Track 1	Track 2	Track 3	Track 4	Track 5
Monday				
Breakfast, catered lunch (8:00 - 9:00)				
Morning Session 1 (9:00 - 10:20)				
Bandit Algorithms	Information Theory 1	Next Generation Wireless Communications: 6G	Federated Learning	Distributed Inference and Task-Oriented Communications (I)
Chair: Christina Fragouli	Chair: Aaron Wagner	Chair: Tingfang Ji	Chair: Cong Shen	Chair: Yuval Cassuto
Compression for Multi-Armed-Bandits, Christina Fragouli	Characterizing the generalization error of Gibbs algorithms using information measures, Yuheng Bu	Selected topics in AI/ML over 5G-Advanced air interface and toward 6G, Taesang Yoo	Gradient-based Clustered Federated Learning for CSI Compression, Hyeji Kim	Reliable distributed inference with erroneous classifiers, Yuval Cassuto
Learning Infinite-Horizon Average-Reward Restless Multi-Action Bandits via Index Awareness, Jian Li	Entropic Causal Inference and Approximate Minimum Entropy Coupling, Murat Kocaoglu	6G Radio Technologies, Tingfang Ji	Differentially Private Wireless Federated Learning Using Orthogonal Sequences, Cong Shen	Semantic-aware communication-resource allocation for distributed inference, Yongjune Kim
Causal Bandits, Ali Tajer	When does feedback increase the capacity of queues?, Sahasranand KR	On the Role of Computer Vision in Millimeter and Terahertz Communications for 6G, Byonghyo Shim	Towards Understanding and Mitigating Dimensional Collapse in Heterogeneous Federated Learning, Vincent Tan	Distributed Deep Neural Networks, Ivan Bajic
Variance-Adaptive Algorithm for Probabilistic Maximum Coverage Bandits with General Feedback, Jinhang Zuo	The Minimax Redundancy of Lossy Compression, Aaron Wagner			
Break, refreshments (10:20 - 10:40)				
Morning Session 2 (10:40-12:00 AM)				
Approximate Message Passing, Graph Codes and Causality	Coding Theory 1	Machine Learning for Wireless Communication	Privacy: Information Retrieval and Computaion	Robust Machine Learning
Chair: Gregory Pottie	Chair: Hessam Mahdavifar	Chair: Wei Yu	Chair: Hulya Seferoglu	Chair: Mayank Bakshi
Graph codes for high-dimensional QKD, Lara Dolecek	Coding for secure communication in distributed wireless networks: a new approach, Hessam Mahdavifar	ProductAE: Toward designing AI-driven error-correction codes of large dimensions, Mohammad Vahid Jamali	Model Privacy: A Statistical Framework to Understand Model Stealing Attack and Defense, Jie Ding	Authentication against limited-view adversaries, Mayank Bakshi
Towards composable distributions of latent space augmentations, Gregory Pottie	Bounds on the density of smooth lattice coverings, Or Ordentlich	State-Augmented Algorithms for Wireless Resource Management with Graph Neural Networks, Navid Naderi Alizadeh	Private Subset Retrieval, Alex Sprintson	How to efficiently catch Byzantine adversaries in distributed gradient descent, Rawad Bitar
Precise asymptotics for spectral methods in mixed Generalized Linear Models, Yihan Zhang	Deletion-detecting codes applied to coded trace reconstruction, Serge Kas Hanna	Machine learning methods for millimeter wave channel modeling, Sundeep Rangan	Efficient coded multi-party computation at edge networks, Hulya Seferoglu	Learning robust neural networks of image reconstruction, Reinhard Heckel
	Concatenated codes for DNA storage, Alexandre Graell i Amat	Active Learning for Communications and Sensing, Wei Yu	Individually-Private Linear Computation, Anoosheh Heidarzadeh	Shaping deep neural networks for robustness: insights from communication theory and neuroscience, Upamanyu Madhow
Know Thy Neighbour: Lighning Introductions, catered lunch (12:00-1:00)				
Plenary: Three Cheers for (and Applications of) Information Theory, Matthieu Bloch, Andrew Barron, Lara Dolecek, Emina Soljann (1:30-2:45)				
Afternoon Session 1 (3:00 - 4:20)				
Reinforcement Learning-1	Channel Coding	Quantum Coding and Information Theory	Innovations in Federated Learning	Federated Learning & Analytics (I)
Chair: Geoffrey Gordon	Chair: Elad Romanov	Chair: Henry Pfister	Chair: Hamid Jafarkhani	Chair: Peter Kairouz
Behavior Design, Geoffrey Gordon	On the generic capacity of K-user linear computation broadcast, Syed Jafar	Entanglement-assisted Quantum Reed-Muller Tensor Product Codes, Shayan Garani	Quantized asynchronous federated learning, Hamid Jafarkhani	Foundations of Collaborative AI: from Algorithms to Platforms to Web3 Marketplace Design, Salman Avestimehr
