A System for Vulnerability Exploit Scoring & Timing	6503
<i>Haipeng Chen, Jing Liu, Rui Liu, Noseong Park, V. S. Subrahmanian</i>	
Demonstration of PerformanceNet: A Convolutional Neural Network Model for Score-to-Audio Music Generation	6506
<i>Yu-Hua Chen, Bryan Wang, Yi-Hsuan Yang</i>	
InterSpot: Interactive Spammer Detection in Social Media	6509
<i>Kaize Ding, Jundong Li, Shivam Dhar, Shreyash Devan, Huan Liu</i>	
Embodied Conversational AI Agents in a Multi-modal Multi-agent Competitive Dialogue	6512
<i>Rahul R. Divekar, Xiangyang Mou, Lisha Chen, Maíra Gatti de Baysar, Melina Alberio Guerra, Hui Su</i>	
A Mobile Application for Sound Event Detection	6515
<i>Yingwei Fu, Kele Xu, Haibo Mi, Huaimin Wang, Dezhi Wang, Boqing Zhu</i>	
The Open Vault Challenge - Learning How to Build Calibration-Free Interactive Systems by Cracking the Code of a Vault	6518
<i>Jonathan Grizou</i>	
Agent-based Decision Support for Pain Management in Primary Care Settings	6521
<i>Xu Guo, Han Yu, Chunyan Miao, Yiqiang Chen</i>	
DISPUTool -- A tool for the Argumentative Analysis of Political Debates	6524
<i>Shohreh Haddadan, Elena Cabrio, Serena Villata</i>	

Contextual Typeahead Sticker Suggestions on Hike Messenger	6527
<i>Mohamed Hanoosh, Abhishek Laddha, Debdoot Mukherjee</i>	
The pywmi Framework and Toolbox for Probabilistic Inference using Weighted Model Integration	6530
<i>Samuel Kolb, Paolo Morettin, Pedro Zuidberg Dos Martires, Francesco Somnavilla, Andrea Passerini, Roberto Sebastiani, Luc De Raedt</i>	
ERICA and WikiTalk	6533
<i>Divesh Lala, Graham Wilcock, Kristiina Jokinen, Tatsuya Kawahara</i>	
Design and Implementation of a Disambiguity Framework for Smart Voice Controlled Devices	6536
<i>Kehua Lei, Tianyi Ma, Jia Jia, Cunjun Zhang, Zhihan Yang</i>	
Neural Discourse Segmentation	6539
<i>Jing Li</i>	
CoTrRank: Trust Evaluation of Users and Tweets	6542
<i>Peiyao Li, Weiliang Zhao, Jian Yang, Jia Wu</i>	
Mappa Mundi: An Interactive Artistic Mind Map Generator with Artificial Imagination	6545
<i>Ruixue Liu, Baoyang Chen, Meng Chen, Youzheng Wu, Zhijie Qiu, Xiaodong He</i>	
CRSRL: Customer Routing System Using Reinforcement Learning	6548
<i>Chong Long, Zining Liu, Xiaolu Lu, Zehong Hu, Yafang Wang</i>	
ACTA A Tool for Argumentative Clinical Trial Analysis	6551
<i>Tobias Mayer, Elena Cabrio, Serena Villata</i>	
A Quantitative Analysis Platform for PD-L1 Immunohistochemistry based on Point-level Supervision Model	6554
<i>Haibo Mi, Kele Xu, Yang Xiang, Yulin He, Dawei Feng, Huaimin Wang, Chun Wu, Yanming Song, Xiaolei Sun</i>	
SAGE: A Hybrid Geopolitical Event Forecasting System	6557
<i>Fred Morstatter, Aram Galstyan, Gleb Satyukov, Daniel Benjamin, Andres Abeliuk, Mehrnoosh Mirtaheri, KSM Tozammel Hossain, Pedro Szekely, Emilio Ferrara, Akira Matsui, Mark Steyvers, Stephen Bennet, David Budescu, Mark Himmelstein, Michael Ward, Andreas Beger, Michele Catasta, Rok Sosic, Jure Leskovec, Pavel Atanasov, Regina Joseph, Rajiv Sethi, Ali Abbas</i>	
Reagent: Converting Ordinary Webpages into Interactive Software Agents	6560
<i>Matthew Peveler, Jeffrey O. Kephart, Hui Su</i>	
AiD-EM: Adaptive Decision Support for Electricity Markets Negotiations	6563
<i>Tiago Pinto, Zita Vale</i>	
Deep Reinforcement Learning for Ride-sharing Dispatching and Repositioning	6566
<i>Zhiwei (Tony) Qin, Xiaocheng Tang, Yan Jiao, Fan Zhang, Chenxi Wang, Qun (Tracy) Li</i>	
GraspSnooker: Automatic Chinese Commentary Generation for Snooker Videos	6569
<i>Zhaoyue Sun, Jiase Chen, Hao Zhou, Deyu Zhou, Lei Li, Mingmin Jiang</i>	
Multi-Agent Visualization for Explaining Federated Learning	6572
<i>Xiguang Wei, Quan Li, Yang Liu, Han Yu, Tianjian Chen, Qiang Yang</i>	
Fair and Explainable Dynamic Engagement of Crowd Workers	6575
<i>Han Yu, Yang Liu, Xiguang Wei, Chuyu Zheng, Tianjian Chen, Qiang Yang, Xiong Peng</i>	
An Online Intelligent Visual Interaction System	6578
<i>Anxiang Zeng, Han Yu, Xin Gao, Kairi Ou, Zhenchuan Huang, Peng Hou, Mingli Song, Jingshu Zhang, Chunyan Miao</i>	
DeepRec: An Open-source Toolkit for Deep Learning based Recommendation	6581
<i>Shuai Zhang, Yi Tay, Lina Yao, Bin Wu, Aixin Sun</i>	
ATTENet: Detecting and Explaining Suspicious Tax Evasion Groups	6584
<i>Qinghua Zheng, Yating Lin, Huan He, Jianfei Ruan, Bo Dong</i>	
Intelligent Decision Support for Improving Power Management	6587
<i>Yongqing Zheng, Han Yu, Kun Zhang, Yuliang Shi, Cyril Leung, Chunyan Miao</i>	
Author Index	

Recommending Links to Maximize the Influence in Social Networks

Federico Corò^{1*}, Gianlorenzo D’Angelo¹ and Yllka Velaj^{2,3}

¹Gran Sasso Science Institute, L’Aquila, Italy

²CWI, Amsterdam, Netherlands

³ISI Foundation, Turin, Italy

{federico.coro, gianlorenzo.dangelo}@gssi.it, yllka.velaj@{cwi.nl, isi.it}

Abstract

Social link recommendation systems, like “People-you-may-know” on Facebook, “Who-to-follow” on Twitter, and “Suggested-Accounts” on Instagram assist the users of a social network in establishing new connections with other users. While these systems are becoming more and more important in the growth of social media, they tend to increase the popularity of users that are already popular. Indeed, since link recommenders aim at predicting users’ behavior, they accelerate the creation of links that are likely to be created in the future, and, as a consequence, they reinforce social biases by suggesting few (popular) users, while giving few chances to the majority of users to build new connections and increase their popularity.

