Signal Processing
Tools
Collaborators
At Work in the World
Theory
Algorithms
Mathematics
Optimization
Computer Vision
Random Matrix Theory
Compressive Sampling
Computational Imaging
Source Coding + Compression
Statistical Estimation + Learning
Machine Learning
Biologists
Statisticians
Pathologists
Oncologists
Radiologists
Geneticists
Psychologists
Mathematicians
Neuroscientists
Environmental Scientist
All Engineeering Disciplines
Computer Scientists
Musicians
Environmental Monitoring
Security
Networking
Medical Imaging & Diagnostics
Personalized Medicine
Sound Technology
Big Data
Construction Site Monitoring
Nuclear Nonproliferation
Image Indexing
Assistive Technology
Assisted Driving
Music, Image & Video Transmission

Signal Processing @ Michigan: Putting Theory to Work for a Better World

What's That Signal?

In the movie *Independence Day*, David Levinson's algorithm detects an unexpected signal buried in atmospheric background noise. When he decodes the meaning embedded in the signal, he realizes that the world is about to be attacked by aliens. Had the original signal not been detected, it definitely would *not* have been a day of independence.

Similarly, though perhaps with less dramatic consequences, signal processing researchers at Michigan are on the lookout for ways to detect and manipulate signals that will help make the world safer (assisted driving, nuclear material detection, surveillance); healthier (medical imaging and diagnosis, environmental monitoring); and more enjoyable (audio, image, and video processing)

Signal processing is the art of generating, transforming, and interpreting information, which we think of as being a signal. Some signals can be detected electronically, while others might be a pattern that arises from data such as a gene pool, social media activity, or economic data. Signal processing delves into nearly as many application areas as can be conceived.

Tools and Collaborations

Michigan faculty bring a variety of theoretical tools to the task of detecting and manipulating signals. These tools include machine learning, random matrix theory, computer vision, compressive sampling, computational imaging, source coding compression, optimization, and statistical estimation & learning.

With the explosion of information available in today's world thanks to pervasive sensing, access to the nanoworld, and increased computational and storage capability of computers, signal processing faculty are approached by potential collaborators from all areas of science and engineering.

"This is not a group that acts in isolation," said Al Hero, R. Jamison and Betty Williams Professor of Engineering, also director of the signal processing area and professor of biomedical engineering and statistics. "The high level of interdisciplinary and cross-disciplinary interactions defines our approach to signal processing."

Michigan signal processing faculty collaborate with radiologists, pathologists, oncologists, geneticists, biologists, statisticians, mathematicians, physicists, and materials scientists, just to name a few. The fact that top researchers in many fields can be found right here at Michigan greatly facilitates their work, yet they still reach out and actively collaborate with others in academia and industry across the country and around the world.

The Theory Behind the Work

Michigan signal processing faculty approach their research with an eye to fundamental theory so their work can be applied to an ever-expanding array of problems. If a project comes along that tests the limits of what can currently be done, they become intrigued. If it seems to require new theory – they get excited.

"While we may solve a specific problem related to economic forecasting or networking," stated Prof. Hero, "this same research can be extended to gaining environmental insights from tree-rings, or earthquake activity."

The theory is highly mathematical. Researchers are trying to extract the maximum possible information from the available data using algorithms and mathematical modeling. An algorithm can be thought of as a computer program that takes the data, in this case a signal, and outputs some kind of decision that will yield the desired information.

Some of the applications to which Michigan faculty are applying their theory are described below.

Flow Cytometry for Diagnosis of Blood-related Illness

Flow cytometry is the process of gathering and quantifying properties of individual blood cells. It is used by hemato-pathologists to diagnose blood-related disorders such as leukemia and lymphoma.

Part of the analysis involves isolating the cells of interest manually. The key challenge lies in the fact that each individual's blood has unique biological markers which must be identified and classified, a process that is currently done by hand. Prof. Clay Scott is applying machine learning techniques to automate this process, saving time and eliminating some forms of human error. He is working with a pathologist in the U-M Medical School who has provided him with blood samples so that he can test his new computational theory.