DeepTOP: Deep Threshold-Optimal Policy for MDPs and RMABs, I-Hong Hou	Variable-length coding over binary asymmetric channels with full feedback, Hengjie Yang	Belief-Propagation with Quantum Messages for Polar Codes on Classical-Quantum Channels, Henry Pfister	Decentralized Federated Learning via Non-Coherent Over-the-Air Consensus, Nicolò Michelusi	A Statistical Framework for Personalized Federated Learning and Estimation, Suhas Diggavi
Towards Decentralized Predictive Quality of Service in Next-Generation Vehicular Networks, Michele Zorzi	On the Role of Channel Capacity in Learning Gaussian Mixture Models, Elad Romanov	Communication complexity of entanglement assisted multi-party computation, Aditya Ramamoorthy	Coding for straggler mitigation in federated learning, Eirik Rosnes	The communication cost of security and privacy in federated frequency estimation, Ayfer Ozgur
			SABRE: Safe Bayesian Peer-to-Peer Federated Learning, Nasimeh Heydaribeni	Federated Learning in Networks: Challenges and Algorithms, Kamalika Chaudhuri
Break, light refereshments (4:20 - 4:30)				
Alumni in Industry: Academia - Industry Collaborations, Christina Fragouli, Tingfang Ji, Peter Kairouz, Sujay Sanghavi, Sreeram Kannan (4:30 - 5:30)				
Reception, food (5:30 - 6:00)				
Tuesday				
Breakfast (8:00 - 9:00)				
Morning Session 1 (9:00 - 10:20)				
Low-Density Parity Check (LDPC) Codes	Information Theoretic Cryptography	Machine Learning	Deep Neural Network Analysis	Estimation and Detection
Chair: Bane Vasic	Chair: Vinod Prabhakaran	Chair: Ivana Maric	Chair: Sujay Sanghavi	Chair: Soheil Mohajer
On Error Propagation Mitigation in Sliding Window Decoding of Spatially Coupled LDPC Codes, David Mitchell	Key-cast over networks, Michael Langberg	Fundamental limits of conditional computation in neural networks, Erdem Koyuncu	Analysis of Deep Neural Subnetwork for Continual Learning and Adversarial Robustness, Salimeh Yasaei Sekeh	A Distributed Framework for Fake News Detection, Soheil Mohajer
Rate-adaptive LDPC codes for high-throughput satellite optical links, Gianluigi Liva	Retractable commitment, Matthieu Bloch	Machine learning for early prediction of pregnancy-related outcomes, Ivana Maric	Studying generalization in deep neural networks, Mehul Motani	Conditional Inference for Communication System Design, Dongwoon Bai
Some results on short block protograph LDPC codes, Dariush Divsalar	Solving the byzantine generals problem using dependent observations, Vinod Prabhakaran	Privacy preserving deep convolutional neural networks over fully homomorphic encryption, Jong-Seon No	Understanding Self-Distillation, Sujay Sanghavi	Learning-Based Channel Estimation and Detection in Space Optical Communications, Zouheir Rezki
Trapping Sets of Quantum Low-Density Parity Check Codes, Bane Vasic	The role of Mrs. Gerber's lemma for evaluating the information leakage of masked computations, Olivier Rioul	Unlimited-range URLLC: Semantic inference approach, Seong-Lyun Kim		Theory and Application of DINE: Directed Information Neural Estimation, Haim Permuter

Break, refreshments (10:20 - 10:40)				
Morning Session 2 (10:40-12:00)				
Error Correcting Codes and List Decoding	Compression, Communication and Capacity	Data-Centric, Online, and Self-Supervised Learning	Optimization and Robustness	Multiclass and Robust Classification
Chair: Ido Tal	Chair: Jeffrey Andrews	Chair: Frederic Sala	Chair: Navid Azizan	Chair: Anant Sahai
Improved Decoding of Expander Codes, Xin Li	Unreliable Relay Channels, Yossef Steinberg	Data Valuation Without Training of a Model, Hye Won Chung	Fundamental Limits of Two-layer Autoencoders and Achieving Them with Gradient Methods, Marco Mondelli	A converse for asymptotic overparameterized interpolating min-norm multiclass classification, Anant Sahai
Linear Independent Sequence Coding for Reliable Communicate, Anders Nilsson Plymoth	The DoF region of the 3-User MIMO Broadcast Channel, Mahesh Varanasi	Optimal competitive control and regret-optimal control with weights, Oron Sabag	Generalized-Smooth Nonconvex Optimization is As Efficient As Smooth Nonconvex Optimization, Yi Zhou	On the consistency of multilabel classification trees for approximate nearest neighbor search, Teemu Roos
