

5th International Conference on Learning Representations (ICLR 2017)

Toulon, France
24 – 26 April 2017

Volume 1 of 4

ISBN: 978-1-7138-7271-9

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© (2017) by International Conference on Learning Representations
All rights reserved.

Printed with permission by Curran Associates, Inc. (2023)

For permission requests, please contact International Conference on Learning Representations
at the address below.

International Conference on Learning Representations
2710 E Corridor Drive
Appleton, WI 54913

www.iclr.cc

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

POSTER PRESENTATIONS

FRACTALNET: ULTRA-DEEP NEURAL NETWORKS WITHOUT RESIDUALS	1
<i>Gustav Larsson, Michael Maire, Gregory Shakhnarovich</i>	
DEEP INFORMATION PROPAGATION.....	12
<i>Samuel S. Schoenholz, Justin Gilmer, Surya Ganguli, Jascha Sohl-Dickstein</i>	
PRUNING CONVOLUTIONAL NEURAL NETWORKS FOR RESOURCE EFFICIENT INFERENCE	30
<i>Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, Jan Kautz</i>	
RECURRENT BATCH NORMALIZATION	47
<i>Tim Cooijmans, Nicolas Ballas, César Laurent, Çağlar Gülçehre, Aaron Courville</i>	
BETA-VAE: LEARNING BASIC VISUAL CONCEPTS WITH A CONSTRAINED VARIATIONAL FRAMEWORK.....	60
<i>Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, Alexander Lerchner</i>	
WORDS OR CHARACTERS? FINE-GRAINED GATING FOR READING COMPREHENSION	82
<i>Zhilin Yang, Bhuwan Dhingra, Ye Yuan, Junjie Hu, William W. Cohen, Ruslan Salakhutdinov</i>	
DEEPPDSL: A COMPILATION-BASED DOMAIN-SPECIFIC LANGUAGE FOR DEEP LEARNING.....	92
<i>Tian Zhao, Xiao Bing Huang, Yu Cao</i>	
HYPERNETWORKS.....	103
<i>David Ha, Andrew M. Dai, Quoc V. Le</i>	
CAPACITY AND TRAINABILITY IN RECURRENT NEURAL NETWORKS.....	121
<i>Jasmine Collins, Jascha Sohl-Dickstein, David Sussillo</i>	
RECURRENT HIDDEN SEMI-MARKOV MODEL.....	138
<i>Hanjun Dai, Bo Dai, Yan-Ming Zhang, Shuang Li, Le Song</i>	
LEARNING CURVE PREDICTION WITH BAYESIAN NEURAL NETWORKS.....	155
<i>Aaron Klein, Stefan Falkner, Jost Tobias Springenberg, Frank Hutter</i>	
A SIMPLE BUT TOUGH-TO-BEAT BASELINE FOR SENTENCE EMBEDDINGS.....	171
<i>Sanjeev Arora, Yingyu Liang, Tengyu Ma</i>	
LEARNING TO REPEAT: FINE GRAINED ACTION REPETITION FOR DEEP REINFORCEMENT LEARNING	187
<i>Sahil Sharma, Aravind S. Lakshminarayanan, Balaraman Ravindran</i>	
IMPROVING NEURAL LANGUAGE MODELS WITH A CONTINUOUS CACHE	211
<i>Edouard Grave, Armand Joulin, Nicolas Usunier</i>	

SNAPSHOT ENSEMBLES: TRAIN 1, GET M FOR FREE.....	220
<i>Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E. Hopcroft, Kilian Q. Weinberger</i>	
SOFT WEIGHT-SHARING FOR NEURAL NETWORK COMPRESSION.....	234
<i>Karen Ullrich, Edward Meeds, Max Welling</i>	
LEARNING TO PERFORM PHYSICS EXPERIMENTS VIA DEEP REINFORCEMENT LEARNING.....	250
<i>Misha Denil, Pulkit Agrawal, Tejas D Kulkarni, Tom Erez, Peter Battaglia, Nando de Freitas</i>	
QUERY-REDUCTION NETWORKS FOR QUESTION ANSWERING	265
<i>Minjoon Seo, Sewon Min, Ali Farhadi, Hannaneh Hajishirzi</i>	
ADVERSARIAL MACHINE LEARNING AT SCALE	278
<i>Alexey Kurakin, Ian J. Goodfellow, Samy Bengio</i>	
VARIATIONAL RECURRENT ADVERSARIAL DEEP DOMAIN ADAPTATION.....	295
<i>Sanjaya Purushotham, Wilka Carvalho, Tanachat Nilanon, Yan Liu</i>	
DISCRETE VARIATIONAL AUTOENCODERS	310
<i>Jason Tyler Rolfe</i>	
SPARSELY-CONNECTED NEURAL NETWORKS: TOWARDS EFFICIENT VLSI IMPLEMENTATION OF DEEP NEURAL NETWORKS	343
<i>Arash Ardakani, Carlo Condo, Warren J. Gross</i>	
RECURRENT MIXTURE DENSITY NETWORK FOR SPATIOTEMPORAL VISUAL ATTENTION.....	357
<i>Loris Bazzani, Hugo Larochelle, Lorenzo Torresani</i>	
EFFICIENT REPRESENTATION OF LOW-DIMENSIONAL MANIFOLDS USING DEEP NETWORKS.....	372
<i>Ronen Basri, David W. Jacobs</i>	
REINFORCEMENT LEARNING THROUGH ASYNCHRONOUS ADVANTAGE ACTOR- CRITIC ON A GPU.....	385
<i>Mohammad Babaeizadeh, Iuri Frosio, Stephen Tyree, Jason Clemons, Jan Kautz</i>	
TOPOLOGY AND GEOMETRY OF HALF-RECTIFIED NETWORK OPTIMIZATION.....	397
<i>C. Daniel Freeman, Joan Bruna</i>	
TIGHTER BOUNDS LEAD TO IMPROVED CLASSIFIERS.....	416
<i>Nicolas Le Roux</i>	
ZONEOUT: REGULARIZING RNNs BY RANDOMLY PRESERVING HIDDEN ACTIVATIONS	427
<i>David Krueger, Tegan Maharaj, Janos Kramar, Mohammad Pezeshki, Nicolas Ballas, Nan Rosemary Ke, Anirudh Goyal, Yoshua Bengio, Aaron Courville, Christopher Pal</i>	
AN INFORMATION-THEORETIC FRAMEWORK FOR FAST AND ROBUST UNSUPERVISED LEARNING VIA NEURAL POPULATION INFOMAX	438
<i>Wentao Huang, Kechen Zhang</i>	
THIRD PERSON IMITATION LEARNING	463
<i>Bradly C Stadie, Pieter Abbeel, Ilya Sutskever</i>	

