Internship and research proposals

Center for Imaging Sciences,
Johns Hopkins University

Main contact information:

- Nicolas Charon, assistant professor
  Department of Applied Mathematics and Statistics
  charon@cis.jhu.edu
- Rene Vidal, associate professor
  Department of Biomedical Engineering, Department of Electrical Engineering
  rvidal@cis.jhu.edu

General information

The University. The Johns Hopkins University was founded in 1876 and counts as one of the leading research university in the world. Its main facilities (Homewood campus), including the Whiting School of Engineering, are located in Baltimore City, in the state of Maryland, 60km North of Washington DC. It has recently ranked in the top 10 universities in the US. Thanks to the close collaboration with Johns Hopkins Hospital, it is also considered as one of the best research centers in the field of medical sciences.

The Center for Imaging Sciences (CIS). CIS is an interdepartmental research center within the university. It brings together faculties, graduate students and research fellows from Biomedical Engineering, Applied Mathematics and Statistics, Computer Science and Electrical Engineering departments. Its essential focuses are image analysis and computer vision, statistical learning as well as computational biology. More information at http://cis.jhu.edu/

Masters degree internships. The following pages contain several research projects for internships with faculty members at CIS. The period and length of these may be flexible but keep in mind that it requires the obtention of a J1 visa from a US embassy, for which the full process can take up to 4 months. All proposed internships come with a minimum stipend pay of $2000/month.

Doing a PhD at Johns Hopkins University. We encourage motivated and capable students to pursue successful internships by entering one of the Ph.D programs proposed at the university. Each department offers one or possibly several Ph.D programs that include financial support for up to 6 years, each with different admission rules and completion requirements. Any student willing to start a PhD must apply to one program of his choice by submitting his candidacy, which shall be then examined with the others by a department admission committee. Note that deadlines for fall 2016 applications (at least for part of the application documents) vary from department but are generally early. Below are the different information links for several of them:

- Applied Mathematics and Statistics. Deadline: January 15th for fall 2016 admission, September 15th for Spring 2016 admission. All infos on the program at http://engineering.jhu.edu/ams/graduate-studies/overview/ams-phd-program/
- Electrical Engineering. Deadline: December 15th for fall 2016 admission. http://engineering.jhu.edu/ece/phd-program/
- Computer Science. Deadline: December 15th for fall 2016 admission. http://www.cs.jhu.edu/graduate-studies/phd-program/
- Biomedical Engineering. Deadline: December 1st for fall 2016 admission. http://www.bme.jhu.edu/graduate/phd/overview

Note also that all applications to PhD programs in the US ask for official TOEFL and Graduate Record Examination (GRE) transcripts. Both these can be obtained by taking a short test in one of the offical test center in France.

Registration of multi-modal brain images

Subject: A very central problem in image analysis, particularly with medical images, is the one of *registration*, which consists in computing an optimal transformation aligning an image onto another one. These two images can for instance correspond to acquisitions of different subjects or of a single individual but at multiple time points. In addition, the impressive development of imaging devices in the biomedical field now also provides us with images of various *modalities* that generally display important local variations in contrast in addition to the geometrical misalignment.

A current important research topic is thus to develop mathematical models and algorithms for large deformation image alignment procedures that are at the same time robust to such local contrast variations between images. In the classical formulation, registration of an image I on another image J writes as the optimization over a deformation ϕ of a functional:

$$d(\mathrm{Id}, \phi)^2 + \|I \circ \phi^{-1} - J\|_{L^2}^2 \tag{1}$$

which is a sum of a regularization term on the deformation (usually a metric on a given group of deformations), and a dissimilarity term between the deformed I and J here measured by the L^2 norm. A lot of work has been done to construct groups of 'large' deformations and associated metrics, in particular with the model of large diffeomorphic deformation metric mapping (LDDMM) proposed by Miller, Trouvé, and Younes.

However, the issue of the dissimilarity term that drives the deformation optimization in such frameworks is not as frequently examined. The simple L^2 metric between images used in equation (1) is an adapted measure of dissimilarity for images of identical modality and similar contrast but fails as soon as local or global contrast changes appear between source and target images. The main objective is to extend the current methodology to efficiently align such images.

