ASA Section on Statistics in Imaging and FSU Workshop on

Recent Advances in Statistical Analysis of Imaging Data

December 4-5, 2020

Online Workshop (Zoom with Pre-Registration, No Fee)

Organizers




Ian Dryden, Xiaoming Huo, Ranjan Maitra, Anuj Srivastava


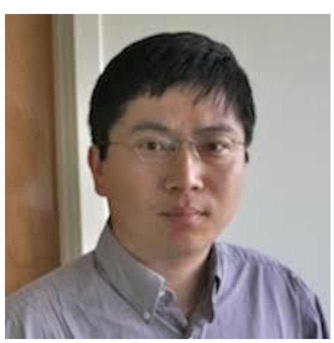

Marina Vanucci, Yingnian Wu, Tingting Zhang, Hongtu Zhu

Abstracts of Talks

Prof. Scott Acton, NSF and University of Virginia

Title: A presentation of opportunities for statisticians and imaging people at NSF.

Prof. Rama Chellappa, Johns Hopkins University

Title: Role of mathematical statistics in deep learning-based computer vision

Abstract: Over nearly forty years, I have applied results from mathematical statistical to develop representations and estimation methods for images using Markov random fields, develop Cramer-Rao bounds for structure from motion algorithms using discrete features and optical flow, apply Fisher-Rao metric for rate-invariant action recognition, etc. Since the reemergence of deep learning methods for solving many problems in computer vision, we have built end-to-end systems for face, object and action detection and recognition and are seeing remarkable performance in some of these tasks. Sadly, I have not been successful in giving any statistical analysis for any of these algorithms. It appears that mathematical statistics is not able to cope up with hierarchical and/or non-linear models. In this talk, I will briefly survey my experiences in using mathematical statistic in my computer vision work and discuss many problems that statisticians can work on to become major players in the era of big data and deep learning. These include simple problems such as regression for pose and landmark extraction using deep networks, domain adaptation and generalization, selecting the best subsets of training data in mini-batch learning, linearized approximations of deep networks and so on.

Prof. Moo Chung, University of Wisconsin

Title: Topological Learning for Brain Networks

Abstract: We present a novel topological learning framework that can integrate networks of different sizes and topology through persistent homology. This is possible through the introduction of a topological loss function that enables such challenging tasks. The use of the proposed loss function bypasses the intrinsic computational bottleneck associated with matching networks. The method is effectively applied to a twin brain imaging study in determining if the functional brain network is genetically heritable. The biggest challenge is in overlaying the functional brain networks obtained from the resting-state functional magnetic resonance imaging (fMRI) onto the structural brain network obtained through diffusion tensor imaging (DTI). While the functional network exhibits more cycles, the structural network is tree-like. Simply overlaying or regressing functional network on top of structural network will destroy the topology of both networks. We found the functional brain networks of identical twins are very similar demonstrating the strong genetic contribution on our though patterns. This is a joint work with PhD student Tananun Songdechakraiwut.

Prof. David Dunson, Duke University

Title: Inferring Low-Dimensional Structure in Brain Networks from MRI Data

Abstract: We are interested in studying how human and model organism brain connectomes vary, both randomly and systematically as a function of covariates. The structural brain connectome consists of the network of inter-connected white matter fiber bundles acting as highways for neural activity and communication. The locations of these fiber bundles can be indirectly inferred from MR Imaging Data. Using state-of-the-art pre-processing of the raw data, we obtain (hopefully) accurate representations of the connectome as an adjacency matrix (for each individual under study) representing the strength of connection between each pair of regions of interest (ROIs) in the brain. One then is faced with the problem of inferences on how covariates and traits of the individuals relate to such replicated graph/network data. This talk describes recent tools ranging from novel multi-scale graph PCA factorizations to graph-structured variational auto-encoders. We show that such methods can infer surprisingly strong relationships between brain connectomes and human traits.

Prof. Aasa Feragon, Denmark Technological University

Title: Graph-valued Models for Dimensionality Reduction and Regression

Abstract: Networks are everywhere! In anatomy and biology, they appear as transportation systems for air, water, nutrients, or signals, and are found both on the large scale of arteries and airways, and on the small scale of neurons in the brain. The structure, geometry and state of the networks affect their function, and therefore also the health of nearby tissue. Conversely, the state of surrounding tissue also affects the networks, making them both first and second order reporters of health, disease and dysfunction. As a consequence, networks are studied extensively in both biology and medicine -- and as a proxy for these, in imaging.