In this paper we measure the popularity of a user by means of its social influence, which is its capability to influence other users’ opinions, and we propose a link recommendation algorithm that evaluates the links to suggest according to their increment in social influence instead of their likelihood of being created. In detail, we give a constant factor approximation algorithm for the problem of maximizing the social influence of a given set of target users by suggesting a fixed number of new connections. We experimentally show that, with few new links and small computational time, our algorithm is able to increase by far the social influence of the target users. We compare our algorithm with several baselines and show that it is the most effective one in terms of increased influence.

1 Introduction

Nowadays, recommendation system has become a key ingredient to boost social networks revenue. These systems help users decide which product to buy, news to read, or which movie to watch. Their combination with online social networks opens up new opportunities for product marketing as well as improving the user experience helping them

to make new friends and therefore be exposed to more information. Indeed, adding new connections between users increase the capabilities of a social network of spreading information which in turn increases the retention rate and the number of new subscriptions. In fact, the capability of a user to spread information, usually called social influence, is directly correlated with both the user engagement and its revenue [Chaoji *et al.*, 2012]. Being able to quickly and effectively reach a large number of people by sharing content helps a user to express and diffuse its own opinion, and to discover novel contents and information. At the same time, this increases the chances of making profits from advertisement.

Most of the existing link recommendation systems do not consider the impact of the new links on social influence, which is crucial for both users and social networks. Instead, they exploit similarity metrics of user profiles and structural network properties to estimate the likelihood that a link is adopted by users. Based on this estimation, they recommend links that are likely to be established; we point the reader to a recent survey [Li *et al.*, 2018]. For example, the Friend-of-Friend algorithm [Liben-Nowell and Kleinberg, 2007] suggests links towards the users that have the highest number of common friends with the receiver of the recommendation. While these approaches are accelerating the growth of social networks, they are also altering their structure, amplifying the popularity of well-known users and reinforcing social biases and under-representation of demographic groups, respect to their natural growth [Sanz-Cruzado and Castells, 2018; Stoica *et al.*, 2018; Su *et al.*, 2016].

In this paper we look at social influence as a measure of users’ popularity and propose a link recommendation algorithm that focuses on links that increase the social influence of a target group of users. We formulate the link recommendation task as an optimization problem that asks to suggest a fixed number of new connections to a subset of users with the aim of maximizing the network portion that is reached by their generated content.

1.1 Related Work

There exist an extensive literature on the problem of recommending links to users of a social networks [Li *et al.*, 2018; Eirinaki *et al.*, 2018]. However, there are only few studies on the problem of adding links in a network considering social influence. In the following, we focus on two widely stud-

*Contact Author

ied diffusion models, the *Independent Cascade Model* (ICM) and the *Linear Threshold Model* (LTM) [Kempe *et al.*, 2015], and review papers that study network modification problems in these models. ICM has been considered in several studies. D’Angelo *et al.* [2019] introduced the *Influence Maximization with Augmentation* problem (IMA) that consists in adding a limited number of edges incident to a given set of nodes in order to maximize their capability of spreading information. They proved that such problem is *NP*-hard to approximate within a factor greater than $1 - 1/2e$ and provide an approximation algorithm that almost matches such upper bound. Sheldon *et al.* [2012] study the problem of adding nodes in a network to maximize the diffusion in a network. They show that their problem is not submodular and propose exact integer programming formulations. Wu *et al.* [2015] considered other types of graph modifications such as modifying the probability of infect other nodes, however, they show that the problem is *NP*-hard and is neither sub- nor super-modular.

We now review papers on network modification problems under LTM. Heuristics for the edge removal problem have been studied in [Kuhlman *et al.*, 2013; Kimura *et al.*, 2008] but without providing an approximation guarantee. Khalil *et al.* [Khalil *et al.*, 2014] consider two types of graph modification, adding/deleting edges in order to minimize the information diffusion showing that this network structure modification problem has a supermodular objective and therefore can be solved by algorithms with proven approximation guarantees. Zhang *et al.* [2015] consider the problem removing edges and nodes with the aim of minimizing the information diffusion and develop algorithms with rigorous performance guarantees and good empirical performance. Experimental studies show that increasing the connectivity or the centrality of a node, by adding edges to the graph, lead also to an increase in the expected number of nodes that the diffusion process is able to reach [Crescenzi *et al.*, 2016; Parotsidis *et al.*, 2016; Papagelis, 2015].

Our Contribution We consider the IMA problem introduced in [D’Angelo *et al.*, 2019]. We first focus on the LTM model of diffusion and show that the objective function is monotone and submodular. We exploit this property to show that a greedy algorithm guarantees a $1 - 1/e$ approximation. We then propose several techniques that heuristically speed up the running time of our algorithm and of that in [D’Angelo *et al.*, 2019]. We test our algorithms in several random and real-world networks, showing that they are able to substantially increase the spreading capabilities of the target nodes and outperform many alternative baselines.

2 Notation and Problem Statement

In this section, we first define the information diffusion models used in the paper, and then introduce the IMA problem and the related notation.

2.1 Information Diffusion Models

A social network is represented by a weighted directed graph $G = (V, E, w)$, where nodes V represent users, edges E represent relationships between users, and the weight function

$w : V \times V \rightarrow [0, 1]$ represents the influence between users. Let N_v denote the set of the in-neighbors of node $v \in V$.