In related work, Prof. Hero was able to render clinical flow cytometry data more interpretable to pathologists through a method called Fisher Information Non-parametric Embedding (FINE). Clinical flow cytometry can reveal up to 12 different properties of a given cell, and about 100,000 individual cells are analyzed for each patient. The pathologist needs to look at all the resulting data and determine what it means. The FINE method, which is currently being used at U-M Hospital, applies a computer algorithm to the data and pulls out the most relevant information needed to determine prognosis of a disease.

More recently, Prof. Hero successfully applied FINE and a novel visualization method called Information Preserving Components Analysis (IPCA) to the problem of diagnosing myelodysplastic

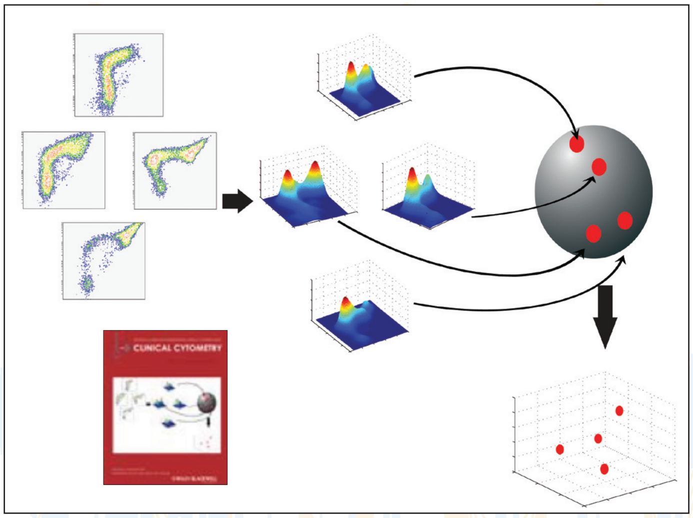

Prof. Hero's research was featured on the cover of Cytometry Part B: Clinical Cytometry, vol. 80B, Issue 5, Sept. 2011. The figure show a schematic overview of FINE.

syndromes (MDS). MDS are a group of diseases of the blood and bone marrow that can lead to anemia or even leukemia.

Flow cytometry data comes in the form of a distribution of the properties of tens of thousands of a patient's blood cells. Under a recent grant from NSF, Al Hero and Clay Scott are teaming up to develop new methods for estimation and classification of distributional data, which will be applied to flow cytometry. They will collaborate with Dr. Lloyd Stoolman at the U-M Medical School.

Predictive Medicine is Big Data

Access to information in today's world is seemingly unlimited. Data coming in from sensors and other information gathering devices now includes the nanoworld, and modern computers seem capable of storing it all. This access to "big data" enables cutting-edge research to be done in areas such as genomics and predictive medicine. It is also pushing researchers to allocate resources at each step so they can sift through the available data.

Prof. Hero is investigating individual gene expressions to determine an individual's relative health as well as their susceptibility to different behavior patterns, such as fatigue or stress. "The idea is you want to take gene sequence data from each individual," explained Hero, "and compare it to their gene expression at a later date to detect changes, or to predict illness." However, the amount of genetic information is so vast it has to be considered sequentially, and much of the information isn't even reliable. His goal is to expend as few resources as possible to get usable baselines for comparison.

Predicting who will get the flu also requires accessing big data sets. Collaborating with faculty from various fields and institutions, Prof. Hero is having success unraveling what in our complex genomic data accounts for why some get sick while others don't. He studied 22,000 genes in 267 blood samples, the largest sampling pool of its kind ever investigated. To find the meaningful data signal amidst the noise, he adapted a pattern recognition algorithm previously developed for satellite imaging of the environment. The algorithm was able to identify the unique genomic signatures associated with immune response and flu symptoms.