DEEP VARIATIONAL BAYES FILTERS: UNSUPERVISED LEARNING OF STATE SPACE MODELS FROM RAW DATA	479
<i>Maximilian Karl, Maximilian Soelch, Justin Bayer, Patrick van der Smagt</i>	
LEARNING INVARIANT REPRESENTATIONS OF PLANAR CURVES.....	492
<i>Gautam Pai, Aaron Wetzler, Ron Kimmel</i>	
DENSITY ESTIMATION USING REAL NVP	503
<i>Laurent Dinh, Jascha Sohl-Dickstein, Samy Bengio</i>	
METACONTROL FOR ADAPTIVE IMAGINATION-BASED OPTIMIZATION.....	535
<i>Jessica B. Hamrick, Andrew J. Ballard, Razvan Pascanu, Oriol Vinyals, Nicolas Heess, Peter W. Battaglia</i>	
THE NEURAL NOISY CHANNEL	556
<i>Lei Yu, Phil Blunsom, Chris Dyer, Edward Grefenstette, Tomas Kocisky</i>	
A COMPARE-AGGREGATE MODEL FOR MATCHING TEXT SEQUENCES.....	569
<i>Shuohang Wang, Jing Jiang</i>	
CALIBRATING ENERGY-BASED GENERATIVE ADVERSARIAL NETWORKS.....	580
<i>Zihang Dai, Amjad Almahairi, Philip Bachman, Eduard Hovy, Aaron Courville</i>	
LEARNING A NATURAL LANGUAGE INTERFACE WITH NEURAL PROGRAMMER.....	597
<i>Arvind Neelakantan, Quoc V. Le, Martin Abadi, Andrew McCallum, Dario Amodei</i>	
LEARNING TO SUPEROPTIMIZE PROGRAMS.....	610
<i>Rudy Bunel, Alban Desmaison, M. Pawan Kumar, Philip H.S. Torr, Pushmeet Kohli</i>	
OPTIMAL BINARY AUTOENCODING WITH PAIRWISE CORRELATIONS.....	624
<i>Akshay Balsubramani</i>	
NEURAL PROGRAM LATTICES.....	644
<i>Chengtao Li, Daniel Tarlow, Alexander L. Gaunt, Marc Brockschmidt, Nate Kushman</i>	
ADVERSARIALY LEARNED INFERENCE.....	661
<i>Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Alex Lamb, Martin Arjovsky, Olivier Mastropietro, Aaron Courville</i>	
FRUSTRATINGLY SHORT ATTENTION SPANS IN NEURAL LANGUAGE MODELING	679
<i>Michał Daniluk, Tim Rocktäschel, Johannes Welbl, Sebastian Riedel</i>	
STEERABLE CNNs.....	689
<i>Taco S. Cohen, Max Welling</i>	
MULTI-VIEW RECURRENT NEURAL ACOUSTIC WORD EMBEDDINGS.....	703
<i>Wanjia He, Weiran Wang, Karen Livescu</i>	
TRANSFER LEARNING FOR SEQUENCE TAGGING WITH HIERARCHICAL RECURRENT NETWORKS.....	716
<i>Zhilin Yang, Ruslan Salakhutdinov, William W. Cohen</i>	
TRAINING AGENT FOR FIRST-PERSON SHOOTER GAME WITH ACTOR-CRITIC CURRICULUM LEARNING	726
<i>Yuxin Wu, Yuandong Tian</i>	

LOSSY IMAGE COMPRESSION WITH COMPRESSIVE AUTOENCODERS 736
Lucas Theis, Wenzhe Shi, Andrew Cunningham, Ferenc Huszár