Both ICM and LTM models distinguish between nodes that spreads information, called *active*, and all the others, called *inactive*. Active nodes are the ones that, with some probability, diffuse the information to their neighbors. The ICM model requires a diffusion probability to be associated with each edge, whereas LTM requires an influence degree to be defined on each edge and an influence threshold on each node. For both models, the diffusion process proceeds iteratively in a synchronous way along a discrete time-axis, starting from an initial set of nodes, usually called *seeds*. The process terminates when no further node gets activated.

Kempe *et al.* showed that the distribution of the set of active nodes in the graph starting from the seed set nodes, usually denoted as A_0 , under both the ICM and LTM process is equivalent to the distribution reachable from the same set A_0 in the set of random graphs called *live-edge graphs* (Theorem 4.5, 4.6 in [Kempe *et al.*, 2015]), since the processes are point-wise identical, also the expected number of activated nodes is the same. Given a graph $G = (V, E, w)$, a live-edge graph $G' = (V, E')$, where $E' \subseteq E$, is built as follows: **(ICM)** Every edge $(u, v) \in E$ is selected to be inserted to E' independently at random with a probability proportional to its weight w_{uv} ; **(LTM)** every node $v \in V$, independently, picks at most one of its incoming edges with probability equal to the weight of that edge: edge $(u, v) \in E$ is selected with probability w_{uv} , and no edge is selected with probability $1 - \sum_{(u,v) \in E} w_{uv}$.

Moreover, evaluating the expected number of active nodes is $\#P$ -complete [Chen *et al.*, 2010], however it can be efficiently approximated by using a polynomial number live-edge graph: Kempe *et al.* proved that, by applying a multiplicative form of the Chernoff bound, we can get $(1 \pm \epsilon)$ -approximation to the expected number of active nodes, with high probability ([Kempe *et al.*, 2015, Proposition 4.1]).

2.2 Problem Statement

Given a directed graph $G = (V, E, w)$, a budget $B \in \mathbb{N}$ and a set $A \subseteq V$ of seed nodes consider a larger graph $\tilde{G} = (V, \tilde{E}, w)$ where \tilde{E} contains all the edges in E and, in addition, it contains all the remaining pairs between nodes in A and any other node in V , namely, $\tilde{E} = E \cup (A \times V)$ with the constraint that $\sum_{(u,v) \in \tilde{E}} w_{uv} \leq 1$ for any node $v \in V$.

For a set S of edges in $\tilde{E} \setminus E$, let us denote by $G(S)$ the graph augmented with the set of edges S , i.e. $G(S) = (V, E \cup S)$. In the *Influence Maximization with Augmentation problem* (IMA) we aim at finding a set $S \subseteq \tilde{E} \setminus E$ of size B in order to maximize the expected number of active nodes in $G(S)$ at the end of the ICM or LTM process. Let us denote by $\mathcal{G}(S)$ the set of all possible live-edges sampled from $G(S)$, then we denote as $\sigma(A, S) := \sum_{G' \in \mathcal{G}(S)} \mathbf{P}(G') |R_A(G')|$ the expected number of activated nodes at the end of the process with seed nodes A in graph $G(S)$, where, $R_A(G')$ is the set of reachable nodes from nodes in A in graph G' , i.e., $R_A(G') = \{u \in V : \exists \text{ path from } v \in A \text{ to } u \in G'\}$. Thus, the IMA problem consists in finding a set S^* such that

$$S^* = \arg \max_{S \subseteq \tilde{E} \setminus E : |S| \leq B} \sigma(A, S).$$

The IMA Problem Under ICM Here we review previous results on the IMA problem under ICM. The problem was originally introduced in [D'Angelo *et al.*, 2016], where the authors defined a preliminary version of the IMA problem in which the set of seed nodes is a singleton and presented a constant factor approximation algorithm for that case. In [D'Angelo *et al.*, 2019] they extended the original problem without restrictions on the seed set size. Namely, they studied the problem of adding a set of edges incident to an arbitrary set of initial seeds, without exceeding a given budget, in order to maximize the number of nodes that eventually become active. Moreover, they consider the case in which the cost of any single edge is a value between $[0, 1]$. They first proved that under ICM and with unitary cost on the edges the IMA problem is NP-hard to be approximated within a constant factor greater than $1 - (2e)^{-1}$. For this case, they provide a greedy based algorithm that guarantees a $1 - e^{-1} - \epsilon$ approximation factor. This result has been obtained by proving that the objective function $\sigma(A, S)$ is monotone and submodular w.r.t. set S , and hence the greedy algorithm in [Nemhauser *et al.*, 1978] can be used to obtain the mentioned approximation ratio. They extend this result to the general cost case by using an enumeration technique.

3 Approximation for IMA under LTM

In this section, we first prove that the function $\sigma(A, S)$ under LTM is monotone and submodular w.r.t. to the set of added edges S and then we exploit this property to provide a constant approximation algorithm for the IMA problem.

Let us denote by \mathcal{G} the set of all possible live-edge graphs sampled from graph G . For every live-edge graph $G' = (V, E') \in \mathcal{G}$ we denote by $\mathbf{P}(G')$ the probability that G' is sampled, and by $p(v, G')$ the probability for a node $v \in V$ of having an incoming edge in G' . Therefore, $p(v, G')$ is either equal to w_{uv} if there exists an edge $(u, v) \in E'$, for some $u \in N_v$, or $p(v, G') = 1 - \sum_{(u,v) \in E} w_{uv}$ if no edge is selected. Then we can easily extend this notation to a set of nodes $V' \subseteq V$ as $p(V', G') = \prod_{v \in V'} p(v, G')$. Thus,

$$\mathbf{P}(G') = \prod_{(u,v) \in E'} w_{uz} \prod_{v \in V: \exists (u,v) \in E'} \left(1 - \sum_{(u,v) \in E} w_{uv}\right).$$

Finally we denote as $\mathbf{P}_S(G') = p(V, G', S)$ the probability of a live-edge graph $G' \in \mathcal{G}(S)$, where $\mathcal{G}(S)$ is the set of all possible live-edge graphs sampled from graph $G(S)$.