Stronger Polarization for the Deletion Channel, Ido Tal	Coverage and Capacity of Joint Communication and Sensing Networks, Jeffrey Andrews	Weak Supervision: A Learning Paradigm for All Seasons, Frederic Sala	One-Pass Learning via Bridging Orthogonal Gradient Descent and Recursive Least-Squares, Navid Azizan	Robust Classification under Sparse Adversarial Attacks, Payam Delgosha
CRC-aided list decoding of convolutional and polar codes for short messages, Richard Wesel	Gaussian Primitive Diamond Channel: Correlated Noise and Applications, Shlomo Shamai	When self-supervised learning recovers spectral embedding: what's next?, Randall Balestriero	Robustness in deep learning: The good (width) the bad (depth) and the ugly (initialization), Volkan Cevher	Statistical span properties of binary run sequences, Hong-Yeop Song
Information Theory Society Town Hall, Matthieu Bloch, IT Society President, Icatered lunch (12:00 - 1:20)				
Break (1:20 - 1:30)				
Plenary - Coding compexlity, Venkatesan Guruswami, Mahdi Cheraghchi, Sivakanth Gopi (1:30 - 2:45)				
Break, light refrehsments (2:45 - 3:00)				
Afternoon Session 1 (3:00 - 4:20)				
Statistics	Compression, Prediction and Information Inequalities	Communication Theory	Coding Theory-2	Statistical Learning and Random Sequences
Chair: Paul Valiant	Chair: Vincent Tan	Chair: Besma Smida	Chair: Sivakanth Gopi	Chair: Christopher Tosh
Correlation-adjusted debiasing: Debiasing the Lasso with inaccurate covariate model, Michael Celentano	Multiple and hierarchical universality, Meir Feder	A feedback aided irregular repetition slotted ALOHA scheme, Tolga M. Duman	Low-complexity near-optimum decoding of G_N-coset codes with dynamic frozen bits, Mustafa Cemil Coskun	Adaptive Power Method: Eigenvector Estimation from Partial Observations, Ilan Shomorony
Improving Pearson's chi-squared test: hypothesis testing of distributions, Paul Valiant	Sequential prediction under log-loss: stochastic and smoothed adversaries, Alankrita Bhatt	Downlink/uplink localization at mmWave bands: learning modelling or both?, Nuria Gonzalez-Prelcic	Partial Geometries Constant Weight and LDPC Codes, Shu Lin	Learning Continuous Exponential Families Beyond Gaussian, Andrey Lokhov
On a phase transition in general order spline regression, Fang Han	Universal Ensembles for Sample-wise Lossy Compression, Neri Merhav	FD Integrated Sensing and Communication, Besma Smida	Long common subsequences between bit-strings and the zero-rate threshold of deletion-correcting codes, Venkatesan Guruswami	Simple and near-optimal algorithms for hidden stratification and multi-group learning, Christopher Tosh
Sampling form the SK model and the binary perceptron, Ahmed El Alaoui	On the behaviour of entropy along semi-groups: a summary of recent results, Chandra Nair	Integrated sensing and communications (ISAC) - A RadBackCom approach, Xiaodong Wang	Improved Maximally Recoverable LRCs using Skew Polynomials, Sivakanth Gopi	
Break, refreshments (4:20 - 4:40)				
Afternoon Session 2 (4:40 - 5:40)				
Blockchains	Coding Theory-3	Wireless Communication	Domain Generalization, Nonparametirc Estimation and Fair Inference	Quantum: Annealing, Detection and Learning
Chair: Sriram Vishwanath	Chair: Mahdi Soleymani	Chair: David Love	Chair: Jimmy Foulds	Chair: Emina Soljanin
Healthcare Blockchain Applications and Use Cases, Tsung-Ting Kuo	Efficient Estimation of Probability Mass Functions from Partially Observed Data, Martin Haardt	Orthogonal Chirp-Division Multiplexing Radar and its Applications, Ramesh Annavajjala	Barycentric-alignment and reconstruction loss minimization for domain generalization, Shuchin Aeron	High-quality Thermal Gibbs Sampling with Quantum Annealing Hardware, Marc Vuffray
A Mean-Field Approach to Cryptocurrency Mining in the Presence of a Rational Adversary, Sriram Vishwanath	On Kronecker separability of multiway covariance, Dogyoon Song	Received Power Maximization with Reconfigurable Intelligent Surfaces, Ender Ayanoglu	Beating the nonparametric curse of dimensionality using multi-view density estimators, Robert Vandermeulen	On optimal discrimination between quantum states with application to coherent optical detection., Yuval Kochman
