Poster Presentations
May 19, 2017

Shaking up Seismology: Data Mining for Earthquake Detection

Presented by
Karianne Bergen, Sixth year ICME PhD

Seismic sensors collect massive quantities of data that contain a wealth of information about processes within the earth. Seismologists are increasingly adopting data mining and machine learning techniques to identify previously unknown earthquakes in large seismic data sets. Our new earthquake detection method, Fingerprint and Similarity Thresholding (FAST), enables waveform-similarity-based earthquake detection in long duration continuous seismic data by leveraging locality sensitive hashing, a data mining technique for efficiently identifying similar items in large data sets. Ongoing research into waveform feature representations and detection over a seismic network will be presented.

Karianne Bergen is a Ph.D. student in Computational & Mathematical Engineering at Stanford University. She holds a B.Sc. in Applied Mathematics from Brown University and a M.Sc. in Computational and Mathematical Engineering from Stanford University. Karianne previously worked at MIT-Lincoln Laboratory as an assistant technical staff member / data scientist . Karianne is graduating soon and is interested in speaking with you about job or postdoc opportunities!

"Is it true?" Deep Learning for Stance Detection in News

Presented by
Shruti Bhargava, First year ICME MS, Data Science Track, and
Neel Rakholia, First year ICME MS, Data Science Track

Stance detection is an important component of fake news detection. In this project we explore different neural net architectures for stance detection in news articles. In particular, we analyze the effectiveness of recurrent neural nets for this problem. We discover that a modified attentive reader model is well suited for the task. While our best deep learning model comfortably exceeds the baseline score set by Fake News Challenge, a simple feedforward network marginally outperforms it. As far as we are aware, our LSTM-based RNN model is the state of the art end-to-end deep learning model for this dataset.

I am currently pursuing my Masters degree in the Data Science track at the Institute for Computational and Mathematical Engineering, Stanford University. My undergraduate degree was in Computer Science. I wish to build innovative utilities which can assist people.

Neel Rakholia is a first year Master's student in the Data Science Track. His research interests are in the area of Machine Learning and Mathematical Modeling. At Stanford, he has worked on research projects in Deep Learning and Social Networks.

FAST LOW-RANK KERNEL MATRIX FACTORIZATION THROUGH SKELETONIZED INTERPOLATION

Presented by
Leopold Cambier, Second year ICME PhD

Integral equations are commonly encountered when solving complex physical problems. Their discretization leads to a kernel matrix that is block or hierarchically low-rank. This paper proposes a new way to build a low-rank factorization of those low-rank blocks at a nearly optimal cost of O (nr^2) for a n Ã— n matrix of rank r. This is done by combining interpolation of the kernel function and skeletonization of the interpolation matrix to further recompress the interpolation. We show the asymptotic accuracy of the algorithm and demonstrate on numerical examples that it performs very well, allowing to obtain rank nearly equal to the optimal rank at a fraction of the cost of the naive algorithm.

I graduated from Universite catholique de Louvain (Belgium) in 2015 with a Master's degree in Mathematical Engineering. I specialized in Optimization and Numerical Linear Algebra. Currently, I am interested in Fast Linear Solvers, either dense or sparse. My current project consists of accelerating BEM equations solve through low-rank approximations of the kernel function.

Fast Unsupervised Object Localization

Presented by
Anjan Dwaraknath, Third year ICME PhD

As opposed to image classification, object localization is a computer vision problem whose solutions have not measured up to human level performance, even with the use of deep learning. Common approaches to address this problem in the deep learning framework, are moving sliding windows across the image, training neural architectures to predict bounding boxes, and using classic image processing techniques such as SIFT and Region Proposals. None of these methods utilize the innards of the classification neural network to localize images. In fact, one could argue that these methods are created with the assumption that neural networks are just abstruse black boxes that give us an output for every input. We posit that a neural network for classification has enough usable spatial information to localize objects. We introduce a powerful and novel technique that harnesses only a pretrained CNN to localize an object in an image. To be clear, we do not train on any bounding box data. We obtained very promising results. In our validation we found the center of the localization methodâ€™s bounding box was within the ground truth bounding box in more than 90% of the cases. The core idea of our technique can be extended to other computer vision problems such as multi-class object localization and image segmentation. In addition, since our technique works so well with only a classification neural network, there is good reason to believe that we can improve the methods that currently train on bounding box data.

I am a third year PhD Student in ICME. My general research interests include Linear algebra, Optimization and numerical solution of PDEs. In particular I am interested in machine learning solutions that improve performance of simulations of nonlinear physical systems such as reservoir models that can also aid in interpretation of results and uncertainty quantification.