State of the art techniques have so far mostly relied on *Mutual information* metrics that are however known to be quite unstable in large dimension registration problems. Alternative ways to explore could consist in constructing different contrast-invariant metrics derived from geometric representations of the image's set of level lines/surfaces and solve the resulting variational problem. Another possibility is to combine deformation group action on images by additional contrast renormalization functions based on histogram matching techniques, and propose a joint estimation formulation and algorithm. The developed methods will be applied on datasets available at CIS and acquired with the very recent technology named *CLARITY*, that provides ultra high-resolution 3D imaging of mice brains with different conditions.

Related references:

- P. Viola, W.M Wells (1997). Alignment by Maximization of Mutual Information, *International Journal of Computer vision* 24(2), 137-154.
- M.F Beg, M.I Miller, A. Trouvé, L. Younes (2005). Computing Large Deformation Metric Mappings via Geodesic Flows of Diffeomorphisms, *International Journal of Computer vision* 61(2), 139-157.
- R. Tomer, L. Ye, B. Hsueh, K. Deisseroth (2014). Advanced CLARITY for rapid and high-resolution imaging of intact tissues, *Nature Protocols* 9, 1682–1697.

Candidate requirements: background in image processing topics and shape analysis, programming skills in MATLAB, Python and/or C++.

Contact info: Nicolas Charon, assistant professor
e-mail: charon@cis.jhu.edu

Combining anatomy and function in medical data analysis

Background and motivation: Biomedical imaging has been an incredibly active field of research in the past decades, which has been triggered in great part by the development of various and more and more accurate acquisition techniques such as computed tomography, MRI, ultrasound or more recently CLARITY, allowing to visualize internal structures in a non-invasive way. The prodigious amount of data now available has lead to important challenges in the analyzis of this data, with notable implications for the understanding of biological processes or the detection/monitoring of pathologies.

One of such issue is related to the study of the anatomy of the imaged structures, in particular the modelling of variability of this anatomy wether among populations of different subjects or along time sequences. This is particularly challenging in the case of the brain which display both very intricate and variable geometry. Such problems have thus found fruitful connections with *shape analysis* and statistical learning, which resulted in the growing domain named *Computational Anatomy*.

Many methods in computational anatomy have been designed to treat large sets of 2D or 3D images (e.g obtained from MRI) or datasets of previously extracted geometrical boundaries like cortical or subcortical triangulated surfaces. Yet, the evolution of data acquisition methods has once again brought new categories of data structures, many of them adding to the anatomy an extra *functional* information. This the case in functional MRI(fMRI) where, after preprocessing, one can can reconstitute activation maps over the cortical surface. In the study of neurodegenerative pathologies, it is often very interesting to study the variations in both the shape of tissues together with potential thickness loss at all vertex locations. In both examples, objects are essentially signals (thickness, brain activity...) but defined over varying shapes.

Internship goals: Attempts to extend the approach of computational anatomy to this class of data is only very recent (cf Charon, Charlier, Trouvé 2015) with the introduction of so called *fshapes*. The subject of this internship is to complete this framework by addressing important remaining issues towards a complete algorithmic pipeline processing shape and function in a joint way (and possibly including that to the current FshapeTk software developed in MATLAB and CUDA). This shall consist more specifically in designing and implementing methods for denoising, reconstruction or inpainting of signal functions defined on potentially large and irregular surface meshes, or treat the situation of datasets having partial observations and uncorresponding boundaries. The developed algorithms could be evaluated on different application datasets such as a cohort of Planum Temporale surfaces used in the study of schizophrenia, or on a dataset of retinal membranes with thickness measurements in relation to the prediction of glaucoma (collaboration with Simon Fraser University in Vancouver).

Related references:

- M.I Miller, A. Qiu, (2009). The emerging discipline of Computational Functional Anatomy, *NeuroImage* 45, 16-39.
- B. Charlier, N. Charon, A. Trouvé (2015). The fshape framework for the variability analysis of functional shapes, to appear in *J. Found. of Comput. Maths.*
- FshapesTk software: https://github.com/fshapes/fshapesTk

Candidate requirements: strong interest and background in computer vision topics, programming skills in MATLAB, C++ and/or CUDA. Some knowledge in discrete geometry is also useful.