In this talk we discuss a well known space of graphs, where networks are modelled as equivalence classes of adjacency matrices modulo the action of the node permutation group. We derive geometric properties of this space and discuss the implications of those geometric properties for potential statistics being developed within the graph space. We move on to define statistical models for dimensionality reduction and graph-valued regression, which combine the intrinsic properties of the space with computational advantages due to the Euclidean nature of the space of adjacency matrices.

Prof. Al Hero, University of Michigan

Title: Sparse Covariance Models for High-Dimensional Image Data Analysis

Abstract: Many imaging applications produce multiway data of exceedingly high dimension. Modeling such multi-way data is important for applications in computer vision, hyperspectral imaging and hyperspectral video where imaging sensors produce data indexed over spatial, frequency, and temporal dimensions. We will introduce a generative sparse model for such imaging problems that is based on extensions of the matrix normal image representation that expresses the spatial-frequential-temporal covariance matrix as a Kronecker product of sparse lower dimensional covariances. This model represents the

Cholesky factors of the precision matrix as a sparse cartesian product graph. A sparsity penalized inference procedure is proposed for inferring the sparse Cholesky factors that we call the Sylvester Graphical Lasso (SyGlasso). We will illustrate the SyGlasso for multiway data arising in EEG spatio-temporal analysis and solar sunspot prediction.

Prof. Sebastian Kurtek, Ohio State University

Title: Variograms for spatial functional data with phase variation

Abstract: Spatial, amplitude and phase variations in spatial functional data are confounded. Conclusions from the popular functional trace-variogram, which quantifies spatial variation, can be misleading when analyzing misaligned functional data with phase variation. To remedy this, we describe a framework that extends amplitude-phase separation methods in functional data to the spatial setting, with a view towards performing clustering and spatial prediction. We propose a decomposition of the trace-variogram into amplitude and phase components and quantify how spatial correlations between functional observations manifest in their respective amplitude and phase components. This enables us to generate separate amplitude and phase clustering methods for spatial functional data, and develop a novel spatial functional interpolant at unobserved locations based on combining separate amplitude and phase predictions. Through simulations and real data analyses, we found that the proposed methods result in more accurate predictions and more interpretable clustering results. This is joint work with Xioahan Guo (Department of Statistics, The Ohio State University) and Karthik Bharath (School of Mathematical Sciences, University of Nottingham).

Prof. Ann Lee, Carnegie Melon University

Title: Using GOES Imagery to Better Understand the Evolution of Tropical Cyclone Convective Structure and Intensity Change

Abstract: Tropical cyclones (TCs) rank among the costliest natural disasters in the United States, and accurate forecasts of track and intensity are critical for emergency response. Intensity guidance has improved steadily but slowly, as processes which drive intensity change are not fully understood. Because most TCs develop far from land-based observing networks, geostationary satellite (GOES) imagery is critical to monitor these storms. However, these complex data can be challenging to analyze and interpret in real time, and off-the-shelf machine learning algorithms have limited applicability on this front due to their "black box" structure. In this talk, I will describe some of our recent and ongoing efforts on building a rich and interpretable functional representation of convective structure patterns which could help scientists and forecasters better understand the relationship between the evolution of TC convective structure and TC short-term intensity change. (This work is joint with Trey McNeely, Kimberley Wood and Dorit Hammerling)

Prof. Xavier Pennec, Université Côte d'Azur and INRIA

Title: Geometric statistics for computational anatomy

Abstract: At the interface of geometry, statistics, image analysis and medicine, computational anatomy aims at analysing and modelling the biological variability of the organs shapes and their dynamics at the population level from databases of images. The goal is to model the mean anatomy, its normal variation, its motion / evolution and to discover morphological differences between normal and pathological groups. However, shapes are usually described by equivalence classes of sets of points, curves, surfaces or images under the action of a transformation group, or directly by the diffeomorphic deformation of a template in diffeomorphometry. This implies that they live in non-linear spaces, while statistics where essentially developed in a Euclidean framework. For instance, adding or subtracting curves or surfaces does not really make sense. Thus, there is a need for redefining a consistent statistical framework for objects living in manifolds and Lie groups, a field which is now called geometric statistics. The objective of this talk is to give an overview of the Riemannian computational tools and of simple statistics in these spaces. The talk is motivated and illustrated by applications in medical image analysis, such as the regression of simple and efficient models of the atrophy of the brain in Alzheimer's disease and the groupwise analysis of the motion of the heart in sequences of images using the parallel transport of surface and image deformations.

