

Analyzing Parallel Programming Models for Magnetic Resonance Imaging

Forest Danford¹(✉), Eric Welch¹, Julio Cárdenas-Rodríguez²,
and Michelle Mills Strout¹

¹ Department of Computer Science, University of Arizona, Tucson, USA
{fdanford,welche,mstrout}@email.arizona.edu

² Department of Medical Imaging, University of Arizona, Tucson, USA
cardenaj@email.arizona.edu

Abstract. The last several decades have been marked by dramatic increases in the use of diagnostic medical imaging and improvements in the modalities themselves. As such, more data is being generated at an ever increasing rate. However, in the case of Magnetic Resonance Imaging (MRI) analysis and reports remain semi-quantitative, despite reported advantages of quantitative analysis (QA), due to prohibitive execution times. We present a collaborator’s QA algorithm for Dynamic Contrast-Enhanced (DCE) MRI data written in MATLAB as a case study for exploring parallel programming in MATLAB and Julia. Parallelization resulted in a 7.66x speedup in MATLAB and a 72x speedup in Julia. To the best of our knowledge, this comparison of Julia’s performance in a parallel, application-level program is novel. On the basis of these results and our experiences while programming in each language, our collaborator now prototypes in MATLAB and then ports to Julia when performance is critical.

Keywords: MATLAB · Julia · Parallel programming languages · Parallel applications · Medical imaging · Dynamic Contrast-Enhanced MRI

1 Introduction

Over the last 30 years there has been a dramatic increase in the usage of Magnetic Resonance Imaging (MRI) and other imaging modalities in the research and medical communities [26,27]. Simultaneously, there have been substantial improvements to the underlying technologies themselves that allow for images to be captured at significantly higher spatial and temporal resolutions. One of the consequences of these improvements has been a massive increase in the amount of data being generated [17]. For scientists and clinicians, this has often translated to making the difficult decision to forgo a truly quantitative analysis (QA) and instead sub-sample the available data to perform analyses and generate reports within an acceptable time frame. One application that currently suffers this

F. Danford and E. Welch—Authors contributed equally

plight is Quantitative Analysis (QA) model-based Dynamic Contrast-Enhanced MRI (DCE MRI) [21]. DCE MRI is used to visualize and characterize cancerous tumors in animal models and humans. The results can be used to predict patient-specific response to anticancer drugs and provide insight for new drug discovery [10,20]. DCE MRI is thus highly valuable to clinical oncologists and scientists. Our work in this study was in collaboration with a medical imaging researcher who had developed a serial implementation of DCE MRI analysis code written in MATLAB that had an unacceptable run time unless it used dramatically sub-sampled datasets ($\geq 50x$). For example, evaluating the QA algorithm on a breast cancer dataset [11] was estimated to take approximately six months, which is not practical.

To understand the volume of data being generated, it is valuable to briefly summarize the workflow of a DCE MRI experiment. Two separate regions, the volume containing the tumor and a volume containing normal vascular tissue (known as the reference region; in our case the leg), must be imaged continuously for several minutes. The MRI machine generates a 2-dimensional double array of intensities (at a specified spatial resolution) for every time point captured during the imaging time frame. These volumes must be imaged twice - first as a baseline to determine the tissue's innate relaxation properties (T_R volume), and then many times after the injection of a gadolinium-based contrast agent (CA). The CA's properties cause a change in the relaxation properties of the tissue, translating to different voxel intensities. These changes in signal intensity over time are used to calculate the CA's concentration, which is run through to a non-linear least squares computation to estimate the permeability of the tumor [20]. For more information, we refer readers to [5].

The computationally expensive portion of the QA DCE MRI algorithm lies in the non-linear fitting step, which must be performed for each voxel containing tissue in the image volume. Even in our heavily sub-sampled dataset, thousands of non-linear least square problems must be solved. Analysis of the original serial MATLAB code revealed that the majority of the workload could be computed in an embarrassingly parallel fashion, as each spatial point in the volume is independent of the others.

Given that the original code was already implemented in MATLAB and the popularity of the language in the biomedical engineering community, the straightforward approach to solving the performance problem was to implement an equivalent parallel implementation using the *Parallel Computing Toolbox* (PCT) in MATLAB. However, this only improved performance on a by 7.66x using 12 cores at best. Therefore, we decided to do a case study where we implemented a version in a newer language called Julia [4]. With Julia, we were able to achieve an approximately 72x performance improvement over the serial MATLAB implementation.

