Syllabus for ECE 566 Detection Theory

Lecturer: Jessica Fridrich

1. This is a bread-and-butter course

The skills you are going to learn in this course will be useful in numerous other courses, including courses on Cryptography, Steganography, Pattern Recognition, Estimation Theory, Machine learning, Information Theory, and general Signal and Image Processing.

2. Wide applications across numerous fields

Detection theory is about deciding what type of signal we are observing through noisy measurements. Let us say, we have a set of observables (a vector of scalar or vector measurements) and we want to know if the data contains a given known signal. This is a classical problem when building any type of a radar (sonar, lidar, etc.). We may also be interested if our observables contain a signal that we do not know exactly but only through its statistical properties (e.g., periodicity, variance, mean, spectrum, etc.) Of course, the main problem is that our observations are inherently noisy and in many cases the signal we are trying to detect is **very weak**. We are thus interested in building the *best possible* detector rather than an *ad hoc* one. To this end, we employ well-developed methods of statistical hypothesis testing, estimation theory, and information theory in order to be able to build the best possible detectors and detectors that are robust.

There are so many applications of detection theory that it is impossible for me to list them all. Astronomers want to detect the presence of periodic dips in star brightness to detect the presence and estimate the size of extra-solar planets similar to the Earth (c.f., NASA's Kepler Project). Forensic analysts want to know if a given image was taken with a specific camera when prosecuting a child pornographer. An analyst wants to know if a given image has been modified and where (identify a digital forgery). In medical imaging, doctors are interested in registering abnormalities that might be indicative of a disease, etc, etc.

3. Core skills needed to succeed in this course

As you might have already guessed, the tool of the trade in this course will be *statistic*. You need to have a **solid** background in basic statistic and be comfortable using it in practice. We will rely on basic calculus, such as integration by parts and integration using a substitution. I will spend three full lectures refreshing your memory on basic statistic concepts to make sure that you are up to the task to handle this course. We will also delve a little into Functional Analysis, which will help us obtain really important and powerful tools. Besides, I hope that providing you with a hint of an abstract approach will help you connect concepts that you are already familiar with but considered them as unrelated. A good engineer knows how the car works so that when it breaks or needs to be updated/modified, s/he knows what to do. Just be able to drive the car is nice but gets you only a limited distance. I hope you see the analogy here.

4. Teaching by hands-on projects

I prefer teaching by examples and hands-on projects as opposed to dry abstract lectures supplied with artificially conceived examples (drills). In the beginning of the course, I will introduce two interesting practical applications that require Detection Theory and that are usually a lot of fun for the students: Digital Forensics and Digital Watermarking. The first concerns problems of detecting very subtle signals in digital images that enable us to uniquely identify the camera that took the image and even tell which parts of the picture were modified. This technology was developed here in Binghamton in late 2000's and has found numerous applications in criminal justice and military intelligence gathering. Students will analyze their own digital cameras and learn about digital image acquisition and processing and signal detection. The course will culminate through a final project in which students will be asked to use the detection and estimation theory to identify their own images among dozens of others.

The second hands-on project is robust digital watermarking. A robust digital watermark is a subtle spread-spectrum signal that is added to an image to either identify the author, the customer, to augment the content with auxiliary information, etc. This is a cutting-edge technology that is currently being used by broadcasters for digital rights management and for various business applications.

I hope you will find these specific applications exciting enough to motivate you to learn and invest in studying the statistical concepts necessary for developing the right tools.

5. Course structure

First, I will introduce the topic, grading policy, the course structure, and outline its content. There will be a motivational show-off to get you excited about the potential of signal detection in digital image forensics and digital watermarking.

The purpose of the next two-three lectures is to get everyone on the same knowledge level of statistics. After this, I begin following the material in the textbook, sometimes adding additional material on digital cameras, digital watermarks, estimation theory, and some special topics.

This course will involve:

(1) Classical lectures
(2) Optional self-study from on-line materials
(3) Classroom discussions and brainstorming
(4) A series of homework assignments
(5) Independent investigations
(6) Final project

6. Grading

The final grade will be computed as an average of grades from homework assignments and the final project.

7. Academic honesty

All students must adhere to the Student Academic Honesty Code of the University and the Watson School. The Department of Electrical and Computer Engineering has adopted a standard policy to enforce these codes for violations involving course work. Category I violations result in a grade of 0 for the graded work plus a one letter course grade reduction. A *Report of Category I Academic Dishonesty* form is filed with the Provost's Office; if a prior report is already on file, the offense is automatically elevated to Category II. Category II violations result in at least a failing grade for the course plus any additional penalties determined by the Watson Academic Integrity Committee.

University Academic Honesty Code:
http://bulletin.binghamton.edu/program.asp?program_id=826
Watson School Academic Honesty Code:
http://www.binghamton.edu/watson/Watson_Academic_Honesty_Policy.pdf
ECE Department Academic Honesty Code Enforcement Policy:
http://www.ece.binghamton.edu/documents/Academic_Honesty_Policy.pdf