LEARNING AND POLICY SEARCH IN STOCHASTIC DYNAMICAL SYSTEMS WITH
BAYESIAN NEURAL NETWORKS 755
Stefan Depeweg, José Miguel Hernández-Lobato, Finale Doshi-Velez, Steffen Udluft

VOLUME 2

NEURAL PHOTO EDITING WITH INTROSPECTIVE ADVERSARIAL NETWORKS 769
Andrew Brock, Theodore Lim, J.M. Ritchie, Nick Weston

BATCH POLICY GRADIENT METHODS FOR IMPROVING NEURAL CONVERSATION
MODELS 784
Kirthevasan Kandasamy, Yoram Bachrach, Ryota Tomioka, Daniel Tarlow, David Carter

QUASI-RECURRENT NEURAL NETWORKS 798
James Bradbury, Stephen Merity, Caiming Xiong, Richard Socher

SIGMA DELTA QUANTIZED NETWORKS 810
Peter O'Connor, Max Welling

HIGHWAY AND RESIDUAL NETWORKS LEARN UNROLLED ITERATIVE ESTIMATION 825
Klaus Greff, Rupesh K. Srivastava, Jürgen Schmidhuber

MODE REGULARIZED GENERATIVE ADVERSARIAL NETWORKS 839
Tong Che, Yanran Li, Athul Jacob, Yoshua Bengio, Wenjie Li

INTROSPECTION: ACCELERATING NEURAL NETWORK TRAINING BY LEARNING
WEIGHT EVOLUTION 852
Abhishek Sinha, Aahitagni Mukherjee, Mausoom Sarkar, Balaji Krishnamurthy

IMPROVING GENERATIVE ADVERSARIAL NETWORKS WITH DENOISING FEATURE
MATCHING 867
David Warde-Farley, Yoshua Bengio

OUTRAGEOUSLY LARGE NEURAL NETWORKS: THE SPARSELY-GATED MIXTURE-OF-
EXPERTS LAYER 878
*Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton,
Jeff Dean*

NORMALIZING THE NORMALIZERS: COMPARING AND EXTENDING NETWORK
NORMALIZATION SCHEMES 897
Mengye Ren, Renjie Liao, Raquel Urtasun, Fabian H. Sinz, Richard S. Zemel

TOWARDS DEEP INTERPRETABILITY (MUS-ROVER II): LEARNING HIERARCHICAL
REPRESENTATIONS OF TONAL MUSIC 913
Haizi Yu, Lav R. Varshney

INCORPORATING LONG-RANGE CONSISTENCY IN CNN-BASED TEXTURE
GENERATION 927
Guillaume Berger, Roland Memisevic

SUPPORT REGULARIZED SPARSE CODING AND ITS FAST ENCODER 946
Yingzhen Yang, Jiahui Yu, Pushmeet Kohli, Jianchao Yang, Thomas S. Huang

TRANSFER OF VIEW-MANIFOLD LEARNING TO SIMILARITY PERCEPTION OF NOVEL OBJECTS	965
<i>Xingyu Lin, Hao Wang, Zhihao Li, Yimeng Zhang, Alan Yuille, Tai Sing Lee</i>	
VARIABLE COMPUTATION IN RECURRENT NEURAL NETWORKS	978
<i>Yacine Jernite, Edouard Grave, Armand Joulin, Tomas Mikolov</i>	
LEARNING VISUAL SERVOING WITH DEEP FEATURES AND FITTED Q-ITERATION	990
<i>Alex X. Lee, Sergey Levine, Pieter Abbeel</i>	
LR-GAN: LAYERED RECURSIVE GENERATIVE ADVERSARIAL NETWORKS FOR IMAGE GENERATION.....	1010
<i>Jianwei Yang, Anitha Kannan, Dhruv Batra, Devi Parikh</i>	
LEARNING RECURRENT REPRESENTATIONS FOR HIERARCHICAL BEHAVIOR MODELING.....	1031
<i>Eyrun Eyjolfsdottir, Kristin Branson, Yisong Yue, Pietro Perona</i>	
EPISODIC EXPLORATION FOR DEEP DETERMINISTIC POLICIES FOR STARCRAFT MICROMANAGEMENT	1043
<i>Nicolas Usunier, Gabriel Synnaeve, Zeming Lin, Soumith Chintala</i>	
LEARNING TO NAVIGATE IN COMPLEX ENVIRONMENTS	1059
<i>Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andy Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, Dharshan Kumaran, Raia Hadsell</i>	
A RECURRENT NEURAL NETWORK WITHOUT CHAOS	1075
<i>Thomas Laurent, James von Brecht</i>	
ENERGY-BASED GENERATIVE ADVERSARIAL NETWORKS	1086
<i>Junbo Zhao, Michael Mathieu, Yann LeCun</i>	
DESIGNING NEURAL NETWORK ARCHITECTURES USING REINFORCEMENT LEARNING.....	1103
<i>Bowen Baker, Otkrist Gupta, Nikhil Naik, Ramesh Raskar</i>	
STRUCTURED ATTENTION NETWORKS	1121
<i>Yoon Kim, Carl Denton, Luong Hoang, Alexander M. Rush</i>	
A STRUCTURED SELF-ATTENTIVE SENTENCE EMBEDDING.....	1142
<i>Zhouhan Lin, Minwei Feng, Cicero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, Yoshua Bengio</i>	
LEARNING TO PLAY IN A DAY: FASTER DEEP REINFORCEMENT LEARNING BY OPTIMALITY TIGHTENING.....	1157
<i>Frank S.He, Yang Liu, Alexander G. Schwing, Jian Peng</i>	
DEEP LEARNING WITH DYNAMIC COMPUTATION GRAPHS.....	1170
<i>Moshe Looks, Marcello Herreshoff, DeLesley Hutchins, Peter Norvig</i>	
DEEP VARIATIONAL INFORMATION BOTTLENECK.....	1182
<i>Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, Kevin Murphy</i>	
GENERALIZING SKILLS WITH SEMI-SUPERVISED REINFORCEMENT LEARNING.....	1201
<i>Chelsea Finn, Tianhe Yu, Justin Fu, Pieter Abbeel, Sergey Levine</i>	

