

Measuring and Improving Code Performance

1 Measuring Performance

Matlab offers two main tools for measuring the performance of code and finding out its bottlenecks. The first are the timing functions, which can be used to measure how long pieces of code and functions take to run. The second is the profiling tool of Matlab, which can measure the performance of large pieces of code (several functions) and displays results in a nice interface.

1.1 Timing

We can time pieces of code with the functions `tic` and `toc`. When `tic` is executed, it starts an internal stopwatch and returns a number (optional) representing the time when `tic` was executed. The function `toc` displays how much time has elapsed in seconds since the function `tic` was executed. Subsequent calls to `toc` display how much time has elapsed since the first `tic` has been executed. We can start a new timer by calling `tic` again.

We can use `tic` and `toc` to time pieces of code.

```
1 % measure how long it takes to execute the code below
2 % start timer
3 tic
4 n = 10000;
5 a = rand(n, n);
6 b = rand(n, 1);
7 a*b;
8 toc
9
10 % we can store the time in a variable
11 t = toc;
12
13 % measure time depending on size of matrix
14 n_values = 1:500:10000;
15 t = zeros(length(n_values), 1);
16 for i = 1:length(n_values)
17     tic;
18     n = n_values(i);
19     a = rand(n, n);
20     b = rand(n, 1);
21     a*b;
22     t(i) = toc;
```

```

23 end
24 plot(n_values, t)

```

We can create more than one timer by saving the output of `tic`, which can then be passed to `toc` to measure the time difference with respect to the respective `tic`.

```

1 total_time = tic;
2 repetitions = 100;
3 times = zeros(length(repetitions), 1);
4 for rep = 1:repetitions
5     t = tic;
6     % do computation
7     mean(randn(10000000, 1));
8     % store time
9     times(rep) = toc(t);
10 end
11 % average time
12 disp(sprintf('Average time: %.2f seconds (%d repetitions)',
13     ...
14     mean(times), length(times)));
15 % total time
16 disp(sprintf('Total time: %.2f seconds', toc(total_time)));

```

The functions `tic` and `toc` are good to time small pieces of code, but to get a more precise time estimation it is necessary to re-run the same code several times and average the `tic` and `toc` results.

A more rigorous timing measure can be obtained with the function `timeit`. We can use `timeit` to measure the typical time it takes to run a function. Given a function, `timeit` runs it several times and reports the median time it takes to complete its execution.

Create a test function:

```

1 % test_timeit.m
2 function [a, b] = test_timeit(X)
3 beta = rand(size(X, 2), 1);
4 a = X*beta;
5 if nargin > 1
6     b = exp(a);
7 end

```

Use `timeit` to time the execution of the function:

```

1 % specify the input to test_timeit
2 X = rand(5000);
3 timeit(@() test_timeit(X))
4 % specify how many outputs the function being timed should
   return
5 timeit(@() test_timeit(X), 2)

```

1.2 Profiling

The command `profile` can be used to profile large pieces of code. We can start the profiler by executing `profile on`. Then, any code that we run afterwards will be automatically timed. We can then call `profile viewer` to stop the profiler and see the results in a new window. Let's first create a test code:

```

1 % test_profiler.m
2 clear all
3 clc
4 % Load stock data
5 data = readmatrix('AAPL.csv');
6 % Compute intraday log-returns
7 prices = reshape(data(:, 3), 78, []);
8 returns = diff(log(prices));
9 % Compute realized variance
10 RV = sum(returns.^2);
11 % Generate bootstrap samples to compute confidence interval
12 total_samples = 10000;
13 RVs = zeros(total_samples, length(RV));
14 pbar = waitbar(0, 'Bootstrapping'); % progress bar
15 for i = 1:total_samples
16     % generate bootstrap indices
17     indices = unidrnd(size(returns, 1), size(returns, 1), 1);
18     % bootstrap sample of the returns
19     returns_sample = returns(indices, :);
20     % compute statistic for each day
21     RVs(i, :) = sum(returns_sample.^2);
22     % generate progress bar
23     waitbar(i/total_samples, pbar);
24 end
25 close(pbar);
26 % Compute confidence interval
27 CI = quantile(RVs, [0.005 0.995]);
28 % Generate timestamps for plotting
29 dates = data(1:78:end, 1);
30 years = floor(dates./10000);
31 months = floor((dates - years.*10^4)./100);
32 days = floor(dates - years.*10^4 - months.*10^2);
33 sdates = datenum(years, months, days);
34 % Plot realized variance with correct x-ticks
35 plot(sdates, RV, 'k');
36 datetick('x', 'yyyy');
37 grid on;
38 % Shade confidence interval region
39 hold on;
40 f = fill([sdates; flip(sdates)], [CI(:, 1); flip(CI(:, 2))],
41         'k');
42 f.FaceAlpha = 0.2;