This paper makes the following contributions:

- We analyze the serial algorithm and identify areas where parallelization is expected to increase performance.

- We provide a more substantial benchmark for the Julia programming language (previous publications only include microbenchmark results) via quantitative performance data for the MATLAB and Julia versions.
- We provide a comparison of the process of writing parallel code in MATLAB and Julia based on our experiences in this case study.

Based on the significant increase in performance observed in the Julia implementations (22x speedup for the serial version and 72x speedup in the parallel version), our medical imaging researcher’s group has begun to translate their performance critical code to Julia.

The remainder of the paper is organized as follows: Sect. 2 contains a thorough description of the serial code and the experimental methodology. Sections 3 and 4 discuss the MATLAB and Julia languages and implementations. Section 5 details our results, Sect. 6 presents related work, and Sect. 7 concludes.

2 Analysis of the Serial MATLAB Code

The QA DCE MRI code consists of the following five tasks:

1. Loading the data
2. Segmenting the tissue from the background
3. Calculating T_1 time (seconds) on a per-voxel basis via non-linear least squares curve fitting to its phenomenological equation
4. Calculating the change in R_1 as a function of time
5. Estimating the relative permeability between the tumor and muscle tissue ($R^{K_{trans}}$) using the linear reference region model and evaluating the goodness-of-fit via the R^2 value.

The original serial algorithm was revised to remove all unnecessary operations (e.g. those done solely to produce human-interpretable images) and the standards recommended by The MathWorks were enforced [31]. This resulted in about a 1.67x speedup (execution time of 46.33 ± 1.33 seconds vs. 77.36 ± 4.22 seconds). This simplified version, whose execution time is profiled in Fig. 1, is used as the baseline for all comparisons presented in the paper. The remainder of this section analyzes the performance of the serial MATLAB implementation and describes each task in more detail.

2.1 Experimental Methodology

The execution times reported are the average of five runs \pm their standard deviation measured after a warm up run whose time was not included. This allows just-in-time (JIT) compilation to complete and parallel workers to initialize in both languages to avoid comparing compiler speeds and efficiencies, and instead provide a comparison of the performance of the computation in each language. Given that these are one-off costs that are amortized over many datasets, the

contribution to the overall run time becomes negligible. All speedups reported are comparisons of the T_1 calculation (see Sect. 2.2).

Execution times were obtained by submitting jobs to a PBS scheduling system on a high-performance computing cluster consisting of Lenovo NeXtScale nx360 M5 compute nodes with Xeon Haswell E5-2695 V3 Dual 14-core processors operating at 2.3 GHz. The dataset used for benchmarking the various implementations consisted of 5 time points, each composed of a 128×128 array of doubles. MATLAB 2015b and Julia v.0.4.3 were used for these experiments.

Timing information in MATLAB was obtained using the built-in `tic/toc` construct. These values were validated against MATLAB’s built-in code profiler and were not found to be significantly different.

Step	Time (seconds)	
	Tumor	Leg
Load data	0.122 ± 0.003	
Segmentation	0.028 ± 0.019 sec	0.014 ± 0.003
T_1 calculation	34.055 ± 1.231	10.932 ± 0.211
ΔR_1 calculation	0.018 ± 0.001	0.007 ± 0.000
K^{trans} via RRM	1.151 ± 0.021	
Total run time	46.329 ± 1.335	

Fig. 1. Execution time for main tasks of serial MATLAB implementation.

2.2 Performance Bottleneck: T_1 Calculation

As shown in Fig. 1, the T_1 calculation accounts for approximately 97% of the total run time on average. The disparity between the run time of the tumor and leg volume arises due to the difference in the number of tissue voxels present post-segmentation (3199 voxels vs. 1083 voxels). If we account for this difference, it appears that the execution time for this function scales linearly, so there is nothing intrinsically different about the tumor. In this task, each voxel from the T_R volume is fit to the following phenomenological equation to solve for T_1 :

$$S(t) = M_z(1 - e^{-T_R/T_1(t)}) \quad (1)$$

via MATLAB’s built-in non-linear least squares solver, `lsqcurvefit`. As the calculations performed for each voxel is independent and the two T_R volumes are independent, they can be computed in an embarrassingly parallel fashion. Parallelization efforts were focused on this function for the MATLAB and Julia implementations and as such speedup results consider this section only.