LOSS-AWARE BINARIZATION OF DEEP NETWORKS	1212
<i>Lu Hou, Quanming Yao, James T. Kwok</i>	
MACHINE COMPREHENSION USING MATCH-LSTM AND ANSWER POINTER	1223
<i>Shuohang Wang, Jing Jiang</i>	
RECURRENT ENVIRONMENT SIMULATORS	1238
<i>Silvia Chiappa, Sébastien Racaniere, Daan Wierstra, Shakir Mohamed</i>	
HYPERBAND: BANDIT-BASED CONFIGURATION EVALUATION FOR HYPERPARAMETER OPTIMIZATION	1299
<i>Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, Ameet Talwalkar</i>	
HIERARCHICAL MULTISCALE RECURRENT NEURAL NETWORKS	1314
<i>Junyoung Chung, Sungjin Ahn, Yoshua Bengio</i>	
TREE-STRUCTURED DECODING WITH DOUBLY-RECURRENT NEURAL NETWORKS	1327
<i>David Alvarez-Melis, Tommi S. Jaakkola</i>	
TRAINING COMPRESSED FULLY-CONNECTED NETWORKS WITH A DENSITY- DIVERSITY PENALTY	1344
<i>Shengjie Wang, Haoran Cai, Jeff Bilmes, William Noble</i>	
DIET NETWORKS: THIN PARAMETERS FOR FAT GENOMICS	1354
<i>Adriana Romero, Pierre Luc Carrier, Akram Erraqabi, Tristan Sylvain, Alex Auvolat, Etienne Dejoie, Marc-André Legault, Marie-Pierre Dubé, Julie G. Hussin, Yoshua Bengio</i>	
WHAT DOES IT TAKE TO GENERATE NATURAL TEXTURES?	1365
<i>Ivan Ustyuzhaninov, Wieland Brendel, Leon Gatys, Matthias Bethge</i>	
LEARNING FEATURES OF MUSIC FROM SCRATCH	1378
<i>John Thickstun, Zaid Harchaoui, Sham Kakade</i>	
AUTOMATIC RULE EXTRACTION FROM LONG SHORT TERM MEMORY NETWORKS	1392
<i>W. James Murdoch, Arthur Szlam</i>	
REASONING WITH MEMORY AUGMENTED NEURAL NETWORKS FOR LANGUAGE COMPREHENSION	1404
<i>Tsendsuren Munkhdalai, Hong Yu</i>	
GEOMETRY OF POLYSEMY	1417
<i>Jiaqi Mu, Suma Bhat, Pramod Viswanath</i>	
UNSUPERVISED CROSS-DOMAIN IMAGE GENERATION	1441
<i>Yaniv Taigman, Adam Polyak, Lior Wolf</i>	
LEARNING INVARIANT FEATURE SPACES TO TRANSFER SKILLS WITH REINFORCEMENT LEARNING	1456
<i>Abhishek Gupta, Coline Devin, YuXuan Liu, Pieter Abbeel, Sergey Levine</i>	
ENTROPY-SGD: BIASING GRADIENT DESCENT INTO WIDE VALLEYS	1470
<i>Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, Riccardo Zecchina</i>	
PAYING MORE ATTENTION TO ATTENTION: IMPROVING THE PERFORMANCE OF CONVOLUTIONAL NEURAL NETWORKS VIA ATTENTION TRANSFER	1489
<i>Sergey Zagoruyko, Nikos Komodakis</i>	

FILTER SHAPING FOR CONVOLUTIONAL NEURAL NETWORKS	1502
<i>Xingyi Li, Fuxin Li, Xiaoli Fern, Raviv Raich</i>	