We first fix the following observation: Consider the probability of a live-edge G' sampled from $\mathcal{G}(S \cup \{e\})$ and a second live-edge graph $G'' \in \mathcal{G}(S)$ when adding a new edge $e = (a, v)$ in $G(S)$. Probabilities $\mathbf{P}_{S \cup e}(G')$ and $\mathbf{P}_S(G'')$ differ only in the calculation concerning node v , since all the other nodes have the same set of possible in-neighbors with the same set of weights. Therefore, we have 3 cases:

1. For any live-edge graph $G' \in \mathcal{G}(S \cup e)$ that selects an edge different from e as incoming edge to v there exists a corresponding $G'' \in \mathcal{G}(S)$ with the same edge set, therefore we have that $\mathbf{P}_{S \cup e}(G') = \mathbf{P}_S(G'')$;
2. For any live-edge graph $G' \in \mathcal{G}(S \cup e)$ in which no incoming edge is selected for the node v

there exists a corresponding $G'' \in \mathcal{G}(S)$ with the same edge set. In this case $\mathbf{P}_{S \cup e}(G') = p(V \setminus \{v\}, G', S)(1 - \sum_{z \in N_v} w_{zv} - w_e)$ and $\mathbf{P}_S(G'') = p(V \setminus \{v\}, G'', S)(1 - \sum_{z \in N_v} w_{zv})$, where $p(V \setminus \{v\}, G', S) = p(V \setminus \{v\}, G'', S)$;

3. For any live-edge graph $G' \in \mathcal{G}(S \cup e)$, $G' = (V, E')$, that selects e as incoming edge to v we have that $\mathbf{P}_{S \cup e}(G') = p(V \setminus \{v\}, G', S)w_e$. Note that, in this case we have $|R_A(G')| \geq |R_A(G'')|$, for any live-edge graph $G'' \in \mathcal{G}(S)$ with edge set $E' \setminus \{e\}$.

We can now prove the following theorem.

Theorem 1. *Given a graph $G = (V, E)$, $\sigma(A, (V, E \cup S))$ is a monotone submodular function of $S \subseteq \bar{E} \setminus E$.*

Proof. We first prove that $\sigma(A, \cdot)$ is a monotonically increasing function, formally $\sigma(A, S \cup \{e\}) \geq \sigma(A, S)$ for any $S \subseteq \bar{E} \setminus E$ and $e = (u, v) \in \bar{E} \setminus E$.

We decompose $\sigma(A, S \cup \{e\})$ as the sum over all the live-edge graphs in which: an edge in E has been selected; v has no incoming edges; and edge e has been selected. Formally,¹

$$\begin{aligned} \sigma(A, S \cup \{e\}) &= \sum_{\substack{G' \in \mathcal{G}(S \cup e) \\ s.t. \exists (z,v) \in E', z \neq u}} \mathbf{P}_{S \cup e}(G') |R_A(G')| \\ &+ \sum_{\substack{G' \in \mathcal{G}(S \cup e) \\ s.t. \nexists (z,v) \in E'}} \mathbf{P}_{S \cup e}(G') |R_A(G')| + \sum_{\substack{G' \in \mathcal{G}(S \cup e) \\ s.t. e \in E'}} \mathbf{P}_{S \cup e}(G') |R_A(G')| \end{aligned}$$

Similarly we have that $\sigma(A, S)$ can be decomposed as: $\sigma(A, S) = \sum_{G'' \in \mathcal{G}(S)} \mathbf{P}_S(G'') |R_A(G'')| + \sum_{\substack{G'' \in \mathcal{G}(S) \\ s.t. \exists (z,v) \in E''}} \mathbf{P}_S(G'') |R_A(G'')|$.

Using observation 1 and 2 we can consider pair of live-edge graphs, one from $\mathcal{G}(S \cup e)$ and one from $\mathcal{G}(S)$, and notice that the two graphs are equivalent in the case in which an edge different from e is selected or node v has no incoming edges. Although, in the latter case the probabilities to sample the live-edge graph are not equal. Thus, we have that

$$\begin{aligned} \sigma(A, S \cup \{e\}) - \sigma(A, S) &= \sum_{\substack{G' \in \mathcal{G}(S \cup e) \\ s.t. \nexists (z,v) \in E'}} p(V \setminus \{v\}, G', S) \left(1 - \sum_{z \in N_v} w_{zv} - w_{uv}\right) |R_A(G')| \\ &+ \sum_{\substack{G' \in \mathcal{G}(S \cup e) \\ s.t. e \in E'}} p(V \setminus \{v\}, G', S) \cdot w_e \cdot |R_A(G')| - \\ &\sum_{\substack{G'' \in \mathcal{G}(S) \\ s.t. \nexists (z,v) \in E''}} p(V \setminus \{v\}, G'', S) \left(1 - \sum_{z \in N_v} w_{zv}\right) |R_A(G'')|. \end{aligned}$$

Thus, we have that $\sigma(A, S \cup \{e\}) - \sigma(A, S) = \sum_{\substack{G' \in \mathcal{G}(S) \\ s.t. \nexists (z,v) \in E'}} p(V \setminus \{v\}, G', S) w_e \left(|R_A(G''')| - |R_A(G')|\right)$, where G''' is the graph G' augmented with the edge e and the

¹ E' and E'' denote the edges sets of graphs G' and G'' , resp.

number of live-edge graphs such that the edge e has been selected is the same as the number of live-edge graphs for which no incoming edge is selected for v . Note that this value is greater or equal than zero because $|R_A(G'')| \geq |R_A(G')|$.

In order to prove that the function is submodular, we show that for each pair of sets S, T such that $S \subseteq T \subseteq \bar{E} \setminus E$ and for each $e = (a, v) \in \bar{E} \setminus (T \cup E)$, $\sigma(A, S \cup \{e\}) - \sigma(A, S) \geq \sigma(A, T \cup \{e\}) - \sigma(A, T)$. Let V' be the set of nodes that have an incoming edge from the set $T \setminus S$, namely, $V' = \{v : (w, v) \in T \setminus S\}$. Observe that for any live-edge graph $G' \in \mathcal{G}(S)$ for which the nodes in V' have no incoming edges there exists $G'_1, \dots, G'_\ell \in \mathcal{G}(T)$ such that $G' \subseteq G'_i$ for any $i = 1, \dots, \ell$ and $R_A(G') \subseteq R_A(G'_i)$, where $\ell = 2^{|T \setminus S|}$. While for all graphs $G' \in \mathcal{G}(S)$ that have at least an edge incoming a node in V' , there exists a corresponding live-edge graph $G'' \in \mathcal{G}(T)$ that is sampled with the same probability as G' . In the former case, instead, we have that the probability for each G'_i , $i = 1, \dots, \ell$, is equal to the probability of the corresponding live-edge G' in $\mathcal{G}(S)$ in which no incoming edge is selected for the nodes in V' . Formally we have that $\mathbf{P}_S(G') = p(V \setminus V', G', S) \prod_{v \in V'} (1 - \sum_{z \in N_v} w_{zv})$ and $\mathbf{P}_T(G'_i) = p(V \setminus V', G', S) \cdot p(V', G', T \setminus S)$, where