```

```
42 f.EdgeAlpha = 0;
```

Now, we can use the profiler to analyze the code for bottlenecks. Clear the workspace before running the profiler.

```
1 clear all
2 profile on
3 test_profiler
4 profile viewer
```

The command `profile viewer` displays the window shown in Figure 1.

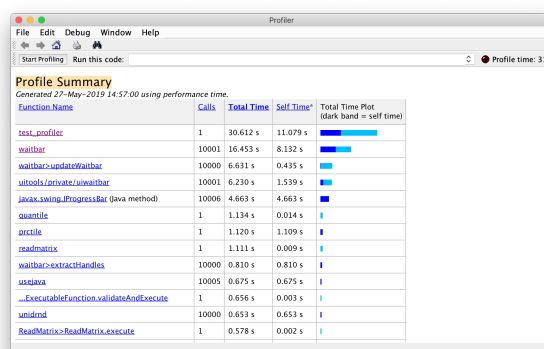


Figure 1: Profile Summary Window.

The profiler window shows all functions called while the profiler was on. The functions are sorted by total executing time. You can click on each function to see a more detailed break down of the execution. Click on the `test_profiler` function, and you should see a window similar to what Figure 2 displays.

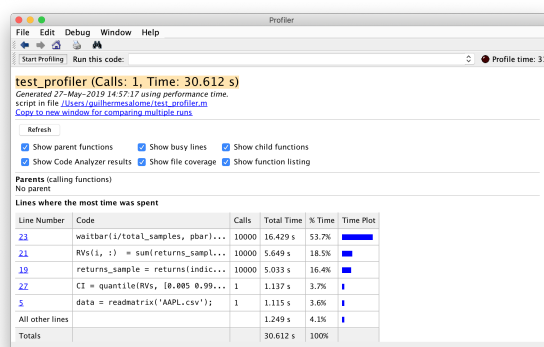


Figure 2: Profiler Detail Window.

We see that the `test_profiler` function took around 30 seconds to finish running. The window also shows the lines where most of the time was spent. We see that out of the 30 seconds, almost 5 seconds were spent on generating bootstrap samples from the original returns. That line was executed 10000 times. You can click on the line number to see the code around it. Click on the button for line 19. You should see a window similar to what is displayed in Figure 3.

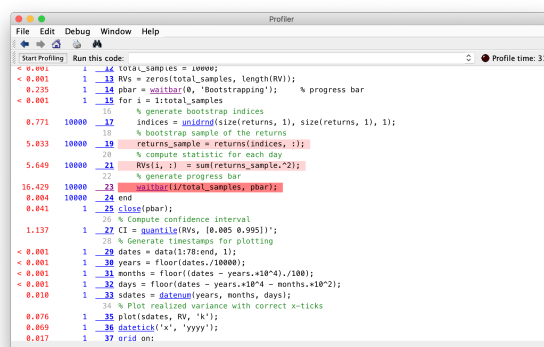


Figure 3: Profiler Line by Line Timing.

The lines displayed in darker red are the ones that take longer to execute. Lines 19 and 21 take about the same time to execute, while updating the progress bar takes the longest. We can use this information to identify lines that are inefficient, update the code and test again.

We can save the profiler results by executing the command `profsave`, which stores the results in a folder named `profile_results`. The file `file0.html` contains the same information as the profiler summary window, displayed in Figure 1.

Using the profiler is useful when working with a data set that is large. The idea is to use a small subset of the data to develop your code, then time it to see how long it takes. Then, you extrapolate how much time it would take to run your code on the entire data set. If the total time is acceptable then you are done. However, if it is not, then you should profile the code and see where you could improve it before running on the entire data set.

2 Efficiency Tips

We will now cover some general tips for efficiency when coding. Let's use the functions `tic` and `toc` to check the difference in execution time.