Prof. Arun Ross, Michigan State University

Title: Altered Biometric Data

Abstract: Biometrics refers to the use of physical and behavioral traits such as fingerprints, face, iris, voice and gait to recognize an individual. The biometric data (e.g., a face image) acquired from an individual may be modified for several reasons. While some modifications are intended to improve the performance of a biometric

system (e.g., face alignment and image enhancement), others may be intentionally adversarial (e.g., spoofing or obfuscating an identity). Furthermore, biometric data may be subjected to a sequence of alterations resulting in a set of near-duplicate data (e.g., applying a sequence of image filters to an input face image). In this talk, we will discuss methods for (a) detecting altered biometric data; (b) determining the relationship between near-duplicate biometric data and constructing a phylogeny tree denoting the sequence in which they were transformed; and (c) deducing the camera or sensor that produced the original biometric data. We will focus on face and iris images. Analyzing the statistics of altered biometric data can lead to the development of robust schemes for evaluating the integrity of biometric data and for detecting malicious images such as DeepFakes.

Prof. Sudeep Sarkar, University of South Florida

Title: Towards Open World Video Event Understanding - Flexible Representations, Commonsense Priors, and Self-Supervised Learning

Abstract: Events are central to the content of human experience. From the constant stream of the sensory onslaught, the brain segments, represents aspect related to events, and stores memory for future comparison, retrieval, and re-storage. Contents of events consist of objects/people (who), location (where), time (when), actions (what), activities (how), and intent (why). Many deep learning-based approaches extract this information from videos. However, most methods cannot adapt much beyond what they were trained and are incapable of recognizing new events beyond what they were explicitly programmed or trained. The main limitation of current event analysis approaches is the implicit closed world assumption. The ability to support open world inference is limited by three main aspects: the underlying representation, the source of semantics, and the ability to learn or adapt continuously.

Pure deep learning methods are evolving into approaches that are a synthesis of symbolic reasoning and neural approaches, i.e., neuro-symbolic approaches. One of the most successful symbolic strategies involves probabilistic reasoning over random variables using graphical structures such as trees, lattices, and directed graphs. The most popular incarnations of these approaches are the Bayesian networks and Markov random fields, for example. Grenander's canonical representations offer a probabilistic symbolic approach that is a more generalized version than these popular models.

In this talk, I will focus on flexible representations, amenable for open-world, and self-supervised learning that is not dependent on a large amount of training data. We will see how Grenander's pattern theory-based canonical representation offers an elegant, flexible, compositional mechanism. It naturally models semantic connections between what is observed directly in the image and prior knowledge in large-scale commonsense

knowledge bases, such as ConceptNet. The use of knowledge bases such as ConceptNet allows expanding the set of primary objects and actions to a very large (not infinite) set without the need for massive annotated training sets. And finally, if we have time, how predictive learning can continuously learn how to segment a video into elementary event segments, again without training annotations.

Prof. Bertrand Thirion, INRIA, France

Title: Statistical inference in high dimension & application to brain imaging

Abstract: In many scientific applications, increasingly large datasets are being acquired to describe more accurately biological or physical phenomena. While the dimensionality of the resulting measures has increased, the number of samples available is often limited, due to physical or financial limits. This results in impressive amounts of complex data observed in small batches of samples. The following question arises: what features in the data are really informative about some outcome of interest? This amounts to inferring the relationships between these variables and the outcome, conditionally to all other variables. Providing statistical guarantees on these associations is needed in many fields of data science, where competing models require rigorous statistical assessment. Yet reaching such guarantees is very hard. In this talk, I will describe some algorithmic efforts to perform inference on structured data, such as images, while giving precise statistical guarantees. Illustrations will be given on brain imaging.