The Merits of Keeping It Smooth

Presented by
Ron Estrin, Third year ICME PhD

We develop an exact penalty function approach with a smooth penalty function for solving constrained nonlinear programs based on work originally proposed by Roger Fletcher in the 1970's. For a sufficiently large penalty parameter, minimizers of the original problem are minimizers of the penalty function. Minimization of problems with difficult nonlinear constraints are replaced by a minimization of a simpler problem with a smooth objective function with either no constraints or simple (linear/bound) constraints.

Ron Estrin is a third year Ph.D. student in ICME, working with Michael Saunders, Michael Friedlander and Dominique Orban, on projects in numerical optimization and linear algebra. Prior to studying at Stanford, he received his B.S. in Mathematics and Computer Science from the University of British Columbia. Outside of academics, he enjoys practicing martial arts and playing tennis.

Towards a Multi-fidelity Hemodynamic Model Pipeline for the Analysis of Cardiovascular Flow Under Uncertainty

Presented by
Casey M. Fleeter, Second year ICME PhD

Deterministic hemodynamic models are successfully employed in various branches of cardiovascular disease research. However, their widespread adoption is hindered by their inability to account for uncertainty stemming from multiple sources. This motivates the transition to a stochastic framework where modeling parameters are defined in probability with distributions either assumed or assimilated from available patient-specific data. A possible approach to performing stochastic analysis on these systems, while maintaining reasonable computational cost, is to leverage multiple varying-fidelity models of the same cardiovascular flow. In this context, we propose an automatic pipeline to convert three-dimensional vascular models into their one-dimensional approximations, and discuss its validation and integration in the SimVascular software suite. Finally, various approaches for non-intrusive uncertainty propagation were first validated on simple one-dimensional parametric models and differential Windkessel circuits, and successively used to characterize output uncertainty in one-dimensional hemodynamics.

Casey is a 2nd year Ph.D. student at ICME currently working in Alison Marsden's lab. Her research interests include uncertainty quantification and computational fluid dynamics. Previously, she received a B.A. in physics from Harvard University.

New MCMC scheme: State Substitution and Its Applications

Presented by
Pengfei Gao, Third year ICME PhD

We describe a new approach to MCMC that uses sequential state substitutions for its Metropolis-Hastings-type transitions. The basic idea is to approximate a target distribution by the empirical distribution of N representative atoms, chosen sequentially by an MCMC scheme so that the distribution converges weakly to the target distribution as the number K of iterations approaches infinity. We also develop the asymptotic theory of MCMC with sequential state substitutions and provide a consistent estimator of its standard error. We conduct some numerical experiments to compare the algorithm with other MCMC scheme and also show how it could be applied to adaptive particle filter.

Pengfei Gao is a third-year PhD candidate in ICME advised by Professor Tze Leung Lai. His research interests are AI applied in quantitative trading. Pengfei received his Bachelorâ€™s degree in Mathematics and Physics (2014) from Tsinghua University.

Neural-Based Method for ATLAS Cluster Splitting

Presented by
Guillaume Genthial, First year ICME MS, Data Science Track

The first step in classifying an event in a particleâ€™s detector such that the LHC is to reconstruct the individual pions that are produced by the collision and decay of other particles. Today's detectors use a clustering algorithm that group cell energy depositions on the different layers of the calorimeter to reconstruct truth particles. This algorithm is imprecise and a gg gglot of reconstructed clusters contain more than one particle. This project studies neural-network based techniques based on computer vision to identify merged clusters.

The main contribution are:

- a proof of concept that there is enough information in calorimeter cells to improve the clustering, without using tracks.
- a baseline, feature-based model that achieves competitive performance on the classification task
- a multi-view model that uses the layers as slices and performs 2d convolutions on top of reconstructed images.
- a model inspired by the PointNet architecture, fully end-to-end and without any hand-crafted features that directly consumes the cells and learns a representation (embedding) of each cell.

I am currently enrolled in my first year as a master's student in Computational and Mathematical Engineering in the Data Science track, where I develop strong computational and programming skills. I also have a masters' degree in Applied Mathematics from the Ecole polytechnique, the France's top university for science and engineering, where I also studied computer science, quantum and statistical physics. I have solid research experience in Natural Language Processing and Deep Learning.

Topology Optimization of Damped Resonators

Presented by
Dustin Gerrard, Fifth year ICME MS, ME PhD

MEMS resonators are used in many applications including inertial navigation and timing. To achieve a strong signal a high quality factor (Q) is needed. Q is often limited by thermoelastic damping. We have developed a method of eliminating this damping using mathematical techniques to create unique topologies.