Contact info: Nicolas Charon, assistant professor
e-mail: charon@cis.jhu.edu

Hardwiring Biological Mechanism into Statistical Learning

Background. Gigantic amounts of data about abnormal perturbations in biological networks are being collected by sequencing and microarray experiments, enabling the inference of systems-level disease signatures. Cancer is perhaps the prototypical systems disease, and the focus of extensive study in quantitative biology, including developing algorithms to predict disease phenotypes, progression and treatment response for individuals. However, translating these programs into personalized clinical care remains elusive, and in particular mature applications of statistical learning to inference at the patient level remain limited.

Realizing this agenda requires recognizing and overcoming the two major limitations of current methodology: (i) the "Black box" decision rules which lack the biological meaning and algorithmic simplicity necessary for transitioning into the clinic; and (ii) the absence of biological mechanism, which is necessary for stable model selection, especially given the unfavorable ratio of samples to variables in the current and any foreseeable learning scenario.

Internship subject. Embedding mechanism in the learning framework is a "win-win" strategy, but requires collaboration with biologists, which is available at Johns Hopkins. Specifically, the student will explore methods for embedding phenotype-dependent mechanisms specific to cancer pathogenesis and progression directly into the mathematical form of the decision rules. These mechanisms concern circuitry involving microRNAs (miR), transcription factors (TF) and their known targets, for example bi-stable feedback loops involving two miRNAs and two mRNAs. Such network motifs can be identified in signaling pathways and biochemical reactions intimately linked to cancer phenotypes.

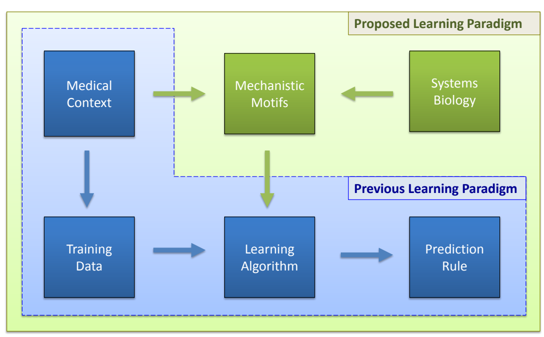

Contact info: Pr. Donald Geman (geman@cis.jhu.edu) and Pr. Laurent Younes (younes@cis.jhu.edu).

Analysis of Diffusion Magnetic Resonance Images

Motivations. Diffusion Magnetic Resonance Imaging (DMRI) is a medical imaging technique that is used to estimate the anatomical network of neuronal fibers in the brain, *in vivo*, by measuring and exploiting the constrained diffusion properties of water molecules in the presence of bundles of neurons. Water molecules will diffuse more readily along fibrous bundles (think of fiber optics cables), then in directions against them. Therefore by measuring the relative rates of water diffusion along different spatial directions, we can estimate the orientations of fibers in the brain. In particular, one important type of DMRI technique that we will analyze is high angular resolution diffusion imaging (HARDI) which measures water diffusion with an increased level of angular resolution in order to better estimate the probability of fiber orientation, known as the Orientation Distribution Function (ODF). HARDI is an advancement over the clinically popular Diffusion Tensor Imaging (DTI) which requires less angular measurements because of a Gaussian assumption which restricts the number of fiber orientations that can be estimated in each voxel. More accurate estimates of ODFs at the voxel level using HARDI lead to more accurate reconstructions of fiber networks. For instance, the extraction of neuronal fibers from HARDI can help understand brain anatomical and functional connectivity in the corpus callosum, cingulum, thalamic radiations, optical nerves, etc. DMRI has been vital in the understanding of brain development and neurological diseases such as multiple sclerosis, amyotrophic lateral sclerosis, stroke, Alzheimer's disease, schizophrenia, autism, and reading disability.

Project Goals. To make DMRI beneficial in both diagnosis and clinical applications, it is of fundamental importance to develop computational and mathematical algorithms for analyzing this complex DMRI data. In this research area, we aim to develop methods for processing and analyzing HARDI data with an ultimate goal of applying these computational tools for robust disease classification and characterization. Possible project areas include:

1. **Sparse HARDI Reconstruction**: To develop advanced algorithms for the sparse representation and reconstruction of HARDI signals with the goals of speeding up HARDI acquisition and compact data compression.

2. **ODF Estimation**: To develop advanced algorithms for computing accurate fields of Orientation Distribution Functions (ODFs) from HARDI data.

3. **Fiber Segmentation**: To develop advanced algorithms for automatically segmenting HARDI volumes into multiple regions corresponding to different structures in the brain.