$$p(V', G', T \setminus S) = \prod_{\substack{z \in V \text{ s.t.} \\ (u, z) \in E' \cap (T \setminus S)}} w_{uz} \prod_{\substack{z \in V \text{ s.t.} \\ \nexists (u, z) \in E' \cap (T \setminus S)}} \left(1 - \sum_{w: (w, z) \in E \cup T} w_{wz}\right).$$

Then, $\sum_{i=1}^{\ell} \mathbf{P}_T(G'_i) = p(V \setminus V', G', S) \prod_{v \in V'} (1 - \sum_{z \in N_v} w_{zv}) = \mathbf{P}_S(G')$.

Finally we can write the difference in the increment when adding the edge $e = (a, v)$ in the set T as follow:

$$\begin{aligned} \sigma(A, T \cup \{e\}) - \sigma(A, T) &= \\ &\sum_{\substack{G' \in \mathcal{G}(S) \\ \text{s.t. } \exists (z, v) \in E'}} \left(\sum_{i=1}^{\ell} \mathbf{P}_T(G'_i) w_e \left(|R_A(G'')| - |R_A(G')| \right) \right) \leq \\ &\sum_{\substack{G' \in \mathcal{G}(S) \\ \text{s.t. } \exists (z, v) \in E'}} \left(p(V \setminus V', G', S) \prod_{v \in V'} \left(1 - \sum_{z \in N_v} w_{zv}\right) \right) w_e \\ &\quad \left(|R_A(G'')| - |R_A(G')| \right) \leq \\ &\sum_{\substack{G' \in \mathcal{G}(S) \\ \text{s.t. } \exists (z, v) \in E'}} p(V \setminus \{v\}, G', S) w_e \left(|R_A(G'')| - |R_A(G')| \right) \end{aligned}$$

that is equal to $\sigma(A, S \cup \{e\}) - \sigma(A, S)$, where G'' is the graph G' augmented with the edge e . \square

Thus, we can use a simple greedy algorithm (reported in Algorithm 1) to find a set S of edges whose value $\sigma(A, S)$ is at least $1 - 1/e$ times the one of an optimal solution for the IMA Problem. The algorithm iterates B times and, at each iteration, it adds to an initially empty solution S an edge $\hat{e} = (\hat{a}, \hat{v})$ s.t. $(\hat{a}, \hat{v}) \in \bar{E} \setminus E$ that gives the maximum marginal increment of the value of $\sigma(A, S)$. Note that we are not able to compute exactly the value of $\sigma(A, S)$ in polynomial time but, with probability $1 - \delta$, we can compute an

Algorithm 1 Greedy algorithm for IMA.

Require: Graph $G = (V, E, w)$; Seed set A ; Budget B

- 1: **for** $i = 1, 2, \dots, B$ **do**
 - 2: **for each** $e \in \bar{E} \setminus (E \cup S)$ **do**
 - 3: Use repeated sampling to estimate a $(1 + \lambda)$ -approx. of $\sigma(A, S \cup \{e\})$ with prob. $1 - \delta$
 - 4: Let $\tilde{\sigma}(A, S \cup \{e\})$ be the estimation
 - 5: $\hat{e} = \arg \max \{ \tilde{\sigma}(A, S \cup \{e\}) \mid e = (a, v) \in \bar{E} \setminus (E \cup S) \}$
 - 6: $S := S \cup \{ \hat{e} \}$
 - 7: **return** S
-

$1 + \lambda$ approximation of it by sampling a polynomial number of live-edge graphs, for any λ and δ [Kempe *et al.*, 2015]. We then exploit the result of Nemhauser *et al.* that allows us to analyze the greedy algorithm in the case of monotone submodular objective functions that can be approximately evaluated [Nemhauser *et al.*, 1978]. The next corollary follows.

Corollary 1. *Algorithm 1 guarantees an approximation factor of $\left(1 - \frac{1}{e} - \epsilon\right)$ for the IMA problem, where ϵ is any positive real number.*

The computational complexity of Algorithm 1 is $O(B \cdot |V| \cdot g(|V|, |E| + B))$, where $g(|V|, |E| + B)$ is the complexity of computing an approximation $\tilde{\sigma}(A, S)$ of $\sigma(A, S)$ in a graph with $|V|$ nodes and $|E| + B$ edges. More precisely, it runs in B iterations, each of which requires estimating the expected spread of $O(|V|)$ node sets. Since $g(|V|, |E| + B) = O(|E| \cdot R)$ where R is the number of simulations, then the complexity of the greedy algorithm is $O(B \cdot |V| \cdot |E| \cdot R)$ which is clearly infeasible, in terms of running time, for very large real networks.

4 Improving the Running Time

In what follows we propose some techniques to heuristically reduce the running time of the greedy algorithm. Note that this techniques can be applied to the greedy algorithm proposed in [D'Angelo *et al.*, 2019] and to Algorithm 1. In Section 5 we evaluate an implementation of the algorithm that exploits a combination of these heuristics.

- **Exploiting submodularity.** Since $\sigma(A, S)$ is submodular, we have that the increment to the expected number of active nodes after adding an edge e to $G(S)$ is monotonic non-increasing. Thus, the increment is upper bounded by any solution $S' \subseteq S$ with the addition of the new edge e , that is $\sigma(A, S' \cup \{e\}) - \sigma(A, S') \geq \sigma(A, S \cup \{e\}) - \sigma(A, S)$. We can exploit this property in Algorithm 1 to reduce the computational complexity of our algorithm. Consider the loop at line 1 for any iteration $i \geq 2$ and for some edge e , we check if the increment found so far is greater than the increment in the previous iteration, i.e., $i - 1$, with the edge e . In this case, in fact, we know that the edge e cannot increase the value of $\sigma(A, S)$ more than the maximum found so far. Therefore, in this case we prune the search.
- **Live-edge graph reduction.** At the end of each iteration of the loop at line 1 of Algorithm 1, we reduce the size

of all the live-edge graphs by removing the nodes that become influenced when adding an edge to the solution. Reducing the size after each iteration reduces the time required to compute $\tilde{\sigma}(A, S \cup \{e\})$ for each new edge e .

