

Poster Sessions : Conference (Floor 1) and Workshop (Floor 3) Track papers presented at each of the poster sessions (on Monday, Tuesday or Wednesday, in the morning or evening).

Monday Morning (April 24th, 10:30am to 12:30pm)

C1: Making Neural Programming Architectures Generalize via Recursion
C2: Learning Graphical State Transitions
C3: Distributed Second-Order Optimization using Kronecker-Factored Approximations
C4: Normalizing the Normalizers: Comparing and Extending Network Normalization Schemes
C5: Neural Program Lattices
C6: Diet Networks: Thin Parameters for Fat Genomics
C7: Unsupervised Cross-Domain Image Generation
C8: Towards Principled Methods for Training Generative Adversarial Networks
C9: Recurrent Mixture Density Network for Spatiotemporal Visual Attention
C10: Paying More Attention to Attention: Improving the Performance of CNN via Attention Transfer
C11: Pruning Filters for Efficient ConvNets
C12: Stick-Breaking Variational Autoencoders
C13: Identity Matters in Deep Learning
C14: On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima
C15: Recurrent Hidden Semi-Markov Model
C16: Nonparametric Neural Networks
C17: Learning to Generate Samples from Noise through Infusion Training
C18: An Information-Theoretic Framework for Unsup. Learning via Neural Population Infomax
C19: Highway and Residual Networks learn Unrolled Iterative Estimation
C20: Soft Weight-Sharing for Neural Network Compression
C21: Snapshot Ensembles: Train 1, Get M for Free
C22: Towards a Neural Statistician
C23: Learning Curve Prediction with Bayesian Neural Networks
C24: Learning End-to-End Goal-Oriented Dialog
C25: Multi-Agent Cooperation and the Emergence of (Natural) Language
C26: Efficient Vector Representation for Documents through Corruption
C27: Improving Neural Language Models with a Continuous Cache
C28: Program Synthesis for Character Level Language Modeling
C29: Tracking the World State with Recurrent Entity Networks
C30: Reinforcement Learning with Unsupervised Auxiliary Tasks
C31: Neural Architecture Search with Reinforcement Learning
C32: Sample Efficient Actor-Critic with Experience Replay
C33: Learning to Act by Predicting the Future
Floor 3 :
W1: Extrapolation and learning equations
W2: Effectiveness of Transfer Learning in EHR data
W3: Intelligent synapses for multi-task and transfer learning
W4: Unsupervised and Efficient Neural Graph Model with Distributed Representations
W5: Accelerating SGD for Distributed Deep-Learning Using an Approximated Hessian Matrix
W6: Accelerating Eulerian Fluid Simulation With Convolutional Networks
W7: Forced to Learn: Discovering Disentangled Representations Without Exhaustive Labels
W8: Dataset Augmentation in Feature Space
W9: Learning Algorithms for Active Learning
W10: Reinterpreting Importance-Weighted Autoencoders
W11: Robustness to Adversarial Examples through an Ensemble of Specialists
W12: (empty)
W13: On Hyperparameter Optimization in Learning Systems
W14: Recurrent Normalization Propagation
W15: Joint Training of Ratings and Reviews with Recurrent Recommender Networks
W16: Towards an Automatic Turing Test: Learning to Evaluate Dialogue Responses
W17: Joint Embeddings of Scene Graphs and Images
W18: Unseen Style Transfer Based on a Conditional Fast Style Transfer Network

Monday Afternoon (April 24th, 4:30pm to 6:30pm)