	Quantitative Group Testing via Plotkin Construction, Mahdi Soleymani	Robust Feedback Coding Using a Data-Driven Architecture, David Love	Fair Inference for Discrete Latent Variable Models, James Foulds	Reconciling (Quantum) Jitters and Memory, Emina Soljanin
Break (5:40 - 6:30) on your own				
Valentine's Day Socilal, Fun Games and Team Building, food served (6:30 - 7:30)				
Wednesday				
Breakfast (8:00 - 9:00)				
Grad Talks 1	Grad Talks 2	Grad Talks 3		
Chair: TBA	Chair: Sivakanth Gopi	Chair: Oron Sabag		
An exponential improvement on the memorization capacity of deep threshold networks, Shashank Rajput	Robustness of gradient-based methods for data-driven decision making, Hesameddin Mohammadi	Contexts can be Cheap: Solving Stochastic Contextual Bandits with Linear Bandit Algorithms, Osama Habib		
Do Neural Networks Compress Manifolds Optimally?, Sourbh Bhadane	Model projection: theory and applications to multi-class fair machine learning, Wael Alghamdi	Offline Reinforcement Learning: Towards Optimal Sample Complexity and Distributional Robustness, Laixi Shi		
Graph Neural Networks Are More Powerful Than we Think, Charilaos Kanatsoulis	Light-Weight Sequential SBL Algorithm: An Alternative to OMP, Rohan Ramchandra Pote	Instance-Adaptive and Optimal Offline Reinforcement Learning, Ming Yin		
A Primal-Dual Framework for Transformers and Neural Networks, Tan Nguyen	Quickest Inference of Network Cascades with Noisy Information, Anirudh Sridhar	Information-Theoretic Characterization of the Generalization Error for Iterative Semi-Supervised Learning, Haiyun He		
Break, refreshments (10:20 - 10:40)				
Grad Talks 4	Grad Talks 5	Grad Talks 6	Quantum Computing and Algorithms	
Chair: Olivier Rioul	Chair: Venkat Guruswami	Chair: Rashmi Vinayak	Chair: Daniel Grier	

Maximal Dissent: a State-Dependent Way to Agree in Distributed Convex Optimization, Ashwin Verma	Parallel Monte Carlo Markov chain decoding of linear codes, Jiun-Ting Huang	Simple Binary Hypothesis Testing: Locally-Private and Communication-Efficient, Ankit Pensia	Advantages and Limitations of the Quantum Approximate Optimization Algorithm, Leo Zhou
Communication-Efficient and Privacy-Preserving Machine Learning, Antonious Girgis	Majority is not Required: a Rational Analysis of the Private Double-Spend Attack from a Sub-Majority Adversary, Yanni Georghiades	A Statistical Learning Theory for High Dimensional Interpolants and Regularized Estimators, Lijia Zhou	Sample-optimal classical shadows for pure states, Daniel Grier
Feature learning in neural networks and kernel machines that recursively learn features, Adityanarayanan Radhakrishnan	Convertible codes: enabling efficient conversion of coded data in distributed storage, Francisco maturana	Tree Learning: Optimal Sample Complexity and Algorithms, Dmitrii Avdiukhin	
Towards trustworthy intelligent online decision-making, Jianyi Yang	Designing Deployable RF Sensing systems, Sai Roshan Ayyalasomayajula	Online k-means Clustering on Arbitrary Data Streams, Robi Bhattacharjee	

Graduation Day Outlinde

Graduation Day is ITA's venue for outstanding graduating students and posdocs seeking research positions to present their work.
Past GD articipants have gone on to prominent academic and industrial research positions.
There are two types of presentations, talks, each lasting 20 minutes, and posters, presented during a 2-hour lunch in a scenic setting.
This year, nine senior judges will award twelve presentation prizes, based on the quality and relevance of the work and clarity of presentation, but ignoring whether the presentation is a poster or a talk.
Please note that non-graduating students can present at Friday's student poster session, with prizes awarded there too.

Graduation Day Posters, catered lunch (12:00 - 2:00)

Yulu Jin, Yuhan Liu, Neophytos Charalambides, Ville Hyvönen, Francisco Maturana, Sebin Gracy, Ming Yin, Umut Demirhan, Debarnab Mitra, Yahya Sattar, Sudeep Salgia, Tao Jiang, Ish Jain
Manideep Dunna, Jacob Imola, Venkatesh Sathyanarayanan, Yifan Wu, Robin Rajamäki, Parthe Pandit, Amrita Roy Chowdhury
Posters will be presented on an A0 paper, 4 feet wide by 3 feet tall.