- **Low probability candidate edge pruning.** The greedy algorithm needs to compute $\tilde{\sigma}(A, S \cup \{e\})$ for each $e \in \bar{E} \setminus (E \cup S)$ to find an edge that maximizes this quantity. However, if we have that $w_{a_1 u} > w_{a_2 u}$ for some nodes $a_1, a_2 \in A$ and $u \in V \setminus A$, then $\sigma(A, S \cup \{a_1, u\}) > \sigma(A, S \cup \{a_2, u\})$. Thus, in the loop of line 2 we only consider, for each $u \in V \setminus A$, the edge $(a, u) \in \bar{E} \setminus (E \cup S)$ with the highest weight.
- **Reduction to Limited Seed Selection.** Given a weighted directed graph with edge weights capturing influence probabilities (in ICM or LTM), an integer B' , and two sets of nodes $A, L \subseteq V$, the *Limited Seed Selection problem* (LSS) aims to find a set of B' users $S \subseteq L$ such that, by targeting $S \cup A$, the expected number of influenced users is maximum. Nodes in S are excluded from the objective function.

Problems IMA and LSS have the same objective function, but the former looks for a set of edges, while the later looks for a set of nodes to be added to a give set of seeds. In what follows we describe how to transform an instance I_{IMA} of the IMA problem into an instance I_{LSS} of the LSS problem and how to transform a solution S_{LSS} for I_{LSS} into a solution S_{IMA} for I_{IMA} with the same value. Given $I_{IMA} = (G, A, B)$, we define $I_{LSS} = (\hat{G}, L, B)$ where $\hat{G} = (\hat{V}, \hat{E}, \hat{w})$ is a graph obtained by adding $|L| = |V \setminus A| \cdot |A|$ nodes and edges to G . Formally, let $L = \cup_{a \in A} L_a$ be the additional nodes, where $|L_i| = |V \setminus A|$ and $L_a \cap L_b = \emptyset$, for each $a, b \in A$, $a \neq b$. Then, $\hat{V} = V \cup L$ and $\hat{E} = E \cup \cup_{a \in A} (L_a \times (V \setminus A))$. The weights of the new edges are equal to that of the corresponding edges in $\bar{E} \setminus \bar{E}$, i.e. $\hat{w}_{a'v} = w_{av}$, for each $a \in A$, $a' \in L_a$, and $v \in V \setminus A$. Any solution S_{LSS} for I_{LSS} is made of nodes in L and each of these nodes corresponds to unique edge in G , we define S_{IMA} accordingly. We denote with σ_G and $\hat{\sigma}_{\hat{G}}$ the expected number of influenced nodes in G and \hat{G} , respectively. The next theorem show that the two solutions have the same value.²

Theorem 2. *In both ICM and LTM, $\sigma_{\hat{G}}(A \cup S_{LSS}, \emptyset) = \sigma_G(A, S_{IMA}) + B$.*

Thanks to Theorem 2, we can use any algorithm for LSS to solve IMA. Problem LSS is different from the influence maximization problem in [Kempe *et al.*, 2015]. However, many algorithm for this latter can be easily adapted for solving LSS. In particular, we adapt the algorithm presented in [Cohen *et al.*, 2014] such that it finds a seed set $S \subseteq L$, given a limited set of nodes L .

5 Experimental Study

In this section we experimentally evaluate the performance of our greedy algorithm and of that in [D'Angelo *et al.*, 2019].

²Note that the objective function of LSS is $\sigma_{\hat{G}}(A \cup S_{LSS}, \emptyset) - B$.

Name	$ V $	$ E $
Software Engineering (SE)	3,141	14,787
Theoretical CS (TCS)	4,172	14,272
High-Performance Comp. (HPC)	4,869	35,036
Wiki-Vote (Wiki)	7,115	103,689
Computer Graphic (CGM)	8,336	41,925
Computer Networks (CN)	9,420	53,003
Artificial Intelligence (AI)	27,617	268,460
Slashdot (SI)	51,083	130,370
Epinions (Epi)	75,879	508,837
Slashdot-Zoo (SI-z)	79,116	515,397
Digg	279,630	1,731,653
Citeseer	384,413	1,751,463
Twitter	465,017	834,797

Table 1: Real-world networks.

For both ICM and LTM, we implemented two versions of these algorithms: GREEDY1 exploits the first three heuristics described in the previous section; and GREEDY2 exploits the reduction to LSS. We compare the number of activated nodes in a graph augmented by using the greedy solution with the number of activated nodes in the original graph and in the graph augmented by using several alternative baselines.

All our experiments have been performed on a computer equipped with two Intel Xeon E5-2643 CPUs (6 cores clocked at 3.4GHz) and 128GB RAM; our programs have been implemented in C++ (gcc compiler v4.8.2 with optimization level O3). We evaluate the performance of the algorithm on four types of randomly generated directed networks which exhibit many of the structural features of complex networks and on real-world graphs that are suitable for our problem, taken from KONECT [Kunegis, 2013], Arnet-Miner [Arnetminer, 2015] and SNAP³ repositories⁴. The size of the graphs are reported in Table 1. For both synthetic and real-world networks, we choose 0.1% of the nodes in V as seeds and we add up to $B = 2 \cdot |A|$ edges. For these experiments, the seed nodes are chosen uniformly at random. The weights on the edges in both models are generated as follows: In ICM we are assigned the probabilities to the edges according to the weighted model, i.e., for each edge (u, v) , assign $w_{uv} = 1/N_v$; In LTM instead we generate for each node $v \in V$ a random variable $\bar{w}_v \in [0, 0.5]$ that represent the probability that v does not select any edge in the live-edge graph, then we assigned for each edge (u, v) in the graph a weight equal to $\frac{1-\bar{w}_v}{N_v}$ and $\frac{\bar{w}_v}{2}$ is assigned to a new edge.⁵ As a measure of the quality of the solution, we adopt the expected number of active nodes $\sigma(A, S)$. As discussed in the preliminaries, it has been proven that evaluating this function is $\#P$ -complete in general. However, by simulating the diffusion process a polynomial number of times and sampling the resulting active sets, it is possible to obtain arbitrarily good approximations to $\sigma(A, S)$. We experimentally tested that 500 samples are enough to obtain a good estima-

³<http://snap.stanford.edu/data>

⁴Here we report only the results on real world networks, those on random instances can be found in the full version of the paper.