4. **HARDI Registration**: To develop advanced algorithms for the registration of HARDI brain volumes to preserve fiber orientation information.

5. **HARDI Feature Extraction**: To develop methods for extracting features from high-dimensional HARDI data that can be exploited for ODF clustering, fiber segmentation, HARDI registration and disease classification.

6. **Disease Classification**: To develop advanced classification techniques using novel HARDI feature representations to robustly classify and characterize neurological disease.

Internship Goals. In our lab, the intern will work with a PhD student to complete a project within an area(s) mentioned above. The intern will read a number of research papers on DMRI and the state-of-the-art, and will learn an understanding of the problem, its applications, and the techniques involved to tackle it. Their are two aspects of the research that may be of interest to the applicant. One is a more theoretical aspect that involves developing mathematical theories to improve existing frameworks. The second is a more computational aspect that involves more image processing, analysis, and algorithm implementation in MATLAB or C++. An applicant with some interest and experience in both areas is most favorable, but it is possible for an applicant to be interested in working on only one of the aspects as well.

At the end of the internship period the student will present their work to other graduate students and professors and will potentially be able to publish their research in medical imaging conferences and journals. As becoming part of the Vision Lab, the intern will experience first-hand a rigorous and rewarding research environment with a highly collaborative and supportive social element.

Experience in MATLAB or C++ and familiarity with image analysis or processing is a plus. Mathematical maturity is also favorable.

More Information. Check out the lab website (http://www.vision.jhu.edu), or contact Prof. René Vidal at rvidal@cis.jhu.edu. Please put the text [Internship Summer 2016] in the beginning of the subject line.

Automated Analysis of Cardiac Action Potentials from Human Embryonic Stem Cells

Motivations. The use of human embryonic stem cell cardiomyocytes (hESC-CMs) in tissue transplantation and repair has led to major advances in cardiac regenerative medicine. However, to avoid potential arrhythmias, it is critical that hESC-CMs used in replacement therapy be electrophysiologically compatible with the adult atrial, ventricular, and nodal phenotypes. One approach to ensuring compatibility is by investigating the electrophysiological signature, or action potential (AP), of a cardiomyocyte.

Project Goals. The goal of this project is to tackle this problem by using machine learning techniques in order to provide objective measures for classifying action potentials derived by human embryonic stem cells by their shape. Our previous work has shown that by using a shape preserving distance framework called *metamorphosis* and computational models of adult APs, one can, with high accuracy, identify the phenotype (atrial, ventricular, or nodal) of the embryonic cardiomyocyte. Further, the framework provides an interpolation between the embryonic and mature APs that may be representative of the maturation process. Our current goal is to optimize the current framework for use in larger populations, as well as use the framework to investigate the efficacy of current biochemical methods for the purification of specific phenotype CMs.

Internship Goals. In this project, the intern will work with PhD students to develop novel mathematical models for representing embryonic and mature cardiac action potentials (APs) and methods for classifying APs from cardiac time-series data. Further, the maturation interpolants will be used to update current, state-of-the-art computational cardiomyocyte models by introducing cell maturation. The intern will implement code to demonstrate their performance on synthetic data as well as a large AP dataset.

The intern will read a number of research papers on cardiac signal acquisition, electrophysiology, cardiac cell models, and machine learning and will learn an understanding of the problem, its applications, and the techniques involved to tackle it. Moreover, the intern will implement novel algorithms in MATLAB and C++ and become familiar with analyzing cardiac time-series data as well as evaluating the developed methods on acquired datasets. The intern will present their work to other graduate students and professors and will potentially be able to publish their research in biomedical engineering conferences and journals. As part of the group, the intern will experience first-hand a rigorous and rewarding multi-disciplinary research environment.

Experience in C++ and MATLAB coding and familiarity with cardiac electrophysiology, signal processing, machine learning, or differential equations (partial and ordinary) is a plus.

More Information. Check out the lab website (http://www.vision.jhu.edu), or contact Prof. René Vidal at rvidal@cis.jhu.edu. Please put the text [Internship Summer 2016] in the beginning of the subject line.

Modeling and Teaching the Language of Surgery

Motivations. To learn the art of surgery, medical students practice in a number of phantoms, but rarely in an actual patient. Recent technological advances have enabled the use of surgical robots for performing certain surgical procedures. However, surgical training still relies on the subjective visual evaluation by an expert. In fact, surgical expertise is more often than not judged by the number of practice hours rather than the actual expertise level. But how can we define and quantify surgical expertise more precisely?