Plenary: **Quantum Computing**, Umesh Vazirani, Geoff Penington, Henry Yuen (2:00 - 3:30)

Break, light refreshments (3:30 - 3:40)

Start Your Own Open-AI, Experienced Experts Discuss Starting Tech Companies, Quinn Li, Meir Feder, Sriram Vishwanath (3:40 - 4:40)

Banquet, ticket required (6:00 - 9:30)

Track 1	Track 2	Track 3	Track 4	Track 5	Track 6
Thursday					
Breakfast (8:00 - 9:00)					
Morning Session 1 (9:00 - 10:20)					
Fairness and Robustness in Machine Learning-1 (I)	Online Learning	Information Theory-2	Learning and Control 1 (I)	Diffusion Models and Distributed Computations	Efficient Algorithms and Lower bounds
Chair: Ahmad Beirami	Chair: Nikhil Karamchandani	Chair: Jonathan Ponniah	Chair: Jorge Poveda	Chair: Kangwook Lee	Chair: Daniel Hsu
Probable domain generalization via quantile risk minimization, Hamed Hassani	Near-optimal packet scheduling with end-to-end deadline constraints in multihop networks, Javad Ghaderi	Information theory in graphs and network alignment, Wayne B. Hayes	Minimum Bitrate Neuromorphic Encoding for Continuous Time Series: A Lower Bound, Takashi Tanaka	Compressed Sensing MRI using Diffusion Generative Models, Ajil Jalal	Estimating the Longest Increasing/Common Subsequence, Alex Andoni
Improving robustness via safe data augmentation, Yao Qin	Identifying the largest communities via random sampling, Nikhil Karamchandani	Second order rate regions for the scalar Gaussian broadcast channel, Daniela Tuninetti	Can direct latent model learning solve linear quadratic Gaussian control?, Kaiqing Zhang	Score-based generative modeling secretly minimizes the Wasserstein distance, Kangwook Lee	Near-optimal statistical query lower bounds for agnostically learning intersections of halfspaces with Gaussian marginals, Daniel Hsu
Multigroup fairness loss minimization and indistinguishability, Parikshit Gopalan	Learning While Scheduling in Multi-Server Systems With Unknown Statistics: MaxWeight with Discounted UCB, Lei Ying	A Framework for Information Theoretic Converses, Jonathan Ponniah	Federated learning and control: Identification & synthesis, James Anderson	Denoising Diffusion Models from I-MMSE Relations, Greg Ver Steeg	Universality of High-Dimensional Estimation and Approximate Message Passing with Nearly Deterministic Matrices, Yue Lu
A data-centric view on reliable generalization: From ImageNet to LAION-5B, Ludwig Schmidt	Stochastic Contextual Bandits with Long Horizon Rewards, Yuzhen Qin	Joint Coding of URLLC and eMBB in the Finite Blocklength Regime, Homa Nikbakht	Convergence and sample complexity of gradient methods for the model-free linear quadratic regulator problem, Mihailo Jovanovic	Age of information: computing updates from updates, Roy Yates	AI-based video compression to accelerate development and deployment, Amir Said
Break, refreshments (10:20 - 10:40)					
Morning Session 2 (10:40 - 11:40)					
Tools from Aoi, Network Science and SIT to Address Online Misinformation (I)	Game Theory: Spectrum Sharing, Reinforcement Learning and Matching Markets	Optimal Transport	EnCORE Session on Distributed Learning & Decision Making 1	Domain Adaptation, Age of Information and Inference Control	Physical Layer Privacy, Secure Communication and Wireless Communication
Chair: Melih Bastopcu	Chair: Roy Fox	Chair: Ziv Goldfeld	Chair: Arya Mazumdar	Chair: George Atia	Chair: Urbashi Mitra
The role of gossiping for information dissemination over networked agents, Melih Bastopcu	Towards a Fundamental Understanding of Advanced Wireless Spectrum Sharing Policies in Next Generation Networks, Alhussein Abouzeid	Gromov-Wasserstein Distances: Entropic Regularization Duality and Sample Complexity, Ziv Goldfeld	Leveraging Bilevel Optimization for Better Distributed Learning, Tianyi Chen	A distributionally robust approach to domain adaptation, George Atia	Statistical Hardness for Private Localization, Urbashi Mitra
Private Status Updating with Erasures: A Case for Retransmission Without Resampling, Ahmed Arafa	Decentralized Learning Algorithms for Competition in Matching Markets, Eric Mazumdar	Normalizing Flows and Knoethe-Rosenblatt Couplings, Zaid Harchaoui	Your prompt attention is needed: on the role of attention in prompt tuningt, Mahdi Soltanolkotabi	Deadline-Aware Distributed Computations with Multiple Resolutions, Homa Esfahanizadeh	Secure communication over 1-2-1 networks, Martina Cardone
Information Freshness for Real-time Inference and Estimation: A Few Recent Results, Yin Sun	Better Together: Population-Based Methods for Single- and Multi-Agent Reinforcement Learning, Roy Fox	Entropy Regularized Optimal Transport Independence Criterion, Lang Liu	Personalized Asynchronous Federated Learning, César A Uribe	Control with multiplicative observation noise: a new MAP-based achievable strategy, Gireeja Ranade	A wideband generalization of the near-field region for extremely large phased-arrays, Robert Heath