2.3 Other Tasks

The raw output from the MRI machine is pre-processed into an array of doubles that are stored in a MATLAB `.mat` file and a `csv` format for the Julia implementation. MATLAB’s built-in command `load()` takes virtually no time (Fig. 1), and the same is true of Julia’s `readdlm` function.

Segmentation is performed using MATLAB’s built-in k -means clustering algorithm (`kmeans`). A 1-D vector containing all voxels is provided as input and the function categorizes each voxel as either tissue or background. MATLAB’s built-in implementation of Otsus method to determine the gray-scale threshold for binarization was also explored, but k -means generated the best mask (results omitted). The k -means algorithm in the clustering package of JuliaStats generated identical masks. The run time of this task is also negligible.

After T_1 is calculated, two equations are combined to describe R_1 (\propto CA concentration) as a function of time, T_1 , and T_R :

$$R_1(t) = -\frac{1}{T_R} \cdot \ln(1 - S(t)e^{-T_R \cdot R_1(0)/S(0)}) \quad (2)$$

As this is a straightforward calculation that takes a negligible amount of time, we do not parallelize it. At this point, concentration as a function of time is known — that is, we have solved the non-linear relationship between voxel intensity and CA concentration for both tissue volumes. The reference region model (RRM) [5] is used to calculate the permeability of the CA, K^{trans} , and the R^2 value on a voxel-by-voxel basis for the tumor tissue via a linear least squares fitting with non-negativity constraints (`lsqnonneg`). Despite each voxel being independent, the function runs for a short enough time that overhead costs would likely outweigh parallelization gains.

2.4 Performance Analysis Summary

To summarize, the tumor and reference region voxels are entirely independent and are processed identically until the calculation of K^{trans} . Additionally, in all tasks that process the segmented tissue voxels, the calculations performed on each voxel are independent, and can therefore be computed in an embarrassingly parallel fashion. Analysis of the computations that occur within these tasks does not reveal anything inherent to the code that would result in one language handling the computation preferentially.

3 MATLAB Implementations

The DCE MRI algorithm was parallelized in MATLAB by a graduate student who had several years of experience working with MATLAB and the *Parallel Computing Toolbox* (PCT) in the biomedical domain. Using the PCT, a 7.66x speedup over the serial code was achieved for the T_1 calculation.

3.1 MATLAB Background

Because most scientists do not receive formal software engineering training [2], MATLAB’s ease of use (e.g. weak and dynamic typing system, lack of need to declare dimensions, etc.) and trusted libraries have made it a popular language for scientific computing applications [9, 23]. Additionally, the core intentions of MATLAB’s parallel programming model (the PCT) were to extend the

forementioned traditional strengths of MATLAB onto the cluster via first-class language constructs to deal with embarrassingly parallel problems [24]. This allows users to easily utilize multicore processors, GPUs and clusters with minimal modification of code or impact on readability.

However, MATLAB is not as performant as other programming languages, especially those used for parallel programming [4]. MATLAB worker threads that execute concurrent computation are heavyweight, and the PCT is proprietary and has limited scalability [15]. Additionally, the dynamic and complex typing system results in significant overhead [1, 8, 14]. Despite these drawbacks, in the realm of scientific computing, time to solution, readability, portability, and maintainability often trump pure performance [2], so MATLAB is utilized significantly [14].

As a bridge, continuous and significant work has been done on static analysis, ahead of time speculation, JIT compilation, and automatically porting existing code to more performant languages [1, 9, 13, 15, 16]. While automatically ported code makes sacrifices in terms of both the highly human-interpretable syntax and interactive nature of MATLAB, The MathWorks has adopted some of the other techniques, and as of release 2015b MATLAB is now entirely JIT compiled [29, 30]. Based on selected case studies [30] and comparing the benchmarks performed by the Julia language creators on MATLAB 2011a [4] to MATLAB 2015b (<http://julialang.org/>), it appears that this has had a predominantly positive impact on performance.