Dustin is in the Tom Kenny lab working on resonators. He plans to go into industry after graduating this year.

A Computational Framework For Supersonic Parachute Inflation Dynamics

Presented by
Daniel Zhengyu Huang, Third year ICME PhD

The discovery of water on Mars and the successful launch and landing of used rockets promise that the future of space exploration will be different from anything we have seen before! However, the recent failures at flight test time of all supersonic, parachute-based decelerators for Mars landing during their inflation are of great concern. To understand how to avoid such failures, predictive high fidelity simulations of the dynamics of a parachute during its inflation are needed. Yet, most if not all relevant computational efforts have focused on developing Computational Fluid Dynamics (CFD) and Fluid-Structure Interaction (FSI) parachute models for the much easier simulations in the postinflation regime. This project develops a multimaterial framework to simulate the parachute inflation process.

My name is Daniel Zhengyu Huang, a third Ph.D. students in ICME.
I am passionate about applying our knowledge to solve practical engineering problems.

On the effects of Partying on subsequent day sports performance

Presented by
Andreas Santucci, Third year ICME MS, Data Science Track
Eric Lax, Third year ICME MS, Data Science Track

It's not unreasonable to think that game performance can be affected (in part) by what takes place off the court. We are interested to see if teams exhibit a decline in performance the day following a game in a party city. We exploit data on bookmaker spreads, the expected score differential between two teams after conditioning on observable performance. We expect a team to meet the spread half the time, since this is how bookmakers minimize risk on their end. We construct a model which attempts to estimate the causal effect of visiting a "party-city" the day before a game on the probability of meeting the spread.

Andreas Santucci and Eric Lax are both part of the Data Science track in ICME and Statistics. They have been working together for several years on coursework, theoretical research such as distributing minimum cut algorithms, and also causal AI. Their recent work is with Guido Imbens, a professor at Stanford GSB and school of economics.

Mini-batch Markov Chain Monte Carlo

Joint work with Rachel Wang, Tung-yu Wu and Wing Hung Wong
Presented by
Dangna Li, Fourth year ICME PhD

In this paper we propose a general framework of performing Markov Chain Monte Carlo(MCMC) with only a mini-batch of data. We show by estimating the Metropolis-Hasting ratio with only a mini-batch of data, one is essentially sampling from the true posterior raised to a certain temperature. We demonstrate through experiments that such a mini-batch MCMC algorithm can be used to efficiently move between modes of a non-convex posterior distribution.

Dangna Li is a forth year PhD in ICME. Her current research focuses on large scale Markov Chain Monte Carlo algorithms.

Temporal Oscillations in the Porous Medium Equation: Why Harmonic Averaging Itself Is Not to Blame

Presented by
Danielle Maddix, Fourth year ICME PhD

Harmonic averaging of the coefficient, $k(p) = p^m$, $m \geq 1$, has been blamed for the nonphysical locking and lagging of numerical pressure solutions of the Porous Medium Equation, $p_t - \nabla \cdot (k(p)\nabla p) = 0$, that occur for small p. This degenerate parabolic equation has applications to plasma heat transfer, gas and groundwater flow. The numerical issues also manifest themselves in spurious temporal oscillations, even with none in space. Arithmetic averaging has been suggested as an alternative. We show that harmonic averaging is not solely to blame and that an improved choice of temporal and spatial discretizations can avoid these issues. The harmonic problems can be traced for standard numerical schemes to a local anti-diffusive term in the modified equation and so can be compensated for. A similar approach works for superslow diffusion, where $k(p) = \exp(-1/p)$.

A more extreme case, arising in foam models, is where k(p) is a step function, e.g. $k(p) = 1$ if $p \geq 0.5$ and 0 otherwise. For this model, both harmonic and arithmetic averaging result in artificial temporal oscillations. To resolve this, we propose a new type of averaging, incorporating the shock position that can be estimated effectively.

Danielle Maddix is a fourth year PhD candidate in the Institute of Computational and Mathematical Engineering (ICME). She has been funded by the NSF graduate research fellowship. Danielle received her bachelor's degree in Applied Mathematics with highest honors from the University of California, Berkeley in 2012. She then continued onto her graduate education at ICME at Stanford and earned her Masters in Computational and Mathematical Engineering in 2015. She has had internships in computational mathematics at the Lawrence Berkeley National Laboratory and NVIDIA. Her current research is on developing new, accurate, stable and conservative numerical methods for mathematical modeling of ocean dynamics.