⁵Note that it is unlikely that more than two edges towards the same node in $V \setminus A$ are added in the solution.




Figure 1: Comparison between GREEDY2 and the baselines on Slashdot-Zoo for ICM (top) and LTM (bottom).

tion. Hence, we run 500 trial to estimate the value of σ in the algorithms and in the final solution. In Table 2, we report: $\sigma(A, \emptyset)$ and $\sigma(A, \emptyset)\%$ that are the absolute and relative initial number of active nodes; $\sigma(A, S)$ and $\sigma(A, S)\%$, are the absolute and relative number of active nodes after the edge addition; the relative increment computed as $I = \frac{\sigma(A, S) - \sigma(A, \emptyset)}{\sigma(A, \emptyset)} \times 100$; and T , the time in seconds. The expected number of active nodes in GREEDY1 and GREEDY2 are similar, except from the time (GREEDY2 is faster), and a small difference is due to the sampling technique used to estimate $\tilde{\sigma}(A, S)$. Thus we reported the results for GREEDY2. From the table we can see that our algorithm is able to highly increase the number of activated nodes with respect to the original graph. Moreover, thanks to the reduction to LSS, the running time is small and this allows us to solve IMA in large networks. We compare GREEDY1 and GREEDY2 with the following alternatives that connect the given seed set to a set of B nodes chosen accordingly. AdamicAdar (AA): nodes with the highest Adamic-Adar [2003] index; PrefAtt (PA): nodes chosen according to the Preferential Attachment model [Bollobás *et al.*, 2003; Newman, 2001]; Jaccard (J): nodes with the highest Jaccard [1901] coefficient; Degree (D): nodes with the highest out-degree; Topk (TopK): nodes with the highest harmonic centrality [Boldi and Vigna, 2014]; Prob (Prob): nodes adding the edges with highest probability; Seed (KKT): nodes chosen by the greedy algorithm proposed in. [Kempe *et al.*, 2015]; Random (R): nodes extracted uniformly at random. In Figure 1, we compare GREEDY2 with

	G	$\sigma(A, \emptyset)$	$\sigma(A, \emptyset)\%$	$\sigma(A, S)$	$\sigma(A, S)\%$	$I\%$	T
ICM	SE	13.38	0.43	103.45	3.29	670.95	0.04
	TCS	9.49	0.23	97.63	2.34	928.76	0.07
	HPC	9.36	0.19	165.75	3.40	1670.83	0.07
	Wiki	10.07	0.14	333.37	4.69	3257.87	0.12
	CGM	20.34	0.24	257.62	3.09	1166.71	0.12
	CN	22.21	0.24	397.95	4.22	1765.86	0.10
	AI	68.94	0.25	1017.82	3.69	1362.71	0.47
	SI	126.11	0.25	612.52	1.20	408.63	1.72
	Epi	352.17	0.46	1230.23	1.62	249.32	3.25
	SI-z	570.84	0.72	3425.87	4.40	505.14	2.63
	Citeseer	683.16	0.18	13072.36	3.40	1813.52	12.20
	Twitter	2861.20	0.62	207061.00	44.53	7136.87	10.23
	Digg	3848.57	1.38	14835.40	5.31	285.48	14.32
LTM	SE	12.40	0.39	59.45	1.89	276.88	0.10
	TCS	8.34	0.20	51.78	1.24	430.62	0.13
	HPC	7.91	0.16	87.98	1.81	935.87	0.39
	Wiki	9.17	0.13	120.93	1.70	1151.57	1.33
	CGM	16.69	0.20	128.96	1.55	525.37	0.29
	CN	18.30	0.19	204.73	2.17	936.94	0.43
	AI	53.15	0.19	530.99	1.92	767.02	4.74
	SI	87.97	0.17	663.43	1.30	592.51	6.82
	Epi	174.98	0.23	2248.09	2.96	999.34	37.42
	SI-z	206.35	0.26	3203.52	4.05	1160.21	36.48
	Citeseer	623.82	0.16	5901.46	1.54	846.03	42.98
	Twitter	1673.07	0.36	127414.00	27.40	7515.56	13.33
	Digg	447.59	0.16	14002.80	5.01	3028.52	128.43

Table 2: Results for real-world networks.

the other approaches on the Slashdot-Zoo network. Results for other networks and for GREEDY1 are similar to this case. The plots show the average number of active nodes as a function of the number of added edges. The experiments clearly show that GREEDY2 outperforms all the alternative baselines in terms of expected number of active nodes. Indeed, all the other competitive algorithms require to add a large number of edges to A in order to significantly increase the expected number of influenced nodes with respect to the initial value ($B = 0$), whereas our algorithm increases $\sigma(A, S)$ by a greater amount with only few added edges.

6 Conclusions and Future Work

We proposed a link recommendation algorithm that, differently from other link recommenders, takes into account the amount of social influence provided by the new connections to the receiver of the recommendation. Our algorithm has a theoretical performance guarantee and, moreover, we have experimentally shown that the algorithm can be used in very large real-world networks and thus can be applied in practice.

Our algorithm aims at maximizing the amount of influence of the users that receive the recommendation. This can be used to improve the popularity of niche users and counter-balance that of famous ones. One interesting open question would be to devise a link recommender that directly considers in its objective function the *balancing* of social influence among users, in such a way that social biases are mitigated.

Acknowledgments

This work has been partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR “Algorithms, Games, and Digital Markets”.