Project Goals. The goal of this project is to develop a more objective way to evaluate the skill of a medical student or surgeon. For that purpose, we are using motion and video data acquired by the Da Vinci robot to build models of surgical tasks that can be used for training and evaluation of medical students. In the same way that speech is divided into sentences, words and phonemes, one can divide a surgical task into subtasks, such as suturing, knot-tying, etc., and each subtask into basic surgical motions, such as grabbing a needle, inserting a needle, etc. Given motion and video data from multiple tasks and performed by different surgeons with different levels of expertise, our goal is to discover the basic surgical motions (the words) and the typical transitions among such basic surgical motions (the grammar). To discover this *language of surgery*, we are using advanced techniques from dynamical systems, machine learning and computer vision. Such techniques automatically segment the data into different surgical motions and determine the most likely sequence of surgical motions being executed. To determine the skill level, we are looking at the way in which basic surgical motions are executed (measured in terms of kinematic and dynamic features) and the way in which transitions between basic surgical motions occur (e.g., novices tend to transition a lot). We thus build models that describe the different motions and the different transitions for different skill levels and use these models to evaluate the skill of new surgeons.

Internship Goals. In this internship, you will develop a set of mathematical tools for modeling skilled human activities (e.g. surgical tasks). The mathematical models investigated will include: language-based models, computer vision classification tools (e.g. bag of features), hybrid dynamical systems, and sparse-representation based approaches. These tools will provide new insights into the relationship between skill, style, and content in human motion.

In addition, you will develop software using these tools for teaching, training and assistance of humans performing these activities. The data of a person who is trying to learn the activity will be compared with the model derived by observing experts, and both physical guidance and information display tools will be developed to provide feedback to the learner based on the expert model.

Experience in MATLAB and C++ coding and familiarity with computer vision, statistical language models or dynamical systems is a plus.

More Information. Check out the lab website (http://www.vision.jhu.edu), or contact Prof. René Vidal at rvidal@cis.jhu.edu. Please put the text [Internship Summer 2016] in the beginning of the subject line.

Human Activity Recognition

Motivations. The human visual system is exquisitely sensitive to an enormous range of human movements. We can differentiate between simple motions (left leg up vs. right hand down), actions (walking vs. running) and activities (playing vs. dancing). We can also identify friends by their walking styles, infer mood and intent from hand or arm gestures, or evaluate the grace and athleticism of a ballerina. Recently, significant progress has been made in automatically recognizing human activities in videos. Such advances have been made possible by the discovery of powerful image descriptors and the development of advanced classification techniques. However, the performance of the "features + classifiers" approach seems to be reaching a limit, which is still well below human performance.

Project Goals. The goal of this project is to develop algorithms for recognizing human movements, actions and activities in unstructured and dynamically changing environments. For instance, recognizing face, hand and arm gestures, human gaits, and daily activities (shaking hands, drinking, eating, etc.). Classical 2D representation will be merged with 3D data (motion capture, Kinect, accelerometers) in order to represent a human performing an action as a collection of 2D/3D pose/shape and 2D/3D dynamical models.

Internship Goals. As part of the project, the intern will work alongside PhD students and develop novel algorithms for 3-D human motion representation for activity recognition. The intern will implement code for these algorithms as well as test them on several databases.

The intern will read research papers on activity recognition, 3D shape modeling, motion capture-based recognition methods, and will learn new techniques to solve the above problem. Moreover, the intern will implement novel algorithms in MATLAB and C++ and become familiar with several computer vision and machine learning concepts. The intern will present their work to other graduate students and professors and will potentially be able to publish their research in computer vision conferences and journals. As part of the group, the intern will experience first-hand a rigorous and rewarding research environment.

Experience in C++ and MATLAB coding and familiarity with classification techniques and dynamical systems is a plus.

More Information. Check out the lab website (http://www.vision.jhu.edu), or contact Prof. René Vidal at rvidal@cis.jhu.edu. Please put the text [Internship Summer 2016] in the beginning of the subject line.