C1: Neuro-Symbolic Program Synthesis
C2: Generative Models and Model Criticism via Optimized Max Mean Discrepancy
C3: Trained Ternary Quantization
C4: DSD: Dense-Sparse-Dense Training for Deep Neural Networks
C5: A Compositional Object-Based Approach to Learning Physical Dynamics
C6: Multilayer Recurrent Network Models of Primate Retinal Ganglion Cells
C7: Improving Generative Adversarial Networks with Denoising Feature Matching
C8: Transfer of View-manifold Learning to Similarity Perception of Novel Objects
C9: What does it take to generate natural textures?
C10: Emergence of foveal image sampling from learning to attend in visual scenes
C11: PixelCNN++: A PixelCNN Imp. with Discretized Logistic Mixture Likelihood
C12: Learning to Optimize
C13: Do Deep Convolutional Nets Really Need to be Deep and Convolutional?
C14: Optimal Binary Autoencoding with Pairwise Correlations
C15: On the Quantitative Analysis of Decoder-Based Generative Models
C16: Adversarial machine learning at scale
C17: Transfer Learning for Sequence Tagging with Hierarchical Recurrent Networks
C18: Capacity and Learnability in Recurrent Neural Networks
C19: Deep Learning with Dynamic Computation Graphs
C20: Exploring Sparsity in Recurrent Neural Networks
C21: Structured Attention Networks
C22: Learning to Repeat: Fine Grained Action Repetition for Deep Reinforcement
C23: Variational Lossy Autoencoder
C24: Learning to Query, Reason, and Answer Questions On Ambiguous Texts
C25: Deep Biaffine Attention for Neural Dependency Parsing
C26: A Compare-Aggregate Model for Matching Text Sequences
C27: Data Noising as Smoothing in Neural Network Language Models
C28: Neural Variational Inference For Topic Models
C29: Bidirectional Attention Flow for Machine Comprehension
C30: Q-Prop: Sample-Efficient Policy Gradient with An Off-Policy Critic
C31: Stochastic Neural Networks for Hierarchical Reinforcement Learning
C32: Learning Invariant Feature Spaces to Transfer Skills with Reinforcement Learning
C33: Third Person Imitation Learning
Floor 3 :
W1: Audio Super-Resolution using Neural Networks
W2: Semantic embeddings for program behaviour patterns
W3: De novo drug design with deep generative models : an empirical study
W4: Memory Matching Networks for Genomic Sequence Classification
W5: Char2Wav: End-to-End Speech Synthesis
W6: Fast Chirplet Transform Injects Priors in Deep L. of Animal Calls and Speech
W7: Weight-averaged consistency targets improve semi-supervised Deep L.
W8: Particle Value Functions
W9: Out-of-class novelty generation: an experimental foundation
W10: Performance guarantees for transferring representations
W11: Generative Adversarial Learning of Markov Chains
W12: Short and Deep: Sketching and Neural Networks
W13: Understanding intermediate layers using linear classifier probes
W14: Symmetry-Breaking Convergence of 2-layered NN with ReLU nonlinearity
W15: Neural Combinatorial Optimization with Reinforcement Learning
W16: Tactics of Adversarial Attacks on Deep Reinforcement Learning Agents
W17: Adversarial Discriminative Domain Adaptation (workshop extended abstract)
W18: Efficient Sparse-Winograd Convolutional Neural Networks
W19: Neural Expectation Maximization

Tuesday Morning (April 25th, 10:30am to 12:30pm)