Avoid creating temporary objects:

```

1 times = zeros(1000, 1);
2 for i = 1:1000
3     tic
4     %% Code to time
5     A = 2*randn(10000, 100);
6     B = A.^2;
7     % store time
8     times(i) = toc;
9 end
10 % report median
11 disp(sprintf('Median execution time: %.4f seconds', median(
    times)));
12
13 times = zeros(1000, 1);
14 for i = 1:1000

```

```

15     tic
16     %% Code to time
17     B = (2*randn(10000, 100)).^2;
18     % store time
19     times(i) = toc;
20 end
21 % report median
22 disp(sprintf('Median execution time: %.4f seconds', median(
    times))));

```

The object stored in A is temporary, created at the beginning of each loop, then used to compute B, and then reassigned when a new loop starts. We can avoid creating and deleting A by moving the `randn` command into B.

Pre-allocate matrices whenever possible.

```

1 %% Append to vector
2 tic
3 a = [];
4 for j = 1:10000
5     a = vertcat(a, randn(1));
6 end
7 disp(sprintf('Time: %.4f seconds', toc));
8 %% Pre-allocate and then modify vector
9 tic
10 a = zeros(10000, 1);
11 for j = 1:10000
12     a(j) = randn(1);
13 end
14 disp(sprintf('Time: %.4f seconds', toc));

```

Creating a matrix and then extending it is slower than pre-allocating the matrix. Extending the matrix means Matlab has to allocate more and more memory at each loop, recreating the vector `a` at each iteration. This reduces the efficiency of the code and makes it slower. Pre-allocating the matrix allows for a single block of memory to be allocated to the vector `a`.

Use sparse matrices to save memory when dealing with a large number of zeros.

```

1 A = eye(10000); % not sparse
2 B = sparse(eye(10000)); % sparse
3 whos A B;

```

Sparse arrays only stores the nonzero elements in memory. It is possible to visualize the sparsity of a matrix with the function `spy`.

Vectorize operations if possible:

```

1 m_size = 1000000;
2 x = randn(m_size, 1);
3 y = rand(m_size, 1);
4 %% Not vectorized
5 z = zeros(m_size, 1);
6 tic

```

```

7  for i = 1:m_size
8      z(i) = x(i)*y(i);
9  end
10 disp(sprintf('Time: %.4f seconds', toc));
11 %% Vectorized
12 clear z
13 tic
14 z = x.*y;
15 disp(sprintf('Time: %.4f seconds', toc));

```

A vectorized operation can run faster than a for-loop if it uses a built-in function. It is also more self-documenting than a for-loop and often leads to fewer lines of code.

3 Generating and Using C Code with Matlab

We can execute code written in C, C++ or Fortran with Matlab. To do so, we need to create a binary file named MEX, which can then be called in Matlab as a regular function. To create MEX files, we will use a Matlab tool that can convert Matlab code to some other language. For this section you will need to install Matlab Coder (or learn how to code in C or C++). You can also execute `coder` on Matlab to get a prompt to install it.

We will learn how to generate the MEX files with an example. Write the following function:

```

1  % euclidean_dist.m
2  function [distances, point_min, point_max] = euclidean_dist(x
    ,points)
3  % Compute Euclidean norm between x and the points
4  distances = sum((x-points).^2).^0.5;
5  % Obtain points that maximize and minimize the distance
6  [~, id_min] = min(distances);
7  [~, id_max] = max(distances);
8  % Save these points
9  point_min = points(:, id_min);
10 point_max = points(:, id_max);
11 end

```

This function takes a column vector `x`, and a matrix containing multiple column vectors `points`, then computes the Euclidean distance between `x` and the other points. The function has three outputs, the distance between all points, the vector closest to `x` and the vector farthest from `x`.

Now, we need to use Matlab's Code Analyzer. To do so, add the comment `%#codegen` right after the definition of the function `euclidean_dist`:

```

1  function [distances, point_min, point_max] = euclidean_dist(x
    ,points) %#codegen

```

This will trigger the Code Analyzer to check for any issues that could prevent the Matlab code from being converted to C/C++. If there are no errors, a green square will appear on the upper right-hand side of the Editor window. If there are errors, then the lines with issues will be underlined in red. You can hover the mouse on top of the line to see a

message explaining the issue and how to fix it. Most of the features of the programming language are supported by the converter, but some are not. For a full list of the supported and unsupported features, see this reference page.

Figure 4 displays the results of the Code Analyzer for the function `euclidean_dist`.




Figure 4: Code Analyzer Without Errors.