Object Recognition

Motivations. When a person is shown an image, he/she is able to immediately identify a variety of things like : the various objects present in the image, their locations, their spatial extent, their categories and the underlying 3D scene of which it is an image. The human visual system uses a combination of prior knowledge about the world and the information present in the image to perform this complicated task. We want to replicate this on a computer. This is broadly called object recognition and it involves object detection (is there an object in this video?where is it located?), segmentation (which pixels contain the object?), categorization (what is the object's class?) and pose estimation (what is the 3D location of object in the scene?). We also want to perform all these tasks jointly rather than a pipieline approach as knowledge of one task helps us perform the others better.

Project Goals. The project aims to develop object represenations (models that capture prior knowledge about how the object looks like under varying viewing conditions) and techniques to perform the tasks of object detection, categorization, image segmentation and pose estimation in a fast and efficient manner. We are developing a graph-theoretic approach in which different levels of abstractions, such as pixels, superpixels, object parts, object categories, their 3d pose, relative configuration in the scene, etc., are represented as nodes in a graph. Contextual relationships among different nodes are captured by an energy defined on this graph. In this energy, bottom-up costs capture local appearance and motion information among neighboring nodes. Each of the tasks corresponds to terms in the energy function (the top-down costs), which is then solved in a joint manner. We are also developing algorithms based on branch and bound (pose estimation task) and graph cuts (image segmentation task) for minimizing the energy, and methods based on structured output learning (structural SVMs) to learn its parameters.

Internship Goals. As part of the project, the intern will help enhance our current framework for object recognition by improving the model to capture more sub-categories, develop models for more object categories and design algorithms to utilize these models for various vision tasks. The intern will be exposed to current research in the area of Object Recognition and Scene Understanding. He/she will read a lot of literature on a variety of topics like image representation, clustering, classification and graphical models. The intern will implement algorithms in Matlab/C++ and test them across various datasets. The intern will present their work to other graduate students and professors and will potentially be able to publish their research in computer vision conferences and journals. This project will help the intern gain a good understanding of challenges and research opportunities in the area of Object Recognition and Scene Understanding. Experience in C++ and MATLAB coding and familiarity with image processing, computer vision, or statistical inference is a plus.

More Information. Check out the lab website (http://www.vision.jhu.edu), or contact Prof. René Vidal at rvidal@cis.jhu.edu. Please put the text [Internship Summer 2016] in the beginning of the subject line.

Subspace Clustering

Motivations. Consider the task of separating different moving objects in a video (e.g., running cars and walking people in a video clip of a street). While human can easily solve this task, a computer would find itself totally clueless since all that it sees is a big chunk of ordered 0's and 1's. Fortunately for the computers, this problem has a specific property that allows an alternative approach which a computer would find more comfortable with. That is, for all the points of the same moving object, the vectors built from their trajectories lie in a common subspace. Thus this problem boils down to a math problem of separating different subspaces in the ambient space.

Project Goals. Given a set of data points that are drawn from multiple subspaces with unknown membership, we want to simultaneously cluster the data into appropriate subspaces and find low-dimensional subspaces fitting each group of points. This problem is known as *subspace clustering*, it has applications in, beside the motion segmentation mentioned above, image segmentation, face clustering, hybrid system identification, etc. The Vision Lab has worked extensively on this topic, and has developed methods of geometric approaches such as the Generalized Principle Component Analysis, and spectral clustering approaches such as the Sparse Subspace Clustering. The performance of the algorithms on a motion segmentation benchmark and a face recognition database has been studied. The goal of the project is thus to further improve the algorithms for subspace clustering, and study the performance the algorithms on tasks that have the multi-subspace structure.

Internship Goals. As part of the project, the intern will work alongside PhD students and develop novel algorithms for subspace clustering. The intern will implement code for these algorithms as well as test them on several databases.

The intern will learn necessary background knowledge in machine learning, computer vision, compressed sensing, and will read research papers on subspace clustering. Moreover, the intern will implement novel algorithms in MATLAB to different datasets. The intern will present their work to other graduate students and professors and will potentially be able to publish their research in computer vision conferences and journals. As part of the group, the intern will experience first-hand a rigorous and rewarding research environment.

Strong background in linear algebra, experience in MATLAB coding is a plus.

More Information. Check out the lab website (http://www.vision.jhu.edu), or contact Prof. René Vidal at rvidal@cis.jhu.edu. Please put the text [Internship Summer 2016] in the beginning of the subject line.