DeepLofting: Building Cardiovascular Models with Convolutional Neural Networks

Presented by
Gabriel Maher, third year ICME PhD

In this work, we tackle the challenging task of constructing 3D cardiovascular models from medical image data. We combine deep learning with a vessel path-line model construction pipeline to form DeepLofting, a new and efficient method for building cardiovascular models. We start by developing two novel neural network architectures, HED-FC and I2I-FC, by combining neural network layers to form a spatial context processor which allows fully convolutional networks to utilize image spatial context to produce accurate localized segmentations. DeepLofting uses convolutional neural networks to compute vessel boundaries along anatomical path-lines and combines these boundaries to form a final cardiovascular model. Given vessel path-lines, DeepLofting is fully-automatic requiring substantially less user-intervention compared to traditional model building workflows. We evaluate our architectures and DeepLofting pipeline on a publicly available dataset of 84 computed tomography (CT) and magnetic resonance (MR) volumes. DeepLofting, combined with HED-FC or I2I-FC, significantly outperforms other neural network architectures and popular cardiovascular model building methods, producing accurate 3D cardiovascular models compared to user constructed 3D models

Gabriel Maher is a PhD student at the Institute for Computational and Mathematical Engineering at Stanford University. For his research Gabriel is applying Deep Learning to cardiovascular medical image analysis with Dr. Alison Marsden at the Cardiovascular Biomechanics Computation Lab.

Dark Fiber Seismology: Repurposing our Telecomm Infrastructure as a Dense, Continuously Recording Seismic Array

Presented by
Eileen Martin, Fifth year ICME PhD

Distributed Acoustic Sensing (DAS) repurposes a fiber optic probed by a laser as a series of strain rate sensors. DAS is being increasingly adopted in the energy industry for permanent monitoring of microseismic events and repeatable time-lapse seismic imaging, and our work suggests it is a promising technology for seismology and near-surface studies. Our previous work tested the utility of fiber buried in trenches as a fit-for-purpose installation, but we can drastically reduce installation costs by using fiber in existing telecommunications conduits.

Since September 2016, we have been continuously recording vibrations along 2.5 km of fiber running in existing telecomm conduits under the Stanford campus with the goals of (i) dense recording of earthquakes and quarry blasts, (ii) ambient noise studies and (iii) active seismic surveys. I will show results from the array and a few practical tools to deal with the challenges of working with densely recorded noisy DAS data in a loud urban recording environment, particularly with regards to near-surface imaging using random noise. Because of the type of sensor we use, we extract coherent signals that are a mix of different types of surface waves, and I show the first steps towards imaging between fiber optic lines.

Eileen Martin is a 5th year ICME PhD student advised by Prof. Biondo Biondi (geophysics). She holds BS degrees in math and physics from the University of Texas at Austin. She has spent the past few summers at Lawrence Livermore National Lab, Shell, and Schlumberger. Her studies have been supported by the DOE CSGF, Schlumberger Innovation Fellowship, and affiliates of the Stanford Exploration Project.

Robust and Efficient Multi-Way Spectral Clustering

Presented by
Victor Minden, Fifth Year ICME PhD

We present a new algorithm for spectral clustering based on a column-pivoted QR factorization that may be directly used for cluster assignment or to provide an initial guess for k-means. Our algorithm is simple to implement, direct, and requires no initial guess. Furthermore, it scales linearly in the number of nodes of the graph and a randomized variant provides signiï¬_cant computational gains. Provided the subspace spanned by the eigenvectors used for clustering contains a basis that resembles the set of indicator vectors on the clusters, we prove that both our deterministic and randomized algorithms recover a basis close to the indicators in Frobenius norm. Finally, we experimentally demonstrate that the performance of our algorithm tracks recent information theoretic bounds for exact recovery in the stochastic block model.

I am a fifth-year PhD student at the Institute for Computational and Mathematical Engineering (ICME) at Stanford University, where I work with Lexing Ying. My research is focused on fast algorithms for scientific computing, in particular fast linear algebra based on data-sparse representations of rank-structured matrices that arise from physical problems in 2D or 3D. More broadly, my interests include numerical linear algebra, numerical optimization, signal processing, spectral graph theory, and high-performance computing.

Persistent Homology of Time-Dependent Filtrations on Temporal Networks

Collaborating with Ron Estrin and Austin Benson
Presented by
Bradley Nelson, Third Year ICME PhD

Many data sets have the structure of a graph with temporal information on the edges between vertices, such as communications on a social network. We investigate the use of persistent homology to compute structural features of these networks as they evolve in time. Persistent homology is a tool from algebraic topology that in this case reveals connected components, cycles, and higher-order structures as information moves through the network. We introduce both local and global filtrations (sequences of graphs) which are used in the persistence algorithm that provide interpretability to the computed features. We demonstrate the use of these filtrations on several real-world communication data sets.