Now, we will use the function `coder.screener` to call the Code Generation Readiness Tool. This tool also checks for possible issues that can arise when generating the C/C++ code.

```
1 coder.screener('euclidean_dist')
```

Figure 5 shows the output of the code above. It indicates the function we created should not lead to issues when being converted to C/C++ code.




Figure 5: Output of Code Generation Readiness Tool.

We can now generate the MEX file with the `codegen` command. The MEX file is based on C code, which is statically typed. This means we need to define the type (class) and size of the inputs to the function when generating the lower-level code. We start with the case where the size of the inputs is fixed.

3.1 Fixed-Size Inputs

The type and size of inputs can be inferred when generating the lower-level code by giving an example to the command `codegen`.

```
1 % create example for codegen
2 x = [1; 2; 3]; % size 3 x 1
3 points = rand(3, 10); % size 3 x 10
```

To create the MEX file, we call `codegen` by passing it the name of the file we want to convert and the example inputs:


```
1 codegen -report euclidean_dist.m -args {x, points}
```

The option `-report` generates a report that we can use to debug any issues that arise during the conversion. After the `-report`, we have the file name, in this case `euclidean_dist.m`, and then the option `-args` and a cell containing the example inputs. The `codegen` command infers the type and size of the inputs from the type and size of `x` and `points`.

When we execute the code above you will see an error. The error is depicted in Figure 6.

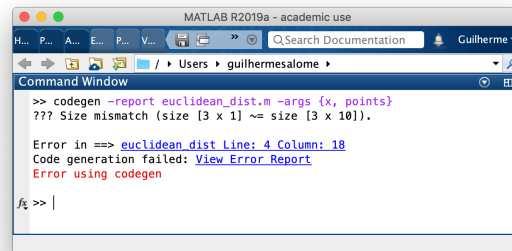


Figure 6: Codegen Error.

The error says there is some type of mismatch in the code. We can get more detail by clicking on `View Error Report`, which will open the window depicted in Figure 7.

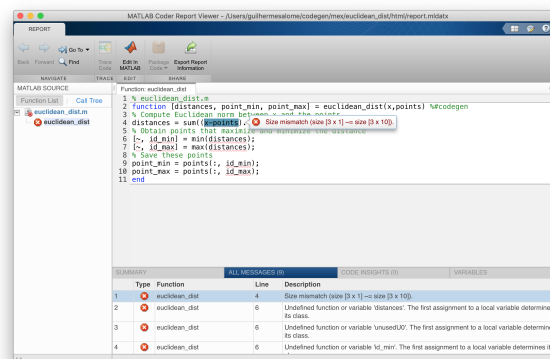


Figure 7: Codegen Report and Error Message.

The error report shows that there is a mismatch in size between `x` and `points`. It is not an issue in the Matlab code, because when we execute `x=points`, the vector `x` is broadcasted to have the same size of `points`. However, it leads to an issue when converting the code to C. We can fix it, by using `repmat` on the vector `x` to avoid the mismatch:

```
1 % Modify the line:
2 distances = sum((x-points).^2).^0.5;
3 % To:
4 distances = sum((repmat(x, 1, size(points, 2))-points).^2)
    .^0.5;
```

We can now compile the code by running:

```
1 codegen -report euclidean_dist.m -args {x, points}
```

The code generation should complete successfully. A file named `euclidean_dist_mex.mexmaci64` will be created. This file contains the binaries that can be executed by Matlab as a regular function. We can compare the time it takes to run the original function and the compiled function:

```
1 % use timeit to compare execution time of the functions
2 sprintf('Original Function: %.6f seconds', timeit(@()
   euclidean_dist(x, points), 3))
3 sprintf('MEX Function: %.6f seconds', timeit(@()
   euclidean_dist_mex(x, points), 3))
```

In this case, there is little gain since vectorized and simple operations in Matlab are already very fast. However, the gains in speed will increase when loops are involved.

Notice that if you try to execute `euclidean_dist_mex` with a vector `points` of a size different than the one specified, you will get an error.

```
1 points = rand(3, 5);
2 euclidean_dist_mex(x, points) % error
```

Next, we modify the MEX file to accept an input that can vary in size.

3.2 Variable-Size Inputs

To have the MEX function accept inputs of varied sizes, we use the function `coder.typeof` to specify the size of the inputs. The function `coder.typeof` takes three inputs:

1. An example for the input (like `x`);
2. A vector that specifies an upper bound for each dimension of the input;
3. A vector that specifies whether each dimension of the input is variable in size or fixed in size.

