Addendum to the syllabus for Math 39000 for students enrolled in Math 59800: Advanced Linear Algebra with applications (Course No: 31655)

Meets: MW 1:30-2:45p in BS 3014
Final Exam: May 5, 1:00-3:00p

Instructor: Carl Cowen **Office:** LD 224P **Phone:** 278-8846
Office Hours: M 10:30-11:30a and 3:00-4:00p, W 10:30a - 12:30p, or by appointment
E-mail: `ccowen@iupui.edu`
URL: `http://www.math.iupui.edu/~ccowen/Math390.html`
URL: `http://www.math.iupui.edu/~ccowen/Math598.html`

This course will be a graduate version of Math 39000, Linear Algebra II, and will be an extension of that course for graduate credit. Students in Math 59800, Advanced Linear Algebra, will complete all the required work for Math 39000, including the homework and tests, and will attend and participate fully in the same classes as the Math 39000 students. In addition, students in Math 59800 will master the material of Chapter 11, Singular Value Decomposition, and do a small project on the application of SVD to image compression. This additional material and other differences for Math 59800 students are described below in more detail. It is important to read this document in conjunction with the Syllabus for Math 39000 to get a complete picture of the course.

The official text will be

Text: *Linear Algebra for Engineering and Science,*

by Carl Cowen (ISBN 0-9650717-4-X)

with supplemental material for the application to cost accounting and possibly other topics. Books on reserve in the library covering some of the topics of the course are the text and:

Linear Algebra and Its Applications, by Gilbert Strang.

Linear Algebra Done Right, by Sheldon Axler.

Math 351 (or Math 511 or the equivalent, such as a first course from another institution or considerable experience in applying linear algebra) is a prerequisite for this course and I will assume you know the material in that course. Some of the important ideas from such courses include linear independence, basis, and rank; the (ten or so) equivalent conditions for invertibility of a matrix; inner products and the Gram–Schmidt algorithm. (Although eigenvalues and eigenvectors are usually introduced in Math 351 and Math 511, we will begin at the beginning of this topic.)

There will be two mid-term tests, the same tests as for the Math 39000 students, each counting about 15–20% of your grade, and about 35% of your grade will come from the two–hour final exam given during Final Exam week (May 5). Your work on Chapter 11 and the short project will count about 15% of your grade and the weekly homework will make up the remaining 10–15% of your grade. Make-up/late homework will **not** be graded for credit.

Chapter 11 covers the Singular Value Decomposition, a very important tool in applications such as image compression (JPEG is based on SVD), search engines (the *Google* search engine is based on SVD), and many other topics. After you have read the material in Chapter 11, we will arrange a time to discuss and clarify this material outside of regular classtime. In addition, there will be separate homework on this material to ensure that you have mastered it (worth about 5% of the course grade). Finally, the small project on image compression, and its accompanying short report, will be worth about 10% of the course grade. You can tackle this material after we have completed our work on Hermitian and normal matrices (Chapter 9), about midway through the course.

The small project that I have in mind, an application of SVD to image compression, is an extraordinarily naive approach to that goal $\cdots$ easy, but surprisingly effective; it is not meant to be an examination of algorithms that are widely used in practice, but rather just enough to give you a feeling for power of the most basic idea. While you will have considerable flexibility in creating your project, here is an outline of a minimal project. First, you should choose or create two (rectangular) black-and-white digital images of about 10,000 pixels each. (Note: my first image in such an experiment was a black block letter 'A' on a white background; it was easy to create but complicated enough to give interesting results.) Then, you associate each image (in the most obvious way) with a matrix such that each black pixel is associated with (say) a 0 in the corresponding matrix position and each white pixel is associated with a 1. If you are given the singular value decomposition of such a matrix, then you can recreate the image by recomputing the matrix and displaying it. Instead, you will use the SVD to create lower rank approximations of the image matrix, and use these approximations to create new images. The project consists of doing this for a significant number of different ranks (say ten or twenty) and then, for each rank, judging the effectiveness of the compression and the amount of compression (as a percentage). Your report should include the original images you used, a description of your work, such as how you obtained the compressed images, and including your conclusions about what rank (and corresponding amount of compression) corresponds to the most compressed image that is nearly indistinguishable from the original, a "good" yet highly compressed image, and the point at which you judge all useful information about the image has been lost. Of course, you should include the compressed images in your report to allow comparison with the originals. Naturally, the report should be 'typed', written in good English, include an appropriate list of references, and so on.