VOLUME 3

UNROLLED GENERATIVE ADVERSARIAL NETWORKS	1516
<i>Luke Metz, Ben Poole, David Pfau, Jascha Sohl-Dickstein</i>	
AN ACTOR-CRITIC ALGORITHM FOR SEQUENCE PREDICTION	1541
<i>Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron Courville, Yoshua Bengio</i>	
LEARNING TO GENERATE SAMPLES FROM NOISE THROUGH INFUSION TRAINING	1558
<i>Florian Bordes, Sina Honari, Pascal Vincent</i>	
DEEP PROBABILISTIC PROGRAMMING	1577
<i>Dustin Tran, Matthew D. Hoffman, Rif A. Saurous, Eugene Brevdo, Kevin Murphy, David M. Blei</i>	
MOLLIFYING NETWORKS	1595
<i>Caglar Gulcehre, Marcin Moczulski, Francesco Visin, Yoshua Bengio</i>	
PROGRAM SYNTHESIS FOR CHARACTER LEVEL LANGUAGE MODELING	1610
<i>Pavol Bielik, Veselin Raychev, Martin Vechev</i>	
IDENTITY MATTERS IN DEEP LEARNING	1627
<i>Moritz Hardt, Tengyu Ma</i>	
HADAMARD PRODUCT FOR LOW-RANK BILINEAR POOLING	1641
<i>Jin-Hwa Kim, Kyoung-Woon On, Woosang Lim, Jeonghee Kim, Jung-Woo Ha, Byoung-Tak Zhang</i>	
INDUCTIVE BIAS OF DEEP CONVOLUTIONAL NETWORKS THROUGH POOLING GEOMETRY	1655
<i>Nadav Cohen, Amnon Shashua</i>	
PRUNING FILTERS FOR EFFICIENT CONVNETS	1683
<i>Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, Hans Peter Graf</i>	
SAMPLERNN: AN UNCONDITIONAL END-TO-END NEURAL AUDIO GENERATION MODEL	1696
<i>Soroush Mehri, Kundan Kumar, Ishaan Gulrajani, Rithesh Kumar, Shubham Jain, Jose Sotelo, Aaron Courville, Yoshua Bengio</i>	
DEEP MULTI-TASK REPRESENTATION LEARNING: A TENSOR FACTORISATION APPROACH	1707
<i>Yongxin Yang, Timothy M. Hospedales</i>	
EPOPT: LEARNING ROBUST NEURAL NETWORK POLICIES USING MODEL ENSEMBLES	1719
<i>Aravind Rajeswaran, Sarvjeet Ghotra, Balaraman Ravindran, Sergey Levine</i>	
HOLSTEP: A MACHINE LEARNING DATASET FOR HIGHER-ORDER LOGIC THEOREM PROVING	1734
<i>Cezary Kaliszky, François Chollet, Christian Szegedy</i>	

TOWARDS A NEURAL STATISTICIAN	1746
<i>Harrison Edwards, Amos Storkey</i>	
EXPLORING SPARSITY IN RECURRENT NEURAL NETWORKS	1759
<i>Sharan Narang, Greg Diamos, Shubho Sengupta, Erich Elsen</i>	
SGDR: STOCHASTIC GRADIENT DESCENT WITH WARM RESTARTS	1769
<i>Ilya Loshchilov, Frank Hutter</i>	
PREDICTING MEDICATIONS FROM DIAGNOSTIC CODES WITH RECURRENT NEURAL NETWORKS.....	1785
<i>Jacek M. Bajor, Thomas A. Lasko</i>	
ONLINE BAYESIAN TRANSFER LEARNING FOR SEQUENTIAL DATA MODELING	1804
<i>Priyank Jaini, Zhitang Chen, Pablo Carbajal, Edith Law, Laura Middleton, Kayla Regan, Mike Schaekermann, George Trimponias, James Tung, Pascal Poupart</i>	
FINE-GRAINED ANALYSIS OF SENTENCE EMBEDDINGS USING AUXILIARY PREDICTION TASKS	1824
<i>Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, Yoav Goldberg</i>	
TRACKING THE WORLD STATE WITH RECURRENT ENTITY NETWORKS.....	1837
<i>Mikael Henaff, Jason Weston, Arthur Szlam, Antoine Bordes, Yann LeCun</i>	
POINTER SENTINEL MIXTURE MODELS.....	1851
<i>Stephen Merity, Caiming Xiong, James Bradbury, Richard Socher</i>	
TRAINING DEEP NEURAL-NETWORKS USING A NOISE ADAPTATION LAYER.....	1866
<i>Jacob Goldberger, Ehud Ben-Reuven</i>	
SAMPLE EFFICIENT ACTOR-CRITIC WITH EXPERIENCE REPLAY	1875
<i>Ziyu Wang, Victor Bapst, Nicolas Heess, Volodymyr Mnih, Remi Munos, Koray Kavukcuoglu, Nando de Freitas</i>	
REVISITING CLASSIFIER TWO-SAMPLE TESTS.....	1895
<i>David Lopez-Paz, Maxime Oquab</i>	
PIXELCNN++: IMPROVING THE PIXELCNN WITH DISCRETIZED LOGISTIC MIXTURE LIKELIHOOD AND OTHER MODIFICATIONS	1910
<i>Tim Salimans, Andrej Karpathy, Xi Chen, Diederik P. Kingma</i>	
CATEGORICAL REPARAMETERIZATION WITH GUMBEL-SOFTMAX.....	1920
<i>Eric Jang, Shixiang Gu, Ben Poole</i>	
CENTRAL MOMENT DISCREPANCY (CMD) FOR DOMAIN-INVARIANT REPRESENTATION LEARNING.....	1932
<i>Werner Zellinger, Thomas Grubinger, Edwin Lughofer, Thomas Natschläger, Susanne Saminger-Platz</i>	
REGULARIZING CNNs WITH LOCALLY CONSTRAINED DECORRELATIONS	1945
<i>Pau Rodríguez, Jordi González, Guillem Cucurull, Josep M. Gonfau, Xavier Roca</i>	
DEEP PREDICTIVE CODING NETWORKS FOR VIDEO PREDICTION AND UNSUPERVISED LEARNING	1956
<i>William Lotter, Gabriel Kreiman, David Cox</i>	