Brad Nelson is a third year ICME PhD student interested in applied and computational topology. He has a BA in mathematics from Dartmouth, and is currently supported by a NDSEG Fellowship. He previously worked at Epic, and has done internships at Ayasdi and Lawrence Livermore National Lab while at Stanford. At ICME, he is involved in computational consulting and helps organize the student seminar. Brad enjoys running and hiking in the Santa Cruz mountains, cooking, and rock climbing.

Bayesian Unidimensional Scaling For Latent Ordering And Uncertainty Estimation

Presented by
Lan Huong Nguyen, Fourth year ICME PhD

Detecting patterns in high-dimensional multivariate datasets is non-trivial. Clustering and dimensionality reduction techniques often help in discerning inherent structures. In biological datasets such as microbial community composition or gene expression data, observations can be generated from a continuous process, often unknown. Estimating data points' `natural ordering' and their corresponding uncertainties can help researchers draw insights about the mechanisms involved.

We introduce a Bayesian Unidimensional Scaling (BUDS) technique which extracts dominant sources of variation in high dimensional datasets and produces their visual data summaries, facilitating the exploration of a hidden continuum. The method maps multivariate data points to latent one-dimensional coordinates along their underlying trajectory, and provides estimated uncertainty bounds. By statistically modeling dissimilarities and applying a DiSTATIS method to their posterior samples, we are able to incorporate visualizations of uncertainties in the estimated data trajectory across different regions using confidence contours for individual data points. We also illustrate the estimated overall data density across different areas by including density clouds. One-dimensional coordinates recovered by BUDS help researchers discover sample attributes or covariates that are factors driving the main variability in a dataset. We demonstrated usefulness and accuracy of BUDS on a set of published microbiome 16S and single cell RNA-seq data.

Our method effectively recovers and visualizes natural orderings present in datasets. Automatic visualization tools for data exploration and analysis are available at: https://github.com/nlhuong/visTrajectory

I am a fourth year ICME PhD student advised by prof. Susan Holmes. My research is on statistical methods for analyzing biological data, with a focus on genomic sequencing data including microbiome and single cell.

Tensor Networks

Presented by
Cindy Orozco, Second year ICME PhD

Tensor Network (TN) diagrams are widely used tools in quantum physics to express density functions of electrons in certain arrangements. Although density functions can be high dimensional (one dimension per electron), they can be understood as contractions of easy-to-understand low dimensional tensors. Once we have a set of densities in TN form, we can extend the set imposing an algebra on it, being able to compute sums, products, reciprocals and even logarithms of these high dimensional objects. Some results for the spin-glass model in 1D periodical arrangement are shown.

I am a second year PhD student at ICME focused in general numerical methods, working with professor Lexing Ying. I did my undergrad in Civil Engineering and Math at Universidad de los Andes, in BogotÃ¡ (Colombia) and I got a masters degree in Applied Mathematics and Computational Sciences at King Abdullah University of Science and Technology (KAUST), in Thuwal (Saudi Arabia). As a side, I am visual artist, interested in the expression of color and shape.

Water Supply Network Optimization

Presented by
Mike Phulsuksombati, second year ICME MS

The goal of this paper is to optimize the global minimum-energy flow of a water supply network at a fixed time. We want to compute a feasible flow that transmits the water from water sources to satisfy the demand and pressure requirement for customers at end nodes while the energy within the network is minimized. The problem features nonconvex objective function and constraints as well as discrete decision variables on water network components such as to turn on/off the pumps and valves. We propose a convex relaxation for the problem and a heuristic method based on the solution of the convex problem to recover the solution of the discrete decision variables.

I am a second-year master student at the ICME. My interests are in linear/quadratic optimization and graph algorithm. I am working with Prof. Yinyu Ye on quadratic constraint quadratic programming (QCQP) problems. I received my BS in computer science and mathematics from Stanford University.

Active Learning for Accurate Estimation of Linear Models

Presented by
Carlos Riquelme, Fifth year ICME PhD

We explore the sequential decision making problem where the goal is to estimate uniformly well a number of linear models, given a shared budget of random contexts independently sampled from a known distribution. The decision maker must query one of the linear models for each incoming context, and receives an observation corrupted by noise levels that are unknown, and depend on the model instance. We present Trace-UCB, an adaptive allocation algorithm that learns the noise levels while balancing contexts accordingly across the different linear functions, and derive guarantees for simple regret in both expectation and high-probability. Finally, we extend the algorithm and its guarantees to high dimensional settings, where the number of linear models times the dimension of the contextual space is higher than the total budget of samples. Simulations with real data suggest that Trace-UCB is remarkably robust, outperforming a number of baselines even when its assumptions are violated.