NSF Opportunities, Alhussein Abouzeid and Alfred Hero	Hello Boss, students and postdocs meet internship mentors, postdoc advisors, and employers		Games	Lunch on your own (11:40 - 1:40)	
Plenary: Deep Learning Theory (1:40 - 2:55)					
Break, light refreshments (2:55 - 3:10)					
Afternoon Session 1 (3:10 - 4:30)					
Differential Privacy-1 (I)	Machine Learning and Applications	Machine Learning and Neural Nets	Robust Learning (I)	Coding Theory 4	Statistical Learning Theory for Control
Chair: Ananda Theertha Suresh	Chair: Peter Gerstoft	Chair: Rayan Saab	Chair: Lalitha Sankar	Chair: Brian Kurkoski	Chair: Nikolai Matni
Regression with Label Differential Privacy, Badih Ghazi	Machine learning for acoustic source localization, Peter Gerstoft	Universal Architectures for Progressive Machine Learning: Model, Performance Evaluation and Applications, John Baras	Realistic estimates of machine learning models classification accuracy from overoptimistic published results, Pouria Saidi	Derivative Descendants of Cyclic Codes and Derivative Decoding, Qin Huang	Independent learning from sequential data, Stephen Tu
Stronger Privacy Amplification by Shuffling for Rényi and Approximate Differential Privacy, Vitaly Feldman	Most resource-efficient constant matrix-vector multiplication on FPGAs, Ralf Muller	Multi-modal prototype learning for interpretable multivariable time series classification, Reza Abbasi-Asl	Smoothly Giving up: Robustness for Simple Machine Learning Models, Lalitha Sankar	CRC-Enabled Lattices for Multiuser Communication, Brian Kurkoski	Transformers as Algorithms: Generalization and Stability in In-context Learning, Samet Oymak

TBA, Vahab Mirrokni	Access differences in fair resource allocation, Mesrob Ohannessian	Looped Transformers are Computer and Prompts are their Programs, Dimitris Papailiopoulos	Learning from Diverse Data in Metric and Preference Learning, Ramya Korlakai Vinayak	The cost of exactness in coding for reliable computing, Mohammad Ali Maddah-Ali	Online switching control with stability and regret guarantees, Yingying Li
Recent Advances in Differentially Private Training of Large Language Models, Janardhan Kulkarni	Error Correction Codes for Testing Chemical/Biological Agents: Using Pooled Testing to Increase Test Reliability, Weiyu Xu	Algorithms for quantizing neural networks with theoretical guarantees, Rayan Saab	Target Conditioned Representation Independence (TCRI); From Domain-Invariant to Domain-General Representations, Sanmi Koyejo	Disk-adaptive coding for distributed storage systems, Rashmi Vinayak	TaSIL: Taylor Series Imitation Learning, Nikolai Matni
Break, refreshments (4:30 - 4:50)					
Afternoon Session 2 (4:50 - 5:50)					
Graphs: Manifold Learning, Compression and Combinatorial Statisitics	Reinforcement Learning-2	Deep Learning	EnCORE Session on Distributed Learning & Decision Making 2	Coding Theory 5	Learning and Control-2 (I)
Chair: Yihong Wu	Chair: Lin Yang	Chair: Pratik Chaudhari	Chair: Aritra Mitra	Chair: Aria Nosratinia	Chair: Yang Zheng
Convergence of Gaussian kernelized graph Laplacian with bi-stochastic normalization, Xiuyuan Cheng	A Policy Gradient Algorithm for the Risk-Sensitive Exponential Cost MDP, Mehrdad Moharrami	Explaining Deep Protein Models in the Spectral Domain, Amirali Aghazadeh	Safe Dual Gradient Method for Resource Allocation over Networks, Mahnoosh Alizadeh	Exploring Properties of Sparse Regression LDPC Codes, JF Chamberland	Optimal Linear-Time Algorithms in Resource Allocation Problems, Jason Marden
Compression with Unlabeled Graph Side Information, Shirin Saeedi Bidokhti	Near-Optimal Sample Complexity Bounds for Constrained MDPs, Lin Yang	A Picture of the Prediction Space of Deep Networks, Pratik Chaudhari	TBA, Tamer Basar	Polar Codes for Low-SNR Regime, Fariba Abbasi	Achieving Logarithmic Regret via Hints in Online Learning of Noisy LQR Systems, Bahman Gharesifard
Random graph matching at Otter's threshold via counting chandeliers, Yihong Wu			Mitigating the Risk Associated with Epistemic and Aleatory Uncertainties in MDPs, Mohammad Ghavamzadeh	An LLR Metric for Empirical Performance Verification of Wiretap Codes, Aria Nosratinia	Certified and Robust Forward Invariance in Neural ODEs, Yuanyuan Shi
					Policy Gradient Algorithm for Multi-Objective Reinforcement Learning, Dileep Kalathil
Friday					
Breakfast (8:00 - 9:00)					
Morning Session 1 (9:00-10:20)					
Privacy and Security in Machine Learning (I)	Fairness and Robustness in Machine Learning-2 (I)	Trustworthy Machine Learning	EnCORE Session on Distributed Learning & Decision Making 3	Distributed Learning	Statistics and Learning Theory