STICK-BREAKING VARIATIONAL AUTOENCODERS.....	1974
<i>Eric Nalisnick, Padhraic Smyth</i>	
WHY DEEP NEURAL NETWORKS FOR FUNCTION APPROXIMATION?	1986
<i>Shiyu Liang, R. Srikant</i>	
MAXIMUM ENTROPY FLOW NETWORKS.....	2003
<i>Gabriel Loaiza-Ganem , Yuanjun Gao , John P. Cunningham</i>	
LATENT SEQUENCE DECOMPOSITIONS.....	2016
<i>William Chan, Yu Zhang, Quoc Le, Navdeep Jaitly</i>	
GENERATIVE MODELS AND MODEL CRITICISM VIA OPTIMIZED MAXIMUM MEAN DISCREPANCY	2028
<i>Danica J. Sutherland, Hsiao-Yu Tung, Heiko Strathmann, Soumyajit De, Aaditya Ramdas, Alex Smola, Arthur Gretton</i>	
EFFICIENT VECTOR REPRESENTATION FOR DOCUMENTS THROUGH CORRUPTION	2039
<i>Minmin Chen</i>	
DEEPCODER: LEARNING TO WRITE PROGRAMS.....	2052
<i>Matej Balog, Alexander L. Gaunt, Marc Brockschmidt, Sebastian Nowozin, Daniel Tarlow</i>	
ON DETECTING ADVERSARIAL PERTURBATIONS	2072
<i>Jan Hendrik Metzen, Tim Genewein, Volker Fischer, Bastian Bischoff</i>	
LEARNING TO REMEMBER RARE EVENTS	2084
<i>Lukasz Kaiser, Ofir Nachum, Aurko Roy, Samy Bengio</i>	
DATA NOISING AS SMOOTHING IN NEURAL NETWORK LANGUAGE MODELS.....	2094
<i>Ziang Xie, Sida I. Wang, Jiwei Li, Daniel Lévy, Aiming Nie, Dan Jurafsky, Andrew Y. Ng</i>	
TRAINED TERNARY QUANTIZATION.....	2106
<i>Chenzhuo Zhu, Song Han, Huizi Mao, William J. Dally</i>	
DO DEEP CONVOLUTIONAL NETS REALLY NEED TO BE DEEP AND CONVOLUTIONAL?.....	2116
<i>Gregor Urban, Krzysztof J. Geras, Samira Ebrahimi Kahou, Ozlem Aslan, Shengjie Wang, Abdelrahman Mohamed, Matthai Philipose, Matt Richardson, Rich Caruana</i>	
VARIATIONAL LOSSY AUTOENCODER.....	2129
<i>Xi Chen, Diederik P. Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, Pieter Abbeel</i>	
GENERATIVE MULTI-ADVERSARIAL NETWORKS.....	2146
<i>Ishan Durugkar, Ian Gemp, Sridhar Mahadevan</i>	
INCREMENTAL NETWORK QUANTIZATION: TOWARDS LOSSLESS CNNs WITH LOW- PRECISION WEIGHTS.....	2160
<i>Aojun Zhou, Anbang Yao, Yiwon Guo, Lin Xu, Yurong Chen</i>	
MULTILAYER RECURRENT NETWORK MODELS OF PRIMATE RETINAL GANGLION CELL RESPONSES	2174
<i>Eleanor Batty, Josh Merel, Nora Brackbill, Alexander Heitman, Alexander Sher, Alan Litke, E.J. Chichilnisky, Liam Paninski</i>	

A COMPOSITIONAL OBJECT-BASED APPROACH TO LEARNING PHYSICAL DYNAMICS.....	2186
<i>Michael Chang, Tomer Ullman, Antonio Torralba, Joshua Tenenbaum</i>	
DECOMPOSING MOTION AND CONTENT FOR NATURAL VIDEO SEQUENCE PREDICTION	2201
<i>Ruben Villegas, Jimei Yang, Seunghoon Hong, Xunyu Lin, Honglak Lee</i>	
AUTOENCODING VARIATIONAL INFERENCE FOR TOPIC MODELS	2223
<i>Akash Srivastava, Charles Sutton</i>	
DELVING INTO TRANSFERABLE ADVERSARIAL EXAMPLES AND BLACK-BOX ATTACKS	2235
<i>Yanpei Liu, Xinyun Chen, Chang Liu, Dawn Song</i>	
LEARNING TO OPTIMIZE	2249
<i>Ke Li, Jitendra Malik</i>	
UNDERSTANDING TRAINABLE SPARSE CODING WITH MATRIX FACTORIZATION.....	2262
<i>Thomas Moreau, Joan Bruna</i>	