Prof. Marina Vannucci, Rice University

Title: Bayesian approaches for inference on brain connectivity networks

Abstract: Statistical methods play a crucial role in understanding and analyzing fMRI data. Bayesian approaches, in particular, allow flexible modeling of spatial and temporal correlations in the data, as well as the integration of multi-modal data. In this talk I will review Bayesian hierarchical models for the estimation of brain connectivity. I will illustrate a variational inference approach to multi-subject vector autoregressive models for inference on effective brain connectivity based on resting-state data. I will describe a user-friendly Matlab software and will show applications to data from TBI patients. The scalability of the methods results in the ability to estimate subject- and group-level brain connectivity networks over whole-brain parcellations of the data. I will also address dynamic models, for the estimation of temporal dynamics of brain networks.

Prof. Jane-Ling Wang, University of California, Davis

Title: Alignment of fMRI time-series and functional Connectivity

Abstract: Due to technology advance, spatially indexed objects are commonly observed across different scientific disciplines. Such object data are typically high-dimensional and pose great challenges to scientists due to the curse of high-dimensionality. While sparsity is commonly adopted as an assumption in high-dimensional settings, its validity is difficult to verify. We propose a new approach for spatially indexed object data by mapping their spatial locations to a targeted one-dimensional interval so objects that are similar are placed near each other on the new target space. The proposed alignment provides a visualization tool to view these complex object data. Moreover, the aligned data often exhibit certain level of smoothness and can be handled by approaches designed for functional data. We demonstrate how to implement such an alignment for fMRI time series and propose a new concept of path length to study functional connectivity, in addition to a new community detection method. The proposed methods are illustrated by simulations and on a study of the Alzheimer's disease. (Joint work with Chun-Jui Chen)

Prof. Xiao Wang, Purdue University

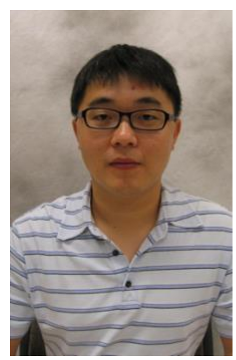

Title: Challenges in Generative Models and Latent Variable Models

Abstract: Latent variable models are an indispensable and powerful tool for uncovering the hidden structure of observed data for image analysis. The well-known generative models such as GANs ,VAEs and flow-based models belong to latent variable models. In this talk, I focus on three important and interconnected research topics related to latent variable models. First, we propose a framework called ALMOND for the inference of latent variable models that overcomes their current limitations. This new framework allows for a fully data-driven latent variable distribution via deep neural networks, and the proposed stochastic gradient method, combined with the Langevin algorithm, is efficient and suitable for complex models and big data. Second, we introduce a novel inferential Wasserstein GAN (iWGAN) model, which is a principled framework to fuse auto-encoders and WGANs. The iWGAN model jointly learns an encoder network and a generator network motivated by the iterative primal dual optimization process. Third, we propose an unbiased version of the contrastive divergence algorithm that completely removes its bias in stochastic gradient methods, based on recent advances on unbiased Markov chain Monte Carlo methods.

Prof. Rebecca Willett, University of Chicago

Title: Model Adaptation for Inverse Problems in Imaging

Abstract: Deep neural networks have been successfully applied to a wide variety of inverse problems arising in computational imaging. These networks are typically trained using a forward model that describes the measurement process to be inverted, which is used to generate the training data and is often incorporated directly into the network itself. However, these approaches lack robustness to misspecification of the forward model: if at test time the forward model varies (even slightly) from the one the network was trained on, the network performance can degrade substantially. Given a network trained to solve an initial inverse problem with a known forward model, we propose two novel retraining procedures that adapt the network to reconstruct measurements from a perturbed forward model, even without full knowledge of the perturbation. Our approaches do not require access to more labeled data (i.e., ground truth images), but only a small set of calibration measurements. We show these simple retraining procedures empirically achieve robustness to changes in the forward model in a variety of settings, including deblurring, super-resolution, and undersampled image reconstruction in magnetic resonance imaging. This is joint work with Davis Gilton and Greg Ongie.