Carlos is a fifth-year PhD student at Stanford, and holds a MSc in Mathematics and Foundations of Computer Science from the University of Oxford. Previously, he obtained his bachelor degree in Computer Science and Mathematics from Universidad Autonoma de Madrid. In 2016, he spent a few months at Facebook Research, working in several projects at the Experimental Design and Causal Inference team. He has also interned at Twitter, Quora, and Adobe Reseach. Carlos is mainly interested in the design of intelligent systems that can make sequential decisions and incorporate feedback, through efficient mechanisms to explore and measure uncertainty. His research lies at the intersection of Reinforcement Learning, Active Learning, and Bayesian Optimization. After finishing his PhD, he will join Google Brain.

Human Interaction with Recommendation Systems

Joint work with Carlos Riquelme
Presented by
Sven Schmit, Fifth year ICME PhD

Many recommendation algorithms rely on user data to generate recommendations. However, these recommendations also affect the data obtained from future users. This work aims to understand the practical effects of this dynamic interaction. We propose a simple model where users with heterogeneous preferences arrive over time. Based on this model, we prove that naive estimators, i.e. those which ignore this feedback loop, perform badly due to selection bias. Furthermore, recent work has focused on efficient exploration with myopic agents. We show that although agents behave myopically, agents' heterogeneous preferences ensure that recommendation systems 'learn' about all alternatives without explicitly incentivizing this exploration. This leads to practical insights that are relevant to designers of a wide range of systems designed to help users make better decisions.

Sven Schmit is a PhD student at ICME working with Prof. Ramesh Johari. He studied Econometrics and Operations Research at the University of Groningen and Mathematics at the University of Cambridge. Previously, he has worked at HP Labs and Stitch Fix.

Feedback Networks

The full list of the poster/paper is Amir R. Zamir, Te-Lin Wu, Lin Sun, William B. Shen, Bertram Shi, Jitendra Malik, Silvio Savarese.
Presented by
William B. Shen, 3rd year BS, Computer Science

Currently, the most successful learning models in computer vision are based on learning successive representations followed by a decision layer. This is usually actualized through feedforward multilayer neural networks, e.g. ConvNets, where each layer forms one of such successive representations. However, an alternative that can achieve the same goal is a feedback based approach in which the representation is formed in an iterative manner based on a feedback received from previous iteration's output.

We establish that a feedback based approach has several core advantages over feedforward: it enables making early predictions at the query time, its output naturally conforms to a hierarchical structure in the label space (e.g. a taxonomy), and it provides a new basis for Curriculum Learning. We observe that feedback develops a considerably different representation compared to feedforward counterparts, in line with the aforementioned advantages. We present a general feedback based learning architecture, instantiated using existing RNNs, with the endpoint results on par or better than current feedforward networks and the addition of the above advantages.

William Shen is a BS student of Computer Science at Stanford University, where he is working at the Vision and Learning Lab. His research interests lies within the intersection of computer vision and cognitive science. His current research focuses on building machines that can percept the 3D world in a more comprehensive way.

Design of an Optimal Frequency Reward Program in the Face of Competition

Presented by
Nolan Skochdopole, fourth year ICME PhD

We optimize the design of a frequency reward program against traditional pricing in a competitive duopoly, where customers measure their utilities in rational economic terms.
We assume two kinds of customers: myopic and strategic. Every customer has a prior loyalty bias toward the reward program merchant, a parameter drawn from a known distribution, indicating an additional probability of choosing the reward program merchant over the traditional pricing merchant.
Under this model, we characterize the customer behavior: the loyalty bias increases the switching costs of strategic customers until a tipping point, after which they strictly prefer and adopt the reward program merchant. Subsequently, we optimize the reward parameters to maximize the revenue objective of the reward program merchant. We show that under mild assumptions, the optimal parameters for the reward program design to maximize the revenue objective correspond exactly to minimizing the tipping point of customers and are independent of the customer population parameters. Moreover, we characterize the conditions for the reward program to be better when the loyalty bias distribution is uniform - a minimum fraction of population needs to be strategic, and the loyalty bias needs to be in an optimal range.
If the bias is high, the reward program creates loss in revenues, as customers effectively gain rewards for "free'', whereas a low value of bias leads to loss in market share to the competing merchant.
In short, if a merchant can estimate the customer population parameters, our framework and results provide theoretical guarantees on the pros and cons of running a reward program against traditional pricing.