Better, Faster Images With Lower X-ray Doses

One of the key ways doctors detect abnormalities in the human body is to take an image through an MRI, CT scan, or other method. These images also track the progress of medical treatment. Prof. Jeff Fessler has spent much of his career improving the quality of images received through these imaging

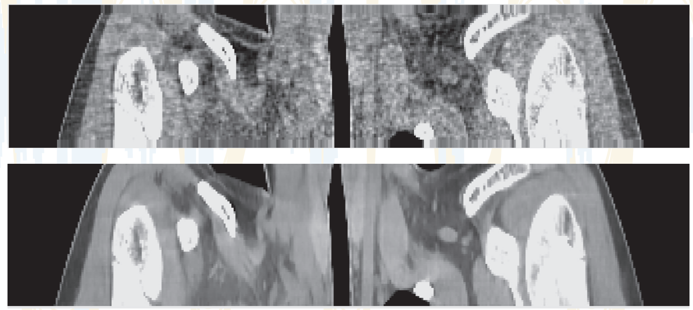

Traditional CT scan vs. improved image using 25% of the X-ray radiation.

techniques, which allows for a reduction in the amount of X-ray some of these machines put into the body. He accomplishes this by employing complex algorithms. He was involved in the early stages of technology known as Veo, manufactured by General Electric, that is currently in use at the University of Michigan Health System. Veo creates quality CT scans using much lower doses of radiation than was previously needed.

With radiation levels now at a minimum, his attention has turned to getting even better images, including 3D images, from the raw data that comes out of a CT or MRI scanner, and getting them faster. He is currently working with computer science colleagues at Michigan to significantly speed up the processing time.

Imaging With Motion

CT scanners are often used to generate 3D photos of coronary arteries, but the fact that the heart keeps beating disturbs the resulting image. CT and other types of medical scans don't operate like cameras which can employ fast shutter speeds and generate instant photos. Instead, signal processing engineers design specialized algorithms that turn the raw data generated by the medical devices into images that can be interpreted by humans, and CT scanners can't get that much faster. MRI scans are even slower.

Prof. Fessler is working with doctors in Radiation/Oncology to generate 3D MRI images of tumors while taking into account motion caused by a patient's breathing. This is called 4D MRI. "One of the holy grail problems that my students and I are tackling is motion in medical imaging," he said.

Tracking Breast Cancer Treatment

Comparing two different medical images of the same area, called image registration, is often necessary to track the progress of a disease and the success of a given treatment. However, these images can't easily be compared when taken at different times, or with differing methods, and they typically don't come with any kind of confidence statement as to the accuracy of their alignment. Professors Fessler, Hero, and Scott are working on this issue as part of a major effort headed by Dr. Charles Meyer, Director of the U-M Radiology Digital Image Processing Laboratory. One of the key applications of the research is to track the progress of breast cancer in an individual.

Nuclear Nonproliferation to Environmental Monitoring

Outside the medical field, Prof. Scott is collaborating with Prof. Sara Pozzi of Nuclear Engineering and Radiological Sciences on a project that has the ultimate goal of nuclear nonproliferation. Their research would enable detection of nuclear material at ports of entry. Prof. Pozzi's group uses an inexpensive detector that reacts mainly with neutrons and gamma rays. Scott is helping to classify which are neutrons and which are gamma rays. However, as Scott explains, "what you really want is a detector that only interacts with neutrons, because those are the particles that are characteristic of nuclear sources." Unfortunately, with this particular detector technology, neutrons cannot be measured alone because they are always contaminated by gamma rays. It becomes an intriguing problem that has no known solution in machine learning. Prof. Scott is working to solve this problem.

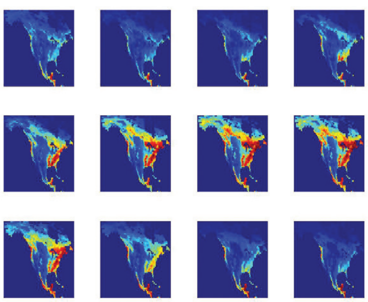

Images reflect carbon-dioxide absorption from the atmosphere over a twelve month period.

He is also collaborating with a professor in Stanford's Department of Global Ecology on a project to determine normal fluctuations in carbon dioxide levels around the world. Taking the data retrieved from satellites, Scott is building a mathematical model of what typical behavior looks like. As deviations arise, earth scientists will determine whether they are from natural (i.e., forest fires, hurricanes) or man-made (i.e., CO^2 emissions) causes.

Better Surveillance From Space

For many decades a technique known as synthetic aperture radar (SAR) has been used to take images of an area from an airplane or satellite. Radar allows for the penetration of clouds, fog and dust particles, but getting a clear picture is challenging.