References

- [Adamic and Adar, 2003] Lada A. Adamic and Eytan Adar. Friends and neighbors on the web. *Social Networks*, 25(3):211–230, 2003.
- [Arnetminer, 2015] Arnetminer. Arnetminer. <http://arnetminer.org>, 2015. Accessed: 2015-01-15.
- [Boldi and Vigna, 2014] Paolo Boldi and Sebastiano Vigna. Axioms for centrality. *Internet Mathematics*, 2014.
- [Bollobás *et al.*, 2003] Béla Bollobás, Christian Borgs, Jennifer T. Chayes, and Oliver Riordan. Directed scale-free graphs. In *14th SODA*, pages 132–139, 2003.
- [Chaoji *et al.*, 2012] Vineet Chaoji, Sayan Ranu, Rajeev Rastogi, and Rushi Bhatt. Recommendations to boost content spread in social networks. In *21st World Wide Web Conference 2012, WWW*, pages 529–538, 2012.
- [Chen *et al.*, 2010] Wei Chen, Yifei Yuan, and Li Zhang. Scalable influence maximization in social networks under the linear threshold model. In *10th International Conference on Data Mining*, pages 88–97, 2010.
- [Cohen *et al.*, 2014] Edith Cohen, Daniel Delling, Thomas Pajor, and Renato F. Werneck. Sketch-based influence maximization and computation: Scaling up with guarantees. In *23th CIKM*, pages 629–638, 2014.
- [Crescenzi *et al.*, 2016] Pierluigi Crescenzi, Gianlorenzo D’Angelo, Lorenzo Severini, and Yllka Velaj. Greedily improving our own closeness centrality in a network. *TKDD*, 11(1):9:1–9:32, 2016.
- [D’Angelo *et al.*, 2016] Gianlorenzo D’Angelo, Lorenzo Severini, and Yllka Velaj. Influence maximization in the independent cascade model. In *17th Italian Conference on Theoretical Computer Science.*, pages 269–274, 2016.
- [D’Angelo *et al.*, 2019] Gianlorenzo D’Angelo, Lorenzo Severini, and Yllka Velaj. Recommending links through influence maximization. *Theor. Comput. Sci.*, 2019.
- [Eirinaki *et al.*, 2018] Magdalini Eirinaki, Jerry Gao, Iraklis Varlamis, and Konstantinos Tserpes. Recommender systems for large-scale social networks: A review of challenges and solutions. *Future Generation Comp. Syst.*, 78:413–418, 2018.
- [Jaccard, 1901] Paul Jaccard. Étude comparative de la distribution florale dans une portion des alpes et des jura. *Bull Soc Vaudoise Sci Nat*, 37, 1901.
- [Kempe *et al.*, 2015] David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. *Theory of Computing*, 11:105–147, 2015.
- [Khalil *et al.*, 2014] Elias Boutros Khalil, Bistra N. Dilkina, and Le Song. Scalable diffusion-aware optimization of network topology. In *20th KDD*, pages 1226–1235, 2014.
- [Kimura *et al.*, 2008] Masahiro Kimura, Kazumi Saito, and Hiroshi Motoda. Solving the contamination minimization problem on networks for the linear threshold model. In *10th PRICAI*, pages 977–984, 2008.
- [Kuhlman *et al.*, 2013] Chris J. Kuhlman, Gaurav Tuli, Samarth Swarup, Madhav V. Marathe, and S. S. Ravi. Blocking simple and complex contagion by edge removal. In *13th ICDM*, pages 399–408, 2013.
- [Kunegis, 2013] Jérôme Kunegis. Konec: the koblenz network collection. In *22nd International Conference on World Wide Web*, pages 1343–1350, 2013.
- [Li *et al.*, 2018] Zhepeng (Lionel) Li, Xiao Fang, and Olivia R. Liu Sheng. A survey of link recommendation for social networks: Methods, theoretical foundations, and future research directions. *ACM Trans. Management Inf. Syst.*, 9(1):1:1–1:26, 2018.
- [Liben-Nowell and Kleinberg, 2007] David Liben-Nowell and Jon M. Kleinberg. The link-prediction problem for social networks. *JASIST*, 58(7):1019–1031, 2007.
- [Nemhauser *et al.*, 1978] George L. Nemhauser, Lawrence A. Wolsey, and Marshall L. Fisher. An analysis of approximations for maximizing submodular set functions - I. *Math. Program.*, 14(1):265–294, 1978.
- [Newman, 2001] Mark EJ Newman. Clustering and preferential attachment in growing networks. *Physical review E*, 64(2):025102, 2001.
- [Papagelis, 2015] Manos Papagelis. Refining social graph connectivity via shortcut edge addition. *TKDD*, 10(2):12:1–12:35, 2015.
- [Parotsidis *et al.*, 2016] Nikos Parotsidis, Evaggelia Pitoura, and Panayiotis Tsaparas. Centrality-aware link recommendations. In *9th WSDM*, pages 503–512, 2016.
- [Sanz-Cruzado and Castells, 2018] Javier Sanz-Cruzado and Pablo Castells. Enhancing structural diversity in social networks by recommending weak ties. In *12th Conference on Recommender Systems, RecSys*, pages 233–241, 2018.
- [Sheldon *et al.*, 2012] Daniel Sheldon, Bistra N. Dilkina, Adam N. Elmachtoub, Ryan Finseth, Ashish Sabharwal, Jon Conrad, Carla P. Gomes, David B. Shmoys, William Allen, Ole Amundsen, and William Vaughan. Maximizing the spread of cascades using network design. *CoRR*, abs/1203.3514, 2012.
- [Stoica *et al.*, 2018] Ana-Andreea Stoica, Christopher J. Riederer, and Augustin Chaintreau. Algorithmic glass ceiling in social networks: The effects of social recommendations on network diversity. In *27th WWW*, pages 923–932, 2018.
- [Su *et al.*, 2016] Jessica Su, Aneesh Sharma, and Sharad Goel. The effect of recommendations on network structure. In *25th International Conference on World Wide Web, WWW*, pages 1157–1167, 2016.
- [Wu *et al.*, 2015] Xiaojian Wu, Daniel Sheldon, and Shlomo Zilberstein. Efficient algorithms to optimize diffusion processes under the independent cascade model. *NIPS Work. on Networks in the Social and Information Sciences*, 2015.
- [Zhang *et al.*, 2015] Yao Zhang, Abhijin Adiga, Anil Vullikanti, and B. Aditya Prakash. Controlling propagation at group scale on networks. In *International Conference on Data Mining, ICDM*, pages 619–628, 2015.