Nolan is a fourth year PhD student in ICME, advised by Amin Saberi. He is currently working to develop models for loyalty reward programs. Other interests include graph theory and approximation algorithms. He studied applied math and chemistry at University of Chicago for undergrad.

Sparse Canonical Correlation Analysis

With Victor Minden, Bradley Nelson
Presented by
Xiaotong Suo, Fifth year ICME PhD

Canonical correlation analysis was proposed by Hotelling and it measures linear relationship between two multidimensional variables. In high dimensional setting, the classical canonical correlation analysis breaks down. We propose a sparse canonical correlation analysis framework by adding ` 1 constraints on the canonical vectors and show how to solve it efficiently using linearized alternating direction method of multipliers (ADMM) and TFOCS (Becker et al.). We illustrate this idea on simulated data.

I am a fifth year Ph.D candidate in ICME. I am interested in high dimensional statistics, and numerical optimization.

Optimal Investment in Online Loans

Presented by
Nina Troha, Credit Analyst @ Colchis Capital Management LP Visiting Student Researcher @ CS Department

Recently institutional investors have started to allocate capital to consumer, small business, and other types of loans originated by online lenders such as Lending Club, Prosper, and OnDeck. Standard investment policies are rule-based; they select loans with certain features without addressing risk/return trade-offs. In this work we harness granular online loan performance data to construct risk/return-optimal loan portfolios using machine learning methods. The distribution of portfolio returns is estimated using a two-step approach: loan-level outcomes are first classified using a random forest algorithm and then loan-level cash flows are modeled conditional on loan outcomes using a beta regression. Fixing a target expected portfolio return, we then select the portfolio with minimum expected shortfall. This approach leads to a tractable linear program with semi-continuous constraints that capture the trading constraints investors face in practice. We implement our approach for Lending Club loans, and show that our optimal portfolios significantly outperform a number of widely-used rule-based portfolios on a risk-adjusted basis. Our approach has also implications for the securitization of online lending loans.

Nina Troha received a BSc degree in Financial Mathematics from University of Ljubljana, Slovenia, and a MSc degree in Quantitative Finance from ETH Zurich, Univeristy of Zurich, Switzerland. She graduated Summa Cum Laude from both, BSc and MSc degrees, and was named best student in class during the BSc. She was a Researcher at Stanford University for a year, where she developed a machine learning model for risk and return prediction and portfolio optimization in online lending. Nina is currently on leave from Ph.D. program at

Stanford. She is working at a leading online lending investment management company, Colchis Capital Management, L.P.

Integrated All-Frequency Physics-Based Sound Synthesis for Physics-Based Animation

Collaborating with Dr. Tim Langlois at Adobe Research
Presented by
Jui-Hsien Wang, Second year ICME PhD

In this project, we aim to build a general finite-difference time-domain wave solver for high-quality sound synthesis for physics-based animation. We will introduce novel algorithms used in our system to support animated objects with various sound source models, including linear modal model, acceleration sound (clicks) model, and water bubble sound model. We will discuss how to optimize the integration of these models onto our system, and challenges in mapping the current system onto state-of-the-art parallel computer architecture, such as GPUs."

I am a second year ICME PhD who is interested broadly in computer animation, computer graphics, and Physics-based sound synthesis.

Feedback Networks

Presented by
Te-Lin Wu, second year EE MS

Currently, the most successful learning models in computer vision are based on learning successive representations followed by a decision layer. This is usually actualized through feedforward multilayer neural networks, e.g. ConvNets, where each layer forms one of such successive representations. However, an alternative that can achieve the same goal is a feedback based approach in which the representation is formed in an iterative manner based on a feedback received from previous iteration's output.

We establish that a feedback based approach has several fundamental advantages over feedforward: it enables making early predictions at the query time, its output naturally conforms to a hierarchical structure in the label space (e.g. a taxonomy), and it provides a new basis for Curriculum Learning. We observe that feedback networks develop a considerably different representation compared to feedforward counterparts, in line with the aforementioned advantages. We put forth a general feedback based learning architecture with the endpoint results on par or better than existing feedforward networks with the addition of the above advantages. We also investigate several mechanisms in feedback architectures (e.g. skip connections in time) and design choices (e.g. feedback length). We hope this study offers new perspectives in quest for more natural and practical learning models.