Suppose we want to allow `x` to have up to 10 rows, but always one column. And we want `points` to have up to 10 rows, and up to 10 columns.

```
1 x = [1; 2; 3];
2 example_x = coder.typeof(x, [10 1], [1 0]);
3 % [10 1] specifies the upper bound for the dimensions of x
4 % [1 0] specifies that the size of the first dimension is
   variable,
5 % while the second is fixed
6 points = rand(3, 10);
7 example_points = coder.typeof(points, [10 10], [1 1]);
8 % [10 10] specifies the upper bound for the dimensions of the
   points
9 % [1 1] specifies that both dimensions are variable
```

We can now compile the code as before with `codegen`, and use `example_x` and `example_points` as the examples, instead of `x` and `points`:

```
1 codegen -report euclidean_dist.m -args {example_x,
    example_points}
```

This will overwrite `euclidean_dist_mex` with the new version. We can test it:

```
1 x = rand(7, 1);
2 points = rand(7, 3);
3 euclidean_dist_mex(x, points)
```

We can further modify the code to allow `x` to have as many rows as we want, but always one column. And `points` to have as many rows and columns as we want.

```
1 example_x = coder.typeof(x, [Inf 1], [1 0]);
2 example_points = coder.typeof(points, [Inf Inf], [1 1]);
3 codegen -report euclidean_dist.m -args {example_x,
    example_points}
```

This will overwrite `euclidean_dist_mex` with the new version. We can test it:

```
1 x = rand(23, 1);
2 points = rand(23, 40);
3 euclidean_dist_mex(x, points)
```

3.3 Generating C Code

If you have a supported C compiler installed on your machine, you can also generate the C code from the Matlab files. You do so by adding the option `-config:lib` when calling `codegen`:

```
1 codegen -d matlab_to_c -report euclidean_dist.m -args {
    example_x, example_points}
```

Notice the use of the option `-d matlab_to_c`, which stores the resulting files in the folder named `matlab_to_c`.

3.4 Observations

The MEX files are platform specific, so if you compile a file in Windows, for example, then you will not be able to re-use it in Linux. However, if you have the original `.m` file you can re-compile it again in a different platform.

It is not possible to generate MEX files from Matlab scripts, only for functions. If you need to convert a script, then you must write it in a function format.

If you are often converting a file to MEX, it is good practice to create a build script that will do the conversion. For example:

```
1 % generate_mex_euclidean_dist.m
2 % Generates .mex file for: euclidean_dist.m
3 % Define inputs
4 example_x = coder.typeof(x, [Inf 1], [1 0]);
5 example_points = coder.typeof(points, [Inf Inf], [1 1]);
6 % Call codegen
```

```
7 | codegen -report euclidean_dist.m -args {example_x,  
   |     example_points}
```

Then, we can execute `generate_mex_euclidean_dist` on Matlab to generate the MEX file.

4 Parallelization

We can speed up the execution of code in Matlab by taking advantage of the multiple cores available in modern computers. Matlab offers a relatively easy way to make for-loops run in parallel on multiple cores using `parfor`. To make use of multiple cores, you will need to install the Parallel Computing Toolbox. It makes sense to make a for-loop run in parallel if:

- Each iteration of the loop takes a long time
- The loop iterations are independent from one another

We can use the `parfor` functionality to work with embarrassingly parallel problems. A problem is called embarrassingly parallel if it can be broken into smaller pieces that have little to no dependency between one another.

The syntax for `parfor` is similar to the syntax of for-loops:

```
1 | parfor i = 1:10  
2 |     % code to be executed at each iteration  
3 |     disp(i)  
4 | end
```

However, the loop is executed in parallel and the order of execution is not deterministic. Additionally, the code executed within the loop must not depend on previous iterations of the loop.

If you execute the code above, Matlab will start a pool of workers and then execute the loop in parallel. The pool of workers takes some time when it is first initialized, but can be used afterwards without slow downs. The workers are associated to the cores in the computer's processor. The more cores available, the more workers can be added to the pool. It is possible to set the number of workers to be used in the settings panel for the Parallel Computing Toolbox.

The code above displays the value of the variable `i` on each iteration. Notice that the values are displayed out of order. This happens because the iterations are executed in a non-deterministic fashion.