The highest quality images from SAR can be generated only by knowing within a fraction of a wavelength (i.e., in the millimeter range) where the antenna is positioned in relation to the center of the scene being recorded. Researchers have tried to do this heuristically, but there are some very difficult autofocusing situations which require a more mathematical foundation.

David C. Munson, Jr., the Robert J. Vlasic Dean of Engineering, has been working with his group to develop an autofocus algorithm that will generate the best focused images possible. "We now have a suite of algorithms that we derived based on

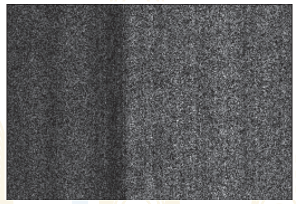

Synthetic aperture radar image that is defocused due to phase errors.

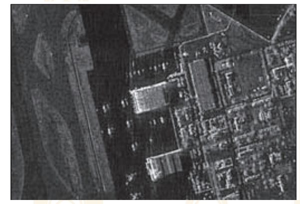

Same image, but properly focused using signal processing techniques.

the fundamental mathematical model of the problem," stated Prof. Munson. "We feel as if we've gotten to the bottom of the autofocus problem."

Computer Vision for Safe Driving and More

The human mind has a remarkable ability to instantly process hundreds of images seen by the human eye, which then impacts individual behavior. When we want a computer to attempt the same thing, we enter the realm of computer vision, and the work of Prof. Silvio Savarese.

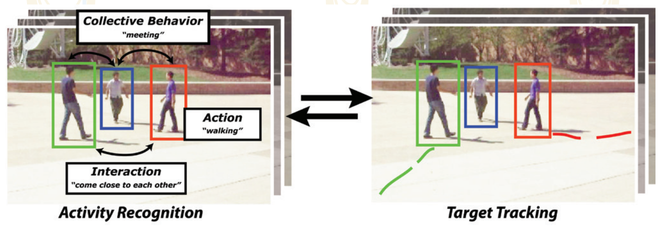

Computer vision algorithms enable tracking of individuals as well as an understanding of what they are doing.

One application of his research is to make cars safer on the road, which has led to several collaborations with the automotive industry. In recent work, Prof. Savarese is helping computers not only track and identify pedestrians with the use of a single uncalibrated and moving camera, but also figure out what they are doing (talking, walking, standing in line, etc). His algorithms allow for tracking even in crowded and changing environments.

Once a computer is able to determine what is going on in a scene, it can use that information for additional applications such as indexing photos and videos online, or surveillance monitoring. Prof. Savarese is also working with Prof. Todd Austin, a computer scientist at Michigan, on a specific application of his research that will help the blind using assistive technologies.

Computer vision techniques are naturally suited to robotic applications, and Prof. Savarese has collaborated with computer science colleague Prof. Benjamin Kuipers as well as individuals in industry to help robotic systems intelligently interact with their environment.

Construction Progress and Safety

Prof. Savarese also applied computer vision techniques to the problem of efficient construction site monitoring. Working with Prof. Feniosky Peña-Mora (Dean of Engineering and Applied Science at Columbia University), and his former student Prof.

Mani Golparvar-Fard (Assistant Professor of Construction Engineering & Management at Virginia Tech), he pioneered a method to automatically track structural changes, and enabled data to be collected simply and inexpensively through a process called Four-Dimensional Augmented Reality (D^4AR). The technology generates a 3D map of the site for a given point in time (time is the fourth dimension). This information can be immediately transmitted to off-site locations, and is expected to greatly facilitate work at construction sites. The three researchers recently co-founded a company to commercialize the technology.

Music Signal Processing

Signal processing technology has transformed the music industry. It is apparent in live performances (microphones, speakers, audio mixers, etc.), and has enabled these performances to be heard on your mobile device or home sound system. In more recent applications, signal processing techniques are being used to transcribe, index, and even classify music.

In related work, Prof. Greg Wakefield has applied signal processing technology to vocal pedagogy. He first mimics the sound of a singer's voice electronically, and then modifies it to approach the sound desired by the teacher. Prof. Wakefield and his collaborator on the project, opera singer and vocal instructor Prof. George Shirley, demonstrated the process at a vocal pedagogy workshop. The vocal instructor said it would have taken six months to get the student to do what was accomplished in minutes using this technique.