VOLUME 4

EMERGENCE OF FOVEAL IMAGE SAMPLING FROM LEARNING TO ATTEND IN VISUAL SCENES	2275
<i>Brian Cheung, Eric Weiss, Bruno Olshausen</i>	
IMPROVING POLICY GRADIENT BY EXPLORING UNDER-APPRECIATED REWARDS.....	2284
<i>Ofir Nachum, Mohammad Norouzi, Dale Schuurmans</i>	
LEARNING TO QUERY, REASON, AND ANSWER QUESTIONS ON AMBIGUOUS TEXTS.....	2301
<i>Xiaoxiao Guo, Tim Klinger, Clemens Rosenbaum, Joseph P. Bigus, Murray Campbell, Ban Kawas, Kartik Talamadupula, Gerry Tesauero, Satinder Singh</i>	
NEURO-SYMBOLIC PROGRAM SYNTHESIS	2313
<i>Emilio Parisotto, Abdel-rahman Mohamed, Rishabh Singh, Lihong Li, Dengyong Zhou, Pushmeet Kohli</i>	
DSD: DENSE-SPARSE-DENSE TRAINING FOR DEEP NEURAL NETWORKS.....	2328
<i>Song Han, Jeff Pool, Sharan Narang, Huizi Mao, Enhao Gong, Shijian Tang, Erich Elsen, Peter Vajda, Manohar Paluri, John Tran, Bryan Catanzaro, William J. Dally</i>	
PIXELVAE: A LATENT VARIABLE MODEL FOR NATURAL IMAGES	2341
<i>Ishaan Gulrajani, Kundan Kumar, Faruk Ahmed, Adrien Ali Taiga, Francesco Visin, David Vazquez, Aaron Courville</i>	
DEEP BIAFFINE ATTENTION FOR NEURAL DEPENDENCY PARSING.....	2357
<i>Timothy Dozat, Christopher D. Manning</i>	
TOWARDS THE LIMIT OF NETWORK QUANTIZATION.....	2365
<i>Yoojin Choi, Mostafa El-Khamy, Jungwon Lee</i>	
STOCHASTIC NEURAL NETWORKS FOR HIERARCHICAL REINFORCEMENT LEARNING.....	2379
<i>Carlos Florensa, Yan Duan, Pieter Abbeel</i>	

DYNAMIC COATTENTION NETWORKS FOR QUESTION ANSWERING	2396
<i>Caiming Xiong, Victor Zhong, Richard Socher</i>	
A BASELINE FOR DETECTING MISCLASSIFIED AND OUT-OF-DISTRIBUTION EXAMPLES IN NEURAL NETWORKS.....	2410
<i>Dan Hendrycks, Kevin Gimpel</i>	
ADVERSARIAL FEATURE LEARNING	2422
<i>Jeff Donahue, Philipp Krähenbühl, Trevor Darrell</i>	
DROPOUT WITH EXPECTATION-LINEAR REGULARIZATION	2440
<i>Xuezhe Ma, Yingkai Gao, Zhiting Hu, Yaoliang Yu, Yuntian Deng, Eduard Hovy</i>	
FASTER CNNs WITH DIRECT SPARSE CONVOLUTIONS AND GUIDED PRUNING.....	2463
<i>Jongsoo Park, Sheng Li, Wei Wen, Ping Tak Peter Tang, Hai Li, Yiran Chen, Pradeep Dubey</i>	
DISTRIBUTED SECOND-ORDER OPTIMIZATION USING KRONECKER-FACTORED APPROXIMATIONS	2474
<i>Jimmy Ba, Roger Grosse, James Martens</i>	
TOPICRNN: A RECURRENT NEURAL NETWORK WITH LONG-RANGE SEMANTIC DEPENDENCY	2491
<i>Adji B. Dieng, Chong Wang, Jianfeng Gao, John Paisley</i>	
ON THE QUANTITATIVE ANALYSIS OF DECODER-BASED GENERATIVE MODELS.....	2504
<i>Yuhuai Wu, Yuri Burda, Ruslan Salakhutdinov, Roger Grosse</i>	
OFFLINE BILINGUAL WORD VECTORS, ORTHOGONAL TRANSFORMATIONS AND THE INVERTED SOFTMAX	2521
<i>Samuel L. Smith, David H. P. Turban, Steven Hamblin, Nils Y. Hammerla</i>	
ATTEND, ADAPT AND TRANSFER: ATTENTIVE DEEP ARCHITECTURE FOR ADAPTIVE TRANSFER FROM MULTIPLE SOURCES IN THE SAME DOMAIN	2531
<i>Janarthan Rajendran, Aravind Lakshminarayanan, Mitesh M. Khapra, Prasanna P, Balaraman Ravindran</i>	
LEARNING THROUGH DIALOGUE INTERACTIONS BY ASKING QUESTIONS.....	2549
<i>Jiwei Li, Alexander H. Miller, Sumit Chopra, Marc'Aurelio Ranzato, Jason Weston</i>	
PALEO: A PERFORMANCE MODEL FOR DEEP NEURAL NETWORKS	2565
<i>Hang Qi, Evan R. Sparks, Ameet Talwalkar</i>	
A LEARNED REPRESENTATION FOR ARTISTIC STYLE	2575
<i>Vincent Dumoulin, Jonathon Shlens, Manjunath Kudlur</i>	
VISUALIZING DEEP NEURAL NETWORK DECISIONS: PREDICTION DIFFERENCE ANALYSIS.....	2601
<i>Luisa M Zintgraf, Taco S Cohen, Tameem Adel, Max Welling</i>	
TYING WORD VECTORS AND WORD CLASSIFIERS: A LOSS FRAMEWORK FOR LANGUAGE MODELING.....	2613
<i>Hakan Inan, Khashayar Khosravi, Richard Socher</i>	
DIALOGUE LEARNING WITH HUMAN-IN-THE-LOOP.....	2626
<i>Jiwei Li, Alexander H. Miller, Sumit Chopra, Marc'Aurelio Ranzato, Jason Weston</i>	