4.1 Converting a For-Loop into a Parfor-Loop

As an example, let's convert a `for`-loop that bootstraps the confidence interval for the mean into a `parfor`-loop. First, we begin with the data generation and bootstrapping in an usual for-loop:

```
1 | % true mean  
2 | mu = 10.35;  
3 | % generate random sample  
4 | sample = normrnd(mu, 2, 1000, 1);
```

```

5 % bootstrap
6 total_iter = 100;
7 means = zeros(total_iter, 1);
8 for i = 1:total_iter
9     % create bootstrap sample
10    b_sample = sample(randi(length(sample), length(sample),
11    1));
12    % compute statistic
13    means(i) = mean(b_sample);
14 end
15 % display 99% confidence interval for the mean
16 CI = quantile(means, [0.005 0.995]);
17 sprintf('99% Confidence Interval for Mean: [%.4f, %.4f]', CI
18 )

```

Let's make the code above into a function, so that we can specify the sample size and the number of bootstrap samples.

```

1 % bootstrap_mean_CI.m
2 function CI = bootstrap_mean_CI(varargin)
3 parser = inputParser;
4 addOptional(parser, 'sample_size', 1000);
5 addOptional(parser, 'bootstrap_samples', 1000);
6 parse(parser, varargin{:});
7 sample_size = parser.Results.sample_size;
8 bootstrap_samples = parser.Results.bootstrap_samples;
9 % true mean
10 mu = 10.35;
11 % generate random sample
12 sample = normrnd(mu, 2, sample_size, 1);
13 % bootstrap loop
14 means = zeros(bootstrap_samples, 1);
15 for i = 1:bootstrap_samples
16     % create bootstrap sample
17     b_sample = sample(randi(length(sample), length(sample),
18     1));
19     % compute statistic
20     means(i) = mean(b_sample);
21 end
22 % display 99% confidence interval for the mean
23 CI = quantile(means, [0.005 0.995]);
24 end

```

Time the function:

```

1 sprintf('Elapsed Time: %.4f seconds', timeit(@
2     bootstrap_mean_CI))
3 % increase sample size
4 fun = @() bootstrap_mean_CI(10000);
5 sprintf('Elapsed Time: %.4f seconds', timeit(fun))

```

```

5 % increase number of bootstrap samples
6 fun = @() bootstrap_mean_CI(10000, 5000);
7 sprintf('Elapsed Time: %.4f seconds', timeit(fun))

```

When we increase the original sample size, the computation of the mean statistic starts taking more time. When we increase the number of bootstrap samples, we increase the number of iterations in the bootstrap loop.

Observe that the iterations in the bootstrap loop are independent from one another. Also, each iteration might take a while, specially if the sample size is big. Therefore, we have an embarrassingly parallel problem. We can use `parfor` to speed up this computation. In this case, all we need to do is change the for-loop into a `parfor`-loop:

```

1 % bootstrap_mean_CI_parallel.m
2 function CI = bootstrap_mean_CI_parallel(varargin)
3 % change only the line
4 for i = 1:bootstrap_samples
5 % into
6 parfor i = 1:bootstrap_samples

```

We can execute the code once to get the pool of workers started. Then, we time the execution and compare to the non-parallel version.

```

1 sprintf('Non-Parallel: %.4f seconds\nParallel: %.4f seconds',
...
2         timeit(@bootstrap_mean_CI), timeit(@
          bootstrap_mean_CI_parallel))
3 % increase sample size
4 fun = @() bootstrap_mean_CI(10000);
5 pfun = @() bootstrap_mean_CI_parallel(10000);
6 sprintf('Non-Parallel: %.4f seconds\nParallel: %.4f seconds',
...
7         timeit(fun), timeit(pfun))
8 % increase number of bootstrap samples
9 fun = @() bootstrap_mean_CI(10000, 5000);
10 pfun = @() bootstrap_mean_CI_parallel(10000, 5000);
11 sprintf('Non-Parallel: %.4f seconds\nParallel: %.4f seconds',
...
12         timeit(fun), timeit(pfun))

```

Notice that there is no improvement in speed for the cases where the sample size is small and the number of bootstrap samples is small. This can happen because there is some overhead involved in making a for-loop run in parallel. Indeed, each worker must have access to the data being re-sampled. If copying this data over takes longer than the computations in the loop, then the parallel loop will run slower. When the sample size increases, the computation in the loop starts becoming slower than copying the data to the worker, and running the loop in parallel starts becoming quicker. When the number of bootstrap samples increases (number of iterations in the loop), the speed gain from the `parfor`-loop becomes relevant.

