

# CS 290H: Sparse matrix algorithms // Homework 0

Due by class time Thursday, September 30

You may do the programming assignments for this course in C or Fortran; I recommend C. In each case, you will set your code up with an interface so that it can be called from Matlab as a “mexFunction”. This will let you use Matlab to debug your code, and it will let me use Matlab to test it. This warmup assignment is just for you to learn how to write a Matlab interface using sparse matrices.

**Problem 1:** Write a Matlab mexFunction, `sparseprint(A)`, which prints the contents of a sparse matrix `A` as a list of nonzero entries like “`A(7,8) = 13.4`”.

Your routine can assume the matrix is real, but should not assume that it’s square. Test it from Matlab with several matrices you make up. See the Matlab functions `sprand` or `sprandn` to generate random sparse test matrices. Or download the Matlab interface to the UF Sparse Matrix Collection, at <http://www.cise.ufl.edu/research/sparse>.

For documentation on mexFunctions, open up a help window in Matlab, and look in the *External Interfaces/API* section, under the *Matlab* section. The C mexFunction syntax is quite a bit nicer than the Fortran. All mexFunctions have the routine name `mexFunction`.

You’ll need to use the following Matlab “mx” and “mex” routines:

---

<code>mxGetM</code>	returns the number of rows of a matrix
<code>mxGetN</code>	returns the number of columns of a matrix
<code>mxGetJc</code>	returns a pointer to the column pointer array ( <code>Ap</code> )
<code>mxGetIr</code>	returns a pointer to the row indices ( <code>Ai</code> )
<code>mxGetPr</code>	returns a pointer to the numerical values ( <code>Ax</code> )

---

There is a more complex but very well-written example of mexFunctions for sparse matrices in Tim Davis’s LDL code, which is at <http://www.cise.ufl.edu/research/sparse/ldl> and also on the CS290H web page. LDL computes a sparse Cholesky factorization; you don’t need to understand how that works (yet), but you can see from that code how a C program can get access to a Matlab sparse matrix.

I recommend the following strategy for writing mexFunctions. Suppose the name of the mexFunction is `foobar`. Put all Matlab-related work in a file called `foobarmex.c` (or `foobarmex.f` for Fortran). The routine name in this file is `mexFunction`. All calls to `mx*` and `mex*` routines should be placed here. This routine then calls a routine in the file `foobar.c` or `foobar.f` that does the actual work. There should be no references to `mx*` or `mex*` routines in the `foobar.c` file, with the exception of calls to Matlab functions that have ANSI C counterparts (such as `malloc`, `free`, and `printf`).

Why bother? Because this makes the `foobar.c` file a stand-alone routine that can then be used in a stand-alone C program (or Fortran, as the case may be), without using Matlab. OK, for such a simple routine as `sparseprint`, this structure is overkill. But it’s a useful habit anyway. Tim’s LDL program follows this structure.