Prof. Yingnian Wu, UCLA

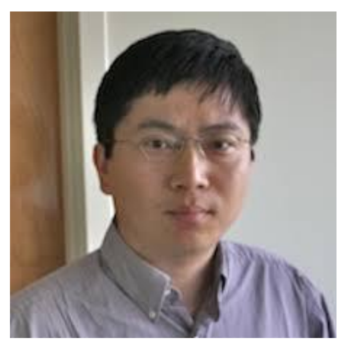

Title: Learning Latent Space Energy-Based Prior Model for Image, Text and Molecule Generation

Abstract: The generator model assumes that the observed example is generated by a low-dimensional latent vector via a top-down network, and the latent vector follows a known prior distribution, such as uniform distribution or isotropic Gaussian distribution. While learning an expressive top-down network, we can also learn an expressive prior model instead of assuming a given prior distribution. This follows the philosophy of empirical Bayes where the prior model is learned from the observed data. We propose to learn an energy-based prior model for the latent vector, where the energy function is parametrized by a simple multi-layer perceptron. We show that the learned model exhibits strong performances in

terms of image, text and molecule generation, anomaly detection, and semi-supervised learning. Joint work with Bo Pang, Tian Han, Erik Nijkamp, and Song-Chun Zhu.

Prof. Bin Yu, University of Berkeley

Title: Veridical Data Science: the practice of responsible data analysis and decision-making

Abstract: Building and expanding on principles of statistics, machine learning, and the sciences, we propose the predictability, computability, and stability (PCS) framework for veridical data science. Our framework is comprised of both a workflow and documentation and aims to provide responsible, reliable, reproducible, and transparent results across the entire data science life cycle, including problem formulation, data cleaning, modeling, and post-hoc interpretation.

We will illustrate the PCS framework in the modeling stage through the development of DeepTune images for characterization of neurons in the difficult V4 area of primary visual cortex.

Prof. Alan Yuille, Johns Hopkins University

Title: Vision as Bayesian Inference: The Importance of Generative Models

Abstract: Much recent progress in vision has been driven by discriminate models, like deep networks. But these models are not robust to significant changes in the images such as the presence of occluding objects or adversarial patch-attacks. We describe a class of generative models, called compositional networks, which are robust to these changes and can also localize the occluders and the attacking patches.

Prof. Tingting Zhang, University of Pittsburgh

Title: High-Dimensional Directional Brain Network Analysis for Focal Epileptic Seizures

Abstract: The brain is a high-dimensional directional network system consisting of many regions that exert influences onto each other. The directional influence from one region to another is referred to as a directional connection. Epilepsy is a directional network disorder, as epileptic activity spreads from a seizure onset zone (SOZ) to many other regions after seizure onset. However, studying epileptic directional brain networks has been limited to low-dimensional directional networks between the SOZ and its contiguous regions, due to the lack of efficient methods for identifying high-dimensional directional brain networks. To address this knowledge gap, we studied high-dimensional directional networks in epileptic brains by using a clustering-enabled multivariate autoregressive state-space model (MARSS) to analyze multi-channel intracranial EEG recordings of focal seizures. This new MARSS considers the SOZ, nearby regions, and many other non-SOZ regions as one integrated high-dimensional system while enabling clustering of regions. With the new MARSS, we identified the SOZ cluster, regions mostly affected by SOZ activity, and directional connections between all the recorded regions. We found that, after seizure onset, the directional connections of the SOZ and the number of regions in the SOZ cluster increased significantly. We revealed, for the first time, that many regions outside the SOZ cluster had no changes in directional connections, although these regions showed ictal activity. Lastly, we used these high-dimensional network results to localize the SOZ for patients with focal seizures and achieved 100% true positive rates and <3%$ false positive rates.

Prof. Hongtu Zhu, University of North Carolina

Title: Clinical and Methodological Advances in Imaging Genetics

Abstract: Recently the UK Biobank study has conducted brain MRI imaging scans of over 40,000 participants. In addition, publicly available imaging genetic datasets also emerge from several other independent studies. We collected massive individual-level MRI data from different data resources, harmonized image processing procedures, and conducted the largest genetic studies so far for various neuroimaging traits from different structural and functional modalities. In this talk, we showcase novel clinical findings from our large-scale analyses, such as the shared genetic influences among brain structures, functions, and a wide spectrum of clinical outcomes. We establish genetic mappings from hundreds

of anatomical brain regions onto their corresponding functional connectivities, and further onto complex mental disorders, which may transverse our understanding of the disease emergence and development. We also discuss methodological challenges we have faced when processing these biobank-scale datasets and highlight opportunities to utilize the learned knowledges in downstream analyses for disease predictions and pathway analyses. This presentation is based on a series of works of the UNC BIG-S2 lab. Our results can be easily browsed through the Brain Imaging Genetics Knowledge Portal (BIG-KP) (https://bigkp.org/).