4.2 Requirements for Parfor-Loop

There are a few requirements to use `parfor`. First, the looping variable must take integer values, and the values must be consecutive and increasing:

```
1 % does not work: non-consecutive integers
2 parfor i = -5:2:5
3     disp(i)
4 end
5 % does not work: not integers
6 parfor i = 0:0.2:1
7     disp(i)
8 end
9 % does not work: not increasing
10 parfor i = 10:-1:1
11     disp(i)
12 end
13 % works: consecutive and increasing integers
14 parfor i = -5:5
15     disp(i)
16 end
```

Second, when Matlab encounters a `parfor`, it classifies all of the variables inside the loop in 5 different categories. The categories describe what each variable is doing inside the loop. The categories are:

- Loop variables: loop indices in the loop;
- Sliced variables: arrays that are sliced inside the loop and used by different iterations;
- Broadcast variables: variables defined outside the loop and used inside it, but that are not assigned inside the loop;
- Reduction variables: variables that accumulate values across iterations of the loop
- Temporary variables: variables created inside the loop and not used outside the loop

If a variable cannot be uniquely classified into one of the five categories, then Matlab will give an error. For example:

```
1 x = [1; 2];
2 out = zeros(10, 1);
3 parfor i = 1:10
4     x(1) = 5;
5     x(2) = 10;
6     out(i) = sum(x);
7 end
```

In this case, it is not clear what is the role of the variable `x`. It is acting like a sliced variable, but at the same time as a temporary variable. The code above will not run and Matlab will display an error. We can fix it, by making `x` a temporary variable:

```

1 out = zeros(10, 1);
2 parfor i = 1:10
3     x = [5 10];
4     out(i) = sum(x);
5 end

```

Now, x is being created inside the loop and it is clear that different iterations of the loop do not depend on one another.

Third, parfor-loops cannot be nested:

```

1 % does not work
2 parfor i = 1:5
3     sprintf('i=%d', i)
4     parfor j = 1:3
5         disp(j)
6     end
7 end
8 % works
9 parfor i = 1:5
10    sprintf('i=%d', i)
11    for j = 1:3
12        disp(j)
13    end
14 end

```

Nesting parfor does not work because we are already using all the workers (cores) in the outermost parfor, so the loop inside should be treated as a regular loop. For a complete set of issues that may arise when using parfor, refer to the reference page for variables in parfor-loops.

4.3 Memory Management

The parfor command can lead to memory issues. Each worker gets a copy of the variables in the workspace, so if you are using 2 GB of memory in your workspace and have a parfor-loop with 4 workers, then Matlab will require access to 8 GB of memory. Therefore, if you are working with big matrices, or your computer has a big number of cores, then you need to be mindful of the memory being used.

5 Assignment

Problem 1 Consider a deterministic growth model, where an agent decides between consumption (c_t) and investment in capital (k_t), while maximizing his utility. We can write this problem as:

$$\begin{aligned}
 & \max \sum_{t=0}^{\infty} \beta^t U(c_t) \\
 & \text{subject to } \begin{cases} k_{t+1} = k_t^\alpha - c_t + (1 - \delta)k_t, \forall t \geq 0 \\ k_0 > 0 \end{cases}
 \end{aligned}$$

Write the problem as a Bellman equation. Let $U(c; \sigma) = \frac{c^{1-\sigma}-1}{1-\sigma}$. Obtain the Euler equation for this problem in terms of the consumption c . Solve the problem by Value Function Iteration. Consider $\sigma = 2$, $\beta = 0.95$, $\delta = 0.1$ and $\alpha = 0.33$. Use the steady state value of k to create a grid for the possible values of k , say 100 points between $0.25k^*$ and $1.75k^*$. Start with a guess for V over the grid, for example $V(k) = 0$ for all k in the grid. Use the Matlab minimization function to solve for k . You may want to add the constraint that c should always be positive.

Use **parfor** to speed up the solution of this problem. What can be easily run in parallel in this problem? Is this problem embarrassingly parallel? If not, are there subsets of the problem that are embarrassingly parallel?

Graph the time it takes to find the solution of the problem when the number of grid points increases. Compare the time using regular loops and using **parfor** loops.