C1: DeepDSL: A Compilation-based Domain-Specific Language for Deep Learning
C2: A Self Attentive Sentence Embedding
C3: Deep Probabilistic Programming
C4: Lie-Access Neural Turing Machines
C5: Learning Features of Music From Scratch
C6: Mode Regularized Generative Adversarial Networks
C7: End-to-end Optimized Image Compression
C8: Variational Recurrent Adversarial Deep Domain Adaptation
C9: Steerable CNNs
C10: Deep Predictive Coding Networks for Video Prediction and Unsupervised Learning
C11: PixelVAE: A Latent Variable Model for Natural Images
C12: A recurrent neural network without chaos
C13: Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer
C14: Tree-structured decoding with doubly-recurrent neural networks
C15: Introspection: Accelerating Neural Network Training By Learning Weight Evolution
C16: Hyperband: Bandit-Based Configuration Evaluation for Hyperparameter Optimization
C17: Quasi-Recurrent Neural Networks
C18: Attend, Adapt and Transfer: Attentive Deep Arch. For Ad. Transfer from multiple sources
C19: A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks
C20: Trusting SVM for Piecewise Linear CNNs
C21: Maximum Entropy Flow Networks
C22: The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables
C23: Unrolled Generative Adversarial Networks
C24: A Simple but Tough-to-Beat Baseline for Sentence Embeddings
C25: Query-Reduction Networks for Question Answering
C26: Machine Comprehension Using Match-LSTM and Answer Pointer
C27: Words or Characters? Fine-grained Gating for Reading Comprehension
C28: Dynamic Coattention Networks For Question Answering
C29: Multi-view Recurrent Neural Acoustic Word Embeddings
C30: Episodic Exploration for Deep Deterministic Policies for StarCraft Micromanagement
C31: Training Agent for First-Person Shooter Game with Actor-Critic Curriculum Learning
C32: Generalizing Skills with Semi-Supervised Reinforcement Learning
C33: Improving Policy Gradient by Exploring Under-appreciated Rewards
Floor 3 :
W1: Programming With a Differentiable Forth Interpreter
W2: Unsupervised Feature Learning for Audio Analysis
W3: Neural Functional Programming
W4: A Smooth Optimisation Perspective on Training Feedforward Neural Networks
W5: Synthetic Gradient Methods with Virtual Forward-Backward Networks
W6: Explaining the Learning Dynamics of Direct Feedback Alignment
W7: Training a Subsampling Mechanism in Expectation
W8: Deep Kernel Machines via the Kernel Reparametrization Trick
W9: Encoding and Decoding Representations with Sum- and Max-Product Networks
W10: Embracing Data Abundance
W11: Variational Intrinsic Control
W12: Fast Adaptation in Generative Models with Generative Matching Networks
W13: Efficient variational Bayesian neural network ensembles for outlier detection
W14: Emergence of Language with Multi-agent Games: Learning to Comm. With Seq. of Symbo
W15: Adaptive Feature Abstraction for Translating Video to Language
W16: Delving into adversarial attacks on deep policies
W17: Tuning Recurrent Neural Networks with Reinforcement Learning
W18: DeepMask: Masking DNN Models for robustness against adversarial samples
W19: Restricted Boltz. Machines provide an accurate metric for retinal responses to visual stimuli

Tuesday Afternoon (April 25th, 2:00pm to 4:00pm)

C1: Sigma Delta Quantized Networks
C2: Paleo: A Performance Model for Deep Neural Networks
C3: DeepCoder: Learning to Write Programs
C4: Topology and Geometry of Deep Rectified Network Optimization Landscapes
C5: Incremental Network Quantization: Towards Lossless CNNs with Low-precision Weights
C6: Learning to Perform Physics Experiments via Deep Reinforcement Learning
C7: Decomposing Motion and Content for Natural Video Sequence Prediction
C8: Calibrating Energy-based Generative Adversarial Networks
C9: Pruning Convolutional Neural Networks for Resource Efficient Inference
C10: Incorporating long-range consistency in CNN-based texture generation
C11: Lossy Image Compression with Compressive Autoencoders
C12: LR-GAN: Layered Recursive Generative Adversarial Networks for Image Generation
C13: Semi-supervised Knowledge Transfer for Deep Learning from Private Training Data
C14: Deep Variational Bayes Filters: Unsup Learning of State Space Models from Raw Data
C15: Mollifying Networks
C16: beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework
C17: Categorical Reparameterization with Gumbel-Softmax
C18: Online Bayesian Transfer Learning for Sequential Data Modeling
C19: Latent Sequence Decompositions
C20: Density estimation using Real NVP
C21: Recurrent Batch Normalization
C22: SGDR: Stochastic Gradient Descent with Restarts
C23: Variable Computation in Recurrent Neural Networks
C24: Deep Variational Information Bottleneck
C25: SampleRNN: An Unconditional End-to-End Neural Audio Generation Model
C26: TopicRNN: A Recurrent Neural Network with Long-Range Semantic Dependency
C27: Frustratingly Short Attention Spans in Neural Language Modeling
C28: Offline Bilingual Word Vectors, Orthogonal Transformations and the Inverted Softmax
C29 : Learning a natural Language Interface with Neural Programmer
C30: Designing Neural Network Architectures using Reinforcement Learning
C31: Metacontrol for Adaptive Imagination-Based Optimization
C32: Recurrent Environment Simulators
C33: EPOpt: Learning Robust Neural Network Policies Using Model Ensembles
Floor 3 :
W1: Lifelong Perceptual Programming By Example
W2: Neu0
W3: Dance Dance Convolution
W4: Bit-Pragmatic Deep Neural Network Computing
W5: On Improving the Numerical Stability of Winograd Convolutions
W6: Fast Generation for Convolutional Autoregressive Models
W7: The Preimage of rectifier network activities
W8: Training Triplet Networks with GAN
W9: On Robust Concepts and Small Neural Nets
W10: PI@ntNet app in the era of deep learning
W11: Exponential Machines
W12: Online Multi-Task Learning Using Biased Sampling
W13: Online Structure Learning for Sum-Product Networks with Gaussian Leaves
W14: A Theoretical Framework for Robustness of (Deep) Classifiers against Adversarial Samples
W15: Compositional Kernel Machines
W16: Loss is its own Reward: Self-Supervision for Reinforcement Learning
W17: REBAR: Low-variance, unbiased gradient estimates for discrete latent variable models
W18: Precise Recovery of Latent Vectors from Generative Adversarial Networks
W19: Arbitrary Style Transfer in Real-time with Adaptive Instance Normalization