Chair: Kamalika Chaudhuri	Chair: Meisam Razaviyayn	Chair: Parinaz Naghizadeh	Chair: AbolFazl Hashemi	Chair: Yu Xiang	Chair: Yizhe Zhu
Measuring data leakage in machine-learning models with Fisher information, Chuan Guo	Fairness in matching under uncertainty, Vatsal Sharan	LightVeriFL: A Lightweight and Verifiable Secure Aggregation for Federated Learning, Baturalp Buyukates	Regret-Optimal Control, Babak Hassibi	Sample Complexity of Q-learning: from Single-agent to Federated Learning, Yuejie Chi	A unified framework for correlation mining in ultra-high dimension, Alfred Hero
Schrödinger's cactus: optimal differential privacy mechanisms in the large-composition regime, Oliver Kosut	Predictive Multiplicity in Machine Learning, Flavio Calmon	Adaptive data debiasing through bounded exploration, Parinaz Naghizadeh	Byzantine Fault-Tolerance in Federated Local SGD under 2f-Redundancy, Thinh Doan	Experimental Design Networks: A Paradigm for Serving Heterogeneous Learners under Networking Constraints, Edmund Yeh	Binary Iterative Hard Thresholding for 1-Bit Compressed Sensing, Arya Mazumdar
Differentially Private Community Detection over Stochastic Block Models, Ravi Tandon	A Unifying Theory of Distance from Calibration, Jaroslaw Blasiok	Demystifying and Mitigating Bias for Graph Node Representation Learning, Yanning Shen	FedAvg with Fine Tuning: Local Updates Lead to Representation Learning, Aryan Mokhtari	On Large-Scale Multiple Testing Over Networks: An Asymptotic Approach, Yu Xiang	Minimum Description Length Interpolation Learning, Nathan Srebro
The curse of overparametrization in adversarial training, Adel Javanmard	A stochastic optimization framework for fair risk minimization, Ahmad Beirami	Fairness without imputation: A decision tree approach for fair prediction with missing values, Haewon Jeong	An Optimal Transport Approach to Federated Learning, Ramtin Pedarsani	Sparse-Sign-SGD with Majority Vote for Communication-Efficient Distributed Learning, Namyoon Lee	Overparameterized random feature regression with nearly orthogonal data, Yizhe Zhu
Break, refreshments (10:20 - 10:40)					
Morning Session 2 (10:40 - 12:00)					
Differential Privacy-2(I)	Game Theory-1	Optimization	Robust Learning in Presence of Adversaries	Optimal Algorithms and Generative Modelling	Control Theory and Reinforcement Learning
Chair: Hilal Asi	Chair: Randall Berry	Chair: Hassan Mansour	Chair: Amin Karbasi	Chair: Farzad Farnoud	Chair: Srinivas Shakkottai
Federated Learning with Formal User-Level Differential Privacy Guarantees, Abhradeep Thakurtha	Comparison of Information Structures in Bayesian Social Learning, Randall Berry	Feature Learning with Gradient Descent, Jason Lee	Learning and Certification under Instance-targeted Poisoning, Amin Karbasi	Data-free Generative Modeling and Its Applications, Mostafa El-Khamy	Optimal Error-Free Gradecast, Jinyuan Chen
Stochastic Differentially Private and Fair Learning, Meisam Razaviyayn	The Price of Anarchy is Fragile in Communication-Denied Games, Philip Brown	Taming Fat-Tailed ("Heavier-Tailed" with Potentially Infinite Variance) Noise in Federated Learning, Jia Liu	Sparse M-estimation of DOA, Christoph Mecklenbrauker	High-Dimensional Geometric Streaming in Polynomial Space, David Woodruff	Competitive control, Gautam Goel
Mean Estimation with User-level Privacy under Data Heterogeneity, Audra McMillan	Cooperation for Better Competition: A Sequential Game Model of Flocking, Mingyan Liu	Deep Proximal Gradient Method for Learned Convex Regularizers, Hassan Mansour	On the Complexity of Adversarial Decision Making, Ayush Sekhari	Optimal testing of discrete distributions with high probability, Themis Gouleakis	A mean-field theory for learning Schoenberg measure of radial basis functions, Masoud Badiei Khuzani
Subset-Based Instance Optimality in Private Estimation, Ananda Theertha Suresh	Reveal or conceal? Signaling capabilities in General Lotto games, Keith Paarporn	Symmetry Teleportation for Accelerated Optimization, Rose Yu	Trustworthy Machine Learning in Complex Environments, Furong Huang	Active Ranking without Strong Stochastic Transitivity, Farzad Farnoud	Structured Reinforcement Learning for Media Streaming at the Wireless Edge, Srinivas Shakkottai
Student Posters, catered lunch (12:00 - 14:00) Jianyi Yang, Yibo Yang, Justin Kang, Xuchuang Wang, Yi Chen, Anthony Thomas, Harit Vishwakarma, Wei-Ning Chen, Simeng Zheng, Gokcan Tatli, Zixiao Zong, Yigit Efe Erginbas, Yinbin Ma, Zhichao Wang, Nathan Stromberg, Atefeh Gilani Travis Cuvelier, Taylan Kargin, Md Kamran Chowdhury Shisher, Sk Miraj Ahmed, Zhichao Wang, Chengshuai Shi, Ahmed Maksud, Yingcong Li, Berivan Isik, Mine Dogan, Tasmeen Zaman Ornee, Cuneyd Ozturk, Yiyun He, Xiaohan Fu, Daniil Vankov, Monica Welfert, Chhavi Yadav, Jesse Friedbaum					
Towards trustworthy intelligent online decision-making, Jianyi Yang					