BIDIRECTIONAL ATTENTION FLOW FOR MACHINE COMPREHENSION	2649
<i>Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, Hannaneh Hajishirzi</i>	
LEARNING TO COMPOSE WORDS INTO SENTENCES WITH REINFORCEMENT LEARNING.....	2662
<i>Dani Yogatama, Phil Blunsom, Chris Dyer, Edward Grefenstette, Wang Ling</i>	
NONPARAMETRIC NEURAL NETWORKS	2672
<i>George Philipp, Jaime G. Carbonell</i>	
TEMPORAL ENSEMBLING FOR SEMI-SUPERVISED LEARNING	2700
<i>Samuli Laine, Timo Aila</i>	
SEMI-SUPERVISED CLASSIFICATION WITH GRAPH CONVOLUTIONAL NETWORKS.....	2713
<i>Thomas N. Kipf, Max Welling</i>	
TRUSTING SVM FOR PIECEWISE LINEAR CNNs	2727
<i>Leonard Berrada, Andrew Zisserman, M. Pawan Kumar</i>	
LIE-ACCESS NEURAL TURING MACHINES.....	2750
<i>Greg Yang, Alexander Rush</i>	
ADVERSARIAL TRAINING METHODS FOR SEMI-SUPERVISED TEXT CLASSIFICATION.....	2763
<i>Takeru Miyato, Andrew M. Dai, Ian Goodfellow</i>	
COMBINING POLICY GRADIENT AND Q-LEARNING.....	2774
<i>Brendan O'Donoghue, Remi Munos, Koray Kavukcuoglu, Volodymyr Mnih</i>	
 <u>ORAL PRESENTATIONS</u>	
MAKING NEURAL PROGRAMMING ARCHITECTURES GENERALIZE VIA RECURSION	2789
<i>Jonathon Cai, Richard Shin, Dawn Song</i>	
TOWARDS PRINCIPLED METHODS FOR TRAINING GENERATIVE ADVERSARIAL NETWORKS.....	2809
<i>Martin Arjovsky, Leon Bottou</i>	
LEARNING END-TO-END GOAL-ORIENTED DIALOG	2826
<i>Antoine Bordes, Y-Lan Boureau, Jason Weston</i>	
NEURAL ARCHITECTURE SEARCH WITH REINFORCEMENT LEARNING	2841
<i>Barret Zoph, Quoc Le</i>	
AMORTISED MAP INFERENCE FOR IMAGE SUPER-RESOLUTION	2857
<i>Casper Kaae Sønderby, Jose Caballero, Lucas Theis, Wenzhe Shi, Ferenc Huszár</i>	
ON LARGE-BATCH TRAINING FOR DEEP LEARNING: GENERALIZATION GAP AND SHARP MINIMA.....	2874
<i>Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, Ping Tak Peter Tang</i>	
SEMI-SUPERVISED KNOWLEDGE TRANSFER FOR DEEP LEARNING FROM PRIVATE TRAINING DATA	2890
<i>Nicolas Papernot, Martín Abadi, Úlfar Erlingsson, Ian Goodfellow, Kunal Talwar</i>	

LEARNING GRAPHICAL STATE TRANSITIONS	2906
<i>Daniel D. Johnson</i>	
MULTI-AGENT COOPERATION AND THE EMERGENCE OF (NATURAL) LANGUAGE.....	2925
<i>Angeliki Lazaridou, Alexander Peysakhovich, Marco Baroni</i>	
OPTIMIZATION AS A MODEL FOR FEW-SHOT LEARNING	2936
<i>Sachin Ravi, Hugo Larochelle</i>	
LEARNING TO ACT BY PREDICTING THE FUTURE.....	2947
<i>Alexey Dosovitskiy, Vladlen Koltun</i>	
END-TO-END OPTIMIZED IMAGE COMPRESSION	2961
<i>Johannes Ballé, Valero Laparra, Eero P. Simoncelli</i>	
Q-PROP: SAMPLE-EFFICIENT POLICY GRADIENT WITH AN OFF-POLICY CRITIC.....	2988
<i>Shixiang Gu, Timothy Lillicrap, Zoubin Ghahramani, Richard E. Turner, Sergey Levine</i>	
UNDERSTANDING DEEP LEARNING REQUIRES RETHINKING GENERALIZATION	3001
<i>Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, Oriol Vinyals</i>	
REINFORCEMENT LEARNING WITH UNSUPERVISED AUXILIARY TASKS	3016
<i>Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, Koray Kavukcuoglu</i>	

Author Index