Wednesday Morning (April 26th, 10:30am to 12:30pm)

C1: Deep Multi-task Representation Learning: A Tensor Factorisation Approach
C2: Training deep neural-networks using a noise adaptation layer
C3: Delving into Transferable Adversarial Examples and Black-box Attacks
C4: Towards the Limit of Network Quantization
C5: Towards Deep Interpretability (MUS-ROVER II): Learn. Hier. Rep. of Tonal Music
C6: Learning to superoptimize programs
C7: Regularizing CNNs with Locally Constrained Decorrelations
C8: Generative Multi-Adversarial Networks
C9: Visualizing Deep Neural Network Decisions: Prediction Difference Analysis
C10: FractalNet: Ultra-Deep Neural Networks without Residuals
C11: Faster CNNs with Direct Sparse Convolutions and Guided Pruning
C12 : Filter Shaping for CNN
C13: The Neural Noisy Channel
C14: Automatic Rule Extraction from Long Short Term Memory Networks
C15: Adversarially Learned Inference
C16: Deep Information Propagation
C17: Revisiting Classifier Two-Sample Tests
C18: Loss-aware Binarization of Deep Networks
C19: Energy-based Generative Adversarial Networks
C20: Central Moment Discrepancy (CMD) for Domain-Invariant Rep. Learning
C21: Temporal Ensembling for Semi-Supervised Learning
C22: On Detecting Adversarial Perturbations
C23: Understanding deep learning requires rethinking generalization
C24: Adversarial Feature Learning
C25: Learning through Dialogue Interactions
C26: Learning to Compose Words into Sentences with Reinforcement Learning
C27: Batch Policy Gradient Methods for Improving Neural Conversation Models
C28: Tying Word Vectors & Word Classifiers: A Loss Frame. for Language Modeling
C29: Geometry of Polysemy
C30: PGQ: Combining policy gradient and Q-learning
C31: Reinforcement Learning through Asynchronous Advantage Actor-Critic on a GPU
C32: Learning to Navigate in Complex Environments
C33: Learning and Policy Search in Stochastic Dyna. Syst. with Bayesian NN
Floor 3 :
W1: Neurogenesis-inspired dico. Learning:online model adaption in a changing world
W2: The High-Dimensional Geometry of Binary Neural Networks
W3: Discovering objects and their relations from entangled scene representations
W4: A Differentiable Physics Engine for Deep Learning in Robotics
W5: Automated Generation of Multilingual Clusters for the Eval. of Distributed Rep.
W6: Development of JavaScript-based deep learning platform, app, to dist. training
W7: Factorization tricks for LSTM networks
W8: Shake-Shake regularization of 3-branch residual networks
W9: Trace Norm Regularised Deep Multi-Task Learning
W10: Deep Learning with Sets and Point Clouds
W11: Deep Nets Don't Learn via Memorization
W12: Multiplicative LSTM for sequence modelling
W13: Learning to Discover Sparse Graphical Models
W14: Revisiting Batch Normalization For Practical Domain Adaptation
W15: Early Methods for Detecting Adversarial Images and a Colorful Saliency Map
W16: Natural Language Gen. In Dialogue using Lexicalized & Delexicalized Data
W17: Coupling Distributed and Symbolic Execution for Natural Language Queries
W18: Adversarial Examples for Semantic Image Segmentation
W19: RenderGAN: Generating Realistic Labeled Data