Prof. Wakefield would also like us to be able to put on headphones and hear the sounds of the New York Philharmonic, or David Bowie, as if each original sound of the orchestra or band were coming from a unique point in space. He and his student successfully rendered 40 unique sound sources – and they did it with the same computational power as needed for a single source.

He is currently collaborating on spatial-audio techniques with Prof. David Kieras, a computer science colleague who developed EPIC, a cognitive architecture that emulates humans doing complicated tasks. "EPIC has a well-developed visual model," explained Wakefield, "but it's never been able to hear very well. We're trying to help EPIC hear."

Processing Textures for Image Compression

As raw photographic images continue to increase in file size, it is increasingly desirable to be able to save and transmit compressed versions of the images quickly and easily. JPEG is a typical method of image compression, but it does a poor job of rendering textures in an image. A texture can be thought of as a pattern that is repeated with enough similarity to seem identical or closely related from one portion of the pattern to another; examples include a roof, sand, a collection of rocks or flowers, or a knitted sweater.

Using a method called Matched Texture Coding (MTC), Prof. David Neuhoff and a team of researchers that includes Prof. Thrasyvoulos Pappas of Northwestern University have

MTC

JPEG

developed a method to compress images containing textures efficiently while retaining good image quality. When a variety of results were shown to human subjects, the MTC method generated images that were rated equal or superior to those compressed by JPEG and other existing algorithms.

A key to the success of MTC is the development by these researchers of a new metric, called the Structural Texture Similarity Metric (STSIM), for judging the perceptual similarity of two patches of texture. STSIM can also be used for automated indexing and retrieval of images containing texture.

Working Around the Holes in Big Data Applications

Many modern signal processing applications require algorithms that can operate despite missing data. For example, consider applications of computer network monitoring and environmental sensor networks, where we need algorithms to detect anomalous events like bot attacks or high levels of industrial run-off in a river. In both cases, isolated measurements may get lost when a network connection goes down, and the prediction algorithms need to take this into account. Similarly, certain pieces of information may be unknown in collaborative filtering applications where algorithms recommend certain actions be taken – for example when an algorithm recommends a movie to a Netflix user, or recommends a genetic experiment to a medical researcher. The existence of missing data changes the way traditional algorithms operate.

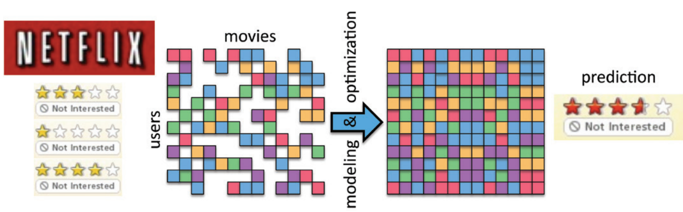

Netflix users enter <1% of the entries that comprise their movie ratings matrix. Algorithms are being designed to predict the rest.

Prof. Laura Balzano is developing novel algorithms that are applicable to these modern signal processing applications, particularly those dealing with massive data sets where some information is not known. Examples include mobile health monitoring, urban sensing, collaborative filtering, computer network monitoring, environmental sensing, electronic medical records analysis, biological networks, and social network mining. She uses tools of optimization and statistical signal processing to develop and analyze algorithms for prediction, detection, and learning from data.

A New Systems Information Theory for Sensing Systems

With today's proliferation of miniature sensors creating sensor networks, there are now a multitude of signals coming from sensors that are communicating with each other as well as a host. New theory is needed to address how information is gathered and interpreted in these current and next-generation sensing systems. As Director of a newly-established Multidisciplinary University Research Initiative (MURI), [see pg. 18] Prof. Hero intends to lay the foundation for a new systems information theory that can be used to design the next generation of autonomous and adaptive sensing systems. Prof. Raj Nadakuditi is an investigator on the project.

