Cornell Information Theory Day (July 7, 2023)
Book of Abstracts

Speaker: Sourbh Bhadane
Advisor(s): Aaron Wagner, Jayadev Acharya
Title: Optimal Neural Network Compression and the Manifold Hypothesis
Abstract: Traditional source coding theory for stationary Gaussian sources shows near-optimality of the linear transform coding paradigm which forms the basis for existing lossy compression algorithms used in practice. However, artificial neural-network-based (ANN-based) compressors have proven to be remarkably effective at compressing sources, such as images, that are nominally high-dimensional but presumed to be concentrated on a low-dimensional manifold. To reconcile these observations, we initiate an entropy-distortion analysis of sources that are uniformly distributed on low-dimensional manifolds. We characterize the optimal one-shot compression performance tradeoffs for these sources and show that previously-optimal ANN-based compressors fail to optimally compress them. We provide insights into the difficulty that ANN-based compressors face in compressing such sources and propose a fix that involves inserting random Fourier feature embeddings to enable learning steep functions.

Speaker: Haiyun He
Advisor(s): Ziv Goldfeld, Christina Lee Yu
Title: Information-Theoretic Characterization of the Generalization Error for Learning Algorithms
Abstract: Using information-theoretic principles, we consider studying the generalization error (gen-error) of learning algorithms. Existing works have shown that for supervised learning, the mutual information between input data and output model bounds the gen-error. To move a step forward, we consider semi-supervised learning (SSL) algorithms that generate pseudo-labels for a large amount of unlabelled data to progressively refine the model parameters. Our first theoretical result suggests that when the class conditional variances are not too large, the gen-error decreases with the number of iterations, but quickly saturates. On the flip side, if the class conditional variances (and so the amount of overlap between the classes) are large, the gen-error increases with the number of iterations. To mitigate this undesirable effect, we show that regularization can reduce the gen-error. These results are corroborated by extensive experiments on the MNIST and CIFAR datasets in which we notice that for easy-to-distinguish classes, the gen-error improves after several pseudo-labelling iterations, but saturates afterwards, and for more difficult-to-distinguish classes, regularization improves the generalization performance. Our second theoretical result suggests that smaller shared information between pseudo-labeled data and true labeled data leads to smaller gen-error, which can serve as a principle for choosing pseudo-labeling methods. Some work in progress on understanding the fairness effect and deep neural networks will also be introduced.

Speaker: Adeel Mahmood
Advisor(s): Aaron B. Wagner
Title: Minimax Rate-Distortion
Abstract: We show the existence of variable-rate rate-distortion codes that meet the distortion constraint almost surely and are minimax, i.e., strongly, universal with respect to an unknown source distribution. We provide an achievable \tilde{O}(1/\sqrt{n}) redundancy rate, which we show is optimal. This is in contrast to prior work on universal lossy compression, which provides O(\log n/n) redundancy guarantees for weakly universal codes under various regularity conditions. We show that either eliminating the regularity conditions or upgrading to strong universality while keeping these regularity conditions entails an inevitable increase in the redundancy to \tilde{O}(1/\sqrt{n}). Our construction involves random coding with non-i.i.d. codewords and a zero-rate uncoded transmission scheme. The proof uses acceptance-rejection sampling and exact asymptotics from large deviations.

Speaker: Hemant K. Mishra
Advisor(s): Mark M. Wilde
Title: Asymptotic error rates for classical and quantum antidistinguishability
Abstract: The concept of antidistinguishability of quantum states has been studied to investigate foundational questions in quantum mechanics. For example, it is used to investigate the reality of quantum states, ruling out Ψ-epistemic ontological models of quantum mechanics [*Pusey et al., Nat. Phys., 8(6):475-478, 2012*]. The emerging importance of antidistinguishability in quantum mechanics warrants a deeper understanding of the concept. In this paper, we prove that the optimal error probability of antidistinguishing finitely many quantum states is bounded above by the pairwise minimum of the optimal error probability of discriminating two states, and in the case of pure states, by the pairwise minimum of inner products of two pure states. We then show that the optimal error probability decreases exponentially in the asymptotic regime, and a lower bound on the asymptotic error exponent (achievable error rate) is given by the maximum of the pairwise Chernoff bound of the states. Furthermore, we provide an exact expression for the optimal asymptotic error exponent in classical antidistinguishability in terms of the Hellinger transform of mutually absolutely continuous probability measures. Joint work with Michael Nussbaum and Mark M. Wilde.

Speaker: Theshani Nuradha
Advisor(s): Ziv Goldfeld and Mark M. Wilde
Title: Quantum Pufferfish Privacy: A Flexible Privacy Framework for Quantum Systems
Abstract: We propose a new and versatile privacy framework for quantum systems, termed *quantum pufferfish privacy* (QPP). Inspired by classical PP, our formulation addresses limitations of quantum differential privacy (QDP) by offering flexibility in specifying private information, feasible measurements, and domain knowledge. We show that QPP can be equivalently formulated in terms of the Datta–Leditzky (DL) information spectrum divergence, thus providing the first operational interpretation thereof. We reformulate the DL divergence as a semi-definite program and derive new properties of it, which are then used to prove convexity, composability, and post-processing of QPP mechanisms. Parameters that guarantee QPP of the depolarization mechanism are also derived. We analyze the privacy-utility tradeoff of general QPP mechanisms and, again, study the depolarization mechanism as an explicit instance. The QPP framework is then applied to privacy auditing for identifying privacy violations via a hypothesis testing pipeline that leverages quantum algorithms. Connections to quantum fairness and other quantum divergences are also explored and several variants of QPP are examined. Joint work with Ziv Goldfeld and Mark M. Wilde and available at https://arxiv.org/abs/2306.13054.

Speaker: Yang Qiu
Advisor(s): Aaron B. Wagner
Title: Unifying Distortion and Realism Constraints
Abstract: In rate-distortion theory, mean squared error (MSE) is commonly used as the distortion measure. However, MSE suffers from certain limitations, including that optimal compressors under MSE distortion lead to blurry reconstructions at low bit-rates. Recently, "realism" constraints, which require that the distribution of the reconstruction is close to the distribution of the source, have become popular as a way of mitigating the limitations of MSE. In this work, we propose a new one-parameter family of distortion measures that unifies MSE distortion and realism constraints. Our distortion recovers MSE and realism constraints under extremal choices of the parameter. We also provide some theoretical results on optimal low-rate compression and experimental results with a trained compressor using the new distortion measure.

Speaker: Gabriel Rioux
Advisor(s): Ziv Goldfeld
Title: Entropic Gromov-Wasserstein Distances: Statistical and Computational Advances
Abstract: The Gromov-Wasserstein (GW) distance quantifies discrepancy between heterogeneous datasets and endows us with a means by which to align them, but suffers from computational hardness. The entropic Gromov-Wasserstein (EGW) distance serves as a computationally efficient proxy for the GW distance. Recently, it was shown that the quadratic GW and EGW distances admit variational forms that tie them to the well-understood optimal transport (OT) and entropic OT (EOT) problems. By leveraging this connection, we derive two notions of stability for the EGW problem. The first stability notion enables us to establish convexity and smoothness of the objective in this variational problem. This results in the first efficient algorithms for solving the EGW problem that are subject to formal guarantees. The second stability notion is used to derive a comprehensive limit distribution theory for the empirical EGW distance and, under additional conditions, asymptotic normality, bootstrap consistency, and semiparametric efficiency thereof. Joint work with Ziv Goldfeld and Kengo Kato.