Wednesday Afternoon (April 26th, 4:30pm to 6:30pm)

C1: Learning recurrent representations for hierarchical behavior modeling
C2: Predicting Medications from Diagnostic Codes with Recurrent Neural Networks
C3: Sparsely-Connected Neural Networks: Towards Efficient VLSI Implementation of Deep NN
C4: HoLStep: A Machine Learning Dataset for Higher-order Logic Theorem Proving
C5: Learning Invariant Representations Of Planar Curves
C6: Entropy-SGD: Biasing Gradient Descent Into Wide Valleys
C7: Amortised MAP Inference for Image Super-resolution
C8: Inductive Bias of Deep Convolutional Networks through Pooling Geometry
C9: Neural Photo Editing with Introspective Adversarial Networks
C10: A Learned Representation For Artistic Style
C11: Learning to Remember Rare Events
C12: Optimization as a Model for Few-Shot Learning
C13: Support Regularized Sparse Coding and Its Fast Encoder
C14: Discrete Variational Autoencoders
C15: Training Compressed Fully-Connected Networks with a Density-Diversity Penalty
C16: Efficient Representation of Low-Dimensional Manifolds using Deep Networks
C17: Semi-Supervised Classification with Graph Convolutional Networks
C18: Understanding Neural Sparse Coding with Matrix Factorization
C19: Tighter bounds lead to improved classifiers
C20: Why Deep Neural Networks for Function Approximation?
C21: Hierarchical Multiscale Recurrent Neural Networks
C22: Dropout with Expectation-linear Regularization
C23: HyperNetworks
C24: Hadamard Product for Low-rank Bilinear Pooling
C25: Adversarial Training Methods for Semi-Supervised Text Classification
C26: Fine-grained Analysis of Sentence Embeddings Using Auxiliary Prediction Tasks
C27: Pointer Sentinel Mixture Models
C28: Reasoning with Memory Augmented Neural Networks for Language Comprehension
C29: Dialogue Learning With Human-in-the-Loop
C30: Zoneout: Regularizing RNNs by Randomly Preserving Hidden Activations
C31: Learning to Play in a Day: Faster Deep Reinforcement Learning by Optimality Tightening
C32: Learning Visual Servoing with Deep Features and Trust Region Fitted Q-Iteration
C33: An Actor-Critic Algorithm for Sequence Prediction
Floor 3 :
W1: Song From PI: A Musically Plausible Network for Pop Music Generation
W2: Charged Point Normalization: An Efficient Solution to the Saddle Point Problem
W3: Towards "AlphaChem": Chemical Synthesis Planning with Tree Search and DNN
W4: CommAI: Evaluating the first steps towards a useful general AI
W5: Joint Multimodal Learning with Deep Generative Models
W6: Transferring Knowledge to Smaller Network with Class-Distance Loss
W7: Regularizing Neural Networks by Penalizing Confident Output Distributions
W8: Adversarial Attacks on Neural Network Policies
W9: Generalizable Features From Unsupervised Learning
W10: Compact Embedding of Binary-coded Inputs and Outputs using Bloom Filters
W11: Semi-supervised deep learning by metric embedding
W12: Changing Model Behavior at Test-time Using Reinforcement Learning
W13: Variational Reference Priors
W14: Gated Multimodal Units for Information Fusion
W15: Playing SNES in the Retro Learning Environment
W16: Unsupervised Perceptual Rewards for Imitation Learning
W17: Perception Updating Networks: Arch. Constraints for interpretable video gen. Models
W18: Adversarial examples in the physical world