Towards Empirical Sandwich Bounds on the Rate-Distortion Function, Yibo Yang					
The Fair Value of Data Under Heterogeneous Privacy Constraints, Justin Kang					
Achieve Near-Optimal Individual Regret & Low Communications in Multi-Agent Bandits, Xuchuang Wang					
Crowdsourced Clustering via Active Querying: Practical Algorithm with Theoretical Guarantees, Yi Chen					
A Theoretical Perspective on Hyperdimensional Computing, Anthony Thomas					
Good Data from Bad Models: Foundations of Threshold-based Auto-labeling}, Harit Vishwakarma					
The Fundamental Price of Secure Aggregation in Differentially Private Federated Learning, Wei-Ning Chen					
Code-Aware Storage Channel Modeling via Machine Learning, Simeng Zheng					
Learning Preference Distributions From Distance Measurements, Gokcan Tatli					
Privacy by Projection: Federated Population Density Estimation by Projecting on Random Features, Zixiao Zong					
Interactive Learning with Pricing for Optimal and Stable Allocations in Markets, Yigit Efe Erginbas					

On Coded Caching Systems with Offline Users, Yinbin Ma					
High-Dimensional Asymptotics of Feature Learning in the Early Phase of Neural Network Training, Zhichao Wang					
Smoothly giving up: robustness for simple models, Nathan Stromberg					
An Alphabet of Leakage Measures, Atefeh Gilani					
Time-invariant prefix coding for MIMO LQG control, Travis Cuvelier					
Thompson sampling achieves O(√T) regret in linear quadratic control, Taylan Kargin					
How Does Data Freshness Affect Real-time Supervised Learning?, Md Kamran Chowdhury Shisher					
Unsupervised Multi-Source Domain Adaptation Without Access to Source Data., Sk Miraj Ahmed					
High-Dimensional Asymptotics of Feature Learning in the Early Phase of Neural Network Training, Zhichao Wang					
Reward teaching for federated multi-armed bandits, Chengshuai Shi					
Secret Key Generation from MIMO Channel With or Without Reciprocity, Ahmed Maksud					
Provable Pathways: Learning Multiple Tasks over Multiple Paths, Yingcong Li					
An Information-Theoretic Justification for Model Pruning, Berivan Isik					
Resilient transmission and scheduling mechanisms for mmWave networks, Mine Dogan					
A Whittle Index Policy for the Remote Estimation of Gauss-Markov Processes through Multiple Channels, Tasmeen Zaman Ornee					
Downlink Throughput Analysis of Mega-Constellations, Cuneyd Ozturk					
Algorithmically Effective Differentially Private Synthetic Data, Yiyun He					
Murfe: Multimodal Biometric Fuzzy Extractor using Lattices, Xiaohan Fu					
Last Iterate Convergence of Popov Method for Non-monotone Stochastic Variational Inequalities, Daniil Vankov					
α-GAN: Convergence and Estimation Guarantees, Monica Welfert					
A Learning-Theoretic Framework for Certified Auditing with Explanations, Chhavi Yadav					
Plenary: Machine Learning Applications (2:00-3:15)					
Break, light refreshments (3:15 - 3:30)					
Afternoon Session 1 (3:30 - 4:50)					
Information-Theoretic Tools for Trustworthy ML(I)	Game Theory-2	Deep Learning and Optimization	EnCORE Session on Distributed Learning & Decision Making 4	Creative AI and Improvisation Control	
Chair: Flavio Calmon	Chair: Marcos Vasconcelos	Chair: Yao Xie	Chair: Aritra Mitra	Chair: Shlomo Dubnov	
DIfferentially private secure multiparty multiplication: hiding information in the rubble of noise, Viveck Cadambe	On Almost-Bayesian Persuasion, Cedric Langbort	Label Noise in Adversarial Training: A Novel Perspective to Study Robust Overfitting, Jingbo Shang	Robust Multi-Agent Multi-Armed Bandits, R. Srikant	Learning Interpretable Creative Representation via Physical Symmetry, Gus Xia	
When Personalization Harms Performance, Berk Ustun	A Finite-Sample Analysis of Independent Learning in Competitive Stochastic Games, Zaiwei Chen	Flat minima generalize for low-rank matrix recovery, Dmitriy Drusvyatskiy	Collaborative linear bandits with adversarial agents: near-optimal regret bounds, Aritra Mitra	Intrinsic motivation as a route to creativity, Stas Tiomkin	
Differential privacy as contraction of f-divergences, Mario Diaz	Extension theorems for General Lotto games with applications to network security, Adel Aghajan	Spatio-temporal point processes with deep non-stationary kernels, Yao Xie	On the Convergence of Differentially Private Federated Learning on Non-Lipschitz Objectives via Clipping and Normalized Client Updates, Abolfazl Hashemi	Causality Measure for co-Creative Machine Improvisation, Shlomo Dubnov	
From Robustness to Privacy and Back, Hilal Asi	Robust estimation over the collision channel in the presence of an intelligent jammer, Marcos Vasconcelos		Trajectory Entropies for Simultaneous Localisation and Navigation, Girish Nair		
Farewell Bash, Food and Goodbye (4:50 - 6:00)					