Finding When Signals Can be Detected in Data

One basic problem in gathering information from sensors is determining at what point we can differentiate meaningful signals from the noise of irrelevant information. The relationship between the signal and the noise is called the signal-to-noise

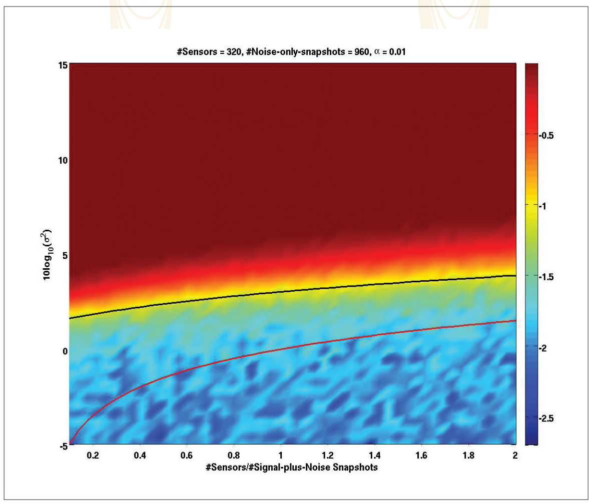

Random Matrix Theory predicts when a signal is detectable in a noisy dataset.

ratio, or SNR. In any given system, there is an SNR threshold above which a signal can be reliably detected, and below which it cannot - but where does that threshold lie? Using random matrix theory, Prof. Nadakuditi has developed a broad theory for predicting where that threshold lies. The theory can also be applied to applications as diverse as radar, sonar, wireless communications, and econometrics, and for designing algorithms that attain this limit.

Focusing Random Matrix Theory on Light and Networks

Prof. Nadakuditi was intrigued when he read published experimental results showing that opaque materials could be used as lenses with tighter foci than cheap glass lenses. Working with colleagues in optics (Prof. Steve Rand) and electromagnetics (Prof. Eric Michielssen), he is currently developing the random matrix theory needed to fully understand the fundamental limits of how much light or energy can be transmitted through mediums that are highly reflective, and is testing his theory in a lab setting. "This is an interesting science problem that could only have been cast and resolved through random matrix theory," stated Prof. Nadakuditi. Potential future applications include biological imaging through tissues.

In addition, he is collaborating with Mark Newman, the Paul Dirac Professor of Physics, on a project that applies random matrix theory to the investigation of complex social, biological, and technological networks. Taking an example from social networking, information about a wide range of individuals' social preferences in Facebook can first be represented in a matrix where a 1 in the matrix represents a connection between two individuals while a zero represents no connection. Prof. Nadakuditi is investigating the limits of when information, such as what types of social communities these individuals might belong to, i.e., book clubs or softball teams, can or cannot be extracted.

In these and a variety of other applications, random matrix theory is proving to be very suitable for these types of high-dimensional streams of information, and is being applied to many disciplines of science, engineering, and finance.

Tech Transfer

Michigan's signal processing research extends far beyond the laboratory walls. Faculty and students are working closely with industry, generating patents, and launching startup companies.

For example, Prof. Savarese recently co-founded the company Vision Construction Monitoring, LLC, with plans to offer his D^4AR modeling technology to the construction industry. Prof. Munson co-founded the company InstaRecon in 2003, which produces software for fast image reconstruction in computer tomography. Their goal is to have the algorithms used in hospital scanners.

Quantum Signal, LLC was co-founded in 2000 by Professor Emeritus Bill Williams and his student Dr. Mitchell Rohde (BSE EE, MSE EE and BioE, PhD BioE). This local company builds products and technologies based on state-of-the-art algorithms and software, and has significant activity in robotics, biometrics, video analysis, and simulation.

Prof. Fessler's research played a part in General Electric's Veo technology for CT images at lower radiation doses, and Prof. Hero's technology allows pathologists at U-M hospital to make improved diagnoses based on an individual's blood signature. In addition, several faculty members publish their software and specially-designed algorithms on their web pages to assist the entire research community.

Finding the Sweet Spot

While signal processing research at Michigan is often driven by real-world applications, our faculty prefer to work on projects that push existing theory into new directions, or even better, that demand new theory. "With the mathematics on our side," said Prof. Hero, "we are able to focus on analysis – predicting performance, and predicting where the next sweet spot is where we can develop algorithms that will do things that haven't been done before."