Te-Lin Wu is a second year Master student in Electrical Engineering at Stanford. Te-Lin currently is a affiliated research assistant at Stanford Computational Vision and Geometry Lab (CVGL) led by Professor Silvio Savarese where he conducted research in representation learning in computer vision using deep learning algorithms. His research interests span from computer vision, general artificial intelligence to robotics.

Local Higher-Order Graph Clustering

Presented by
Hao Yin, second year ICME PhD

Local graph clustering methods are attractive because they enable targeted clustering around a given seed node and are faster than traditional global graph clustering methods. However, current local graph partitioning methods are not designed to account for the higher-order structures crucial to the network, nor can they effectively handle directed networks. Here we introduce a new class of local graph clustering methods that address these issues by incorporating higher-order network information captured by small subgraphs, a.k.a. network motifs. We first show how to adapt the approximate personalized PageRank algorithm to find clusters containing a seed node with minimal motif conductance, a generalization of the conductance metric for network motifs. Then we generalize existing theory to maintain the properties of fast running time (independent of the size of the graph) and cluster quality (in terms of motif conductance). For community detection tasks on both synthetic and real-world networks, our new framework outperforms the current edge-based personalized PageRank methodology.

I am a second-year PhD student in ICME. My main research interest lies in data mining, especially social network analysis. I also work on machine learning methodology and optimization.

Generic 3D Representation

Amir R Zamir, Tilman Wekel, Pulkit Agrawal, Colin Wei, Jitendra Malik, Silvio Savarese.
Presented by
Amir R. Zamir, Postdoc, Computer Science

Though a large body of computer vision research has investigated developing generic semantic representations, efforts towards developing a similar representation for 3D has been limited. In this paper, we learn a generic 3D representation through solving a set of foundational proxy 3D tasks: object-centric camera pose estimation and wide baseline feature matching. Our method is based upon the premise that by providing supervision over a set of carefully selected foundational tasks, generalization to novel tasks and abstraction capabilities can be achieved. We empirically show that the internal representation of a multi-task ConvNet trained to solve the above core problems generalizes to novel 3D tasks (e.g., scene layout estimation, object pose estimation, surface normal estimation) without the need for fine-tuning and shows traits of abstraction abilities (e.g., cross modality pose estimation).

In the context of the core supervised tasks, we demonstrate our representation achieves state-of-the-art wide baseline feature matching results without requiring apriori rectification (unlike SIFT and the majority of learnt features). We also show 6DOF camera pose estimation given a pair local image patches. The accuracy of both supervised tasks come comparable to humans. Finally, we contribute a large-scale dataset composed of object-centric street view scenes along with point correspondences and camera pose information, and conclude with a discussion on the learned representation and open research questions.

Additional material: http://3drepresentation.stanford.edu/

Amir R. Zamir is a postdoctoral researcher at Stanford and UC Berkeley. His research interests lie broadly in computer vision and machine learning with a focus on self-supervised/unsupervised learning, open-world vision, and video understanding. He is particularly passionate about going beyond task-specific models and developing the theoretical and experimental grounds of a human-like comprehensive and generalizable perception. Amir is the recipient of CVPR Best Student Paper Award (2016), UCF Graduate Research Forum Award (2014), National Geospatial-Intelligence Agency (NGA) NARP-SW Award (2013), and UCF Research Excellence Award (2013).

Maximum Improvement in Level Set Estimation using Gaussian processes

This poster is collaborated by Andrea Zanette and Junzi Zhang.
Presented by
Andrea Zanette, second year ICME PhD, and
Junzi Zhang, second year ICME PhD

This work focuses on the problem of determining as large a region as possible where a function exceeds a given threshold with high probability. We assume we only have access to a noise-corrupted version of the function and that function evaluations are costly. To select the next query point we propose to maximize the expected increase in the area identified as above the threshold as predicted by a Gaussian process. Our numerical results indicate that our approach can significantly outperform existing heuristics in the literature, especially when there is a limited budget of function evaluations.

Andrea Zanette is a second year master student in ICME. He is broadly interested in theoretical artificial intelligence with an emphasis on reinforcement learning and decision making under uncertainty. He holds a BS in mechanical engineering.

Junzi Zhang is a second year Ph.D. student in Institute for Computational & Mathematical Engineering at Stanford. Before coming to Stanford, he obtained my B.S. degree in School of

Mathematical Sciences, Peking University (China). His current research interest mainly lies in the joint of optimization and machine learning, especially convex optimization, Bayesian optimization and reinforcement learning. Apart from this, he is also interested in application of these techniques to (renewable) energy problems and data mining problems.