3.2 Parallelization Using the Parallel Computing Toolbox

For the DCE MRI application, the T_1 calculation was rewritten so that every voxel was fit in parallel, while the tissue was processed serially (version 1). Modifications to the serial code were minimal and consisted of slicing variables [18] and changing the `for` to a `parfor` in the function shown in Fig. 2. As nested parallelism is not supported in MATLAB, version 2 consisted of restructuring the code to process both the tissues and voxels in parallel.

Table 1 contains the execution time, speedup, and efficiency for the parallelized versions as a function of the number of cores provided to the MATLAB

Table 1. Summary of parallel MATLAB implementations

# cores	Version					
	1			2		
	Runtime (seconds)	Speedup	Efficiency	Runtime (seconds)	Speedup	Efficiency
1	49.58 ± 1.90	0.93x	93.17%	50.09 ± 2.42	0.92x	92.21%
2	26.96 ± 0.18	1.71x	85.66%	25.26 ± 0.50	1.83x	91.44%
4	14.57 ± 0.14	3.17x	79.25%	13.86 ± 0.55	3.33x	83.30%
8	8.37 ± 0.20	5.52x	69.01%	7.95 ± 0.19	5.81x	72.64%
12	6.03 ± 0.14	7.66x	63.85%	6.17 ± 0.10	7.49x	62.43%

parallel pool for the T_1 calculation. Versions 1 and 2 using 12 cores were the most performant and achieved a 7.66x speedup at 63.85% efficiency and a 7.49x speedup at 62.43% efficiency respectively. The execution time of versions 1 and 2 were not found to be statistically different as determined by a Two-Sample Kolmogorov-Smirnov test (`kstest2`) at a p -value of .05. On average, it took 10.27 ± 1.48 seconds to initialize the parallel pool, which is nearly 45% of the total execution time (for all sections) of the fastest version.

4 Julia Implementations

The DCE MRI algorithm was converted from MATLAB to Julia by a graduate student who had never worked with MATLAB or Julia, had no experience with medical imaging, and had no prior experience with parallel computing. Thanks to several features of the two languages, this process was straightforward, as was the parallelization of the T_1 calculation, which ultimately resulted in a 72x speedup over the serial MATLAB code. Julia v0.4.3 was used and timing information was obtained using the `@time` macro as per the recommended best practices.

4.1 Julia Background

Julia was designed specifically for numerical and scientific computing, and has a syntax similar to languages such as MATLAB and R, but has been shown to outperform such dynamically-typed languages on microbenchmarks, often achieving performance comparable to C and Fortran [4]. Julia’s creators have indicated that the language’s speed is accounted for by its robust type inference system, multiple dispatch, and high-performance LLVM-based JIT compiler that generates optimized, on-the-fly native machine code directly [3].

Julia features built-in parallel capability [3], and although it is still in pre-release, its user base has created a significant number of native libraries for the language [25], as well as interfaces to commonly utilized libraries from other languages such as NLopt.

4.2 Serial Julia Implementation

The run time of the T_1 calculation in the serial Julia implementation was 2.07 ± 0.01 seconds, roughly a 22x speedup over the serial MATLAB version. Because MATLAB and Julia were designed for programmability and are thoroughly documented, it was easy to investigate MATLAB functions and determine how to implement them in Julia. Indeed, Julia seems to liberally “borrow” features from many other languages, including MATLAB, and we found that many MATLAB functions, such as the matrix manipulation function `reshape`, have been implemented in Julia with nearly identical syntax and semantics (Fig. 5).

When required MATLAB functions had no equivalent in the Julia standard library, it was easy to locate third-party, open-source libraries providing the needed functionality. This occurred in the segmentation task — Julia did not

MATLAB

```

% Normalize signal
Signal = masked_T1_vTR ./
    repmat(max(masked_T1_vTR, [], 2), 1,
        numTimePoints);

% Set curve fitting parameters
x0 = [1.2, 3.0];
lb = [1, 1];
ub = [2, 5];

% Define model function
T1vTRfunc = @(pars, xdata) pars(1) .*
    (1 - exp(-xdata./pars(2)));

% Do curve fitting on per voxel basis
for q = 1:numSignalPixels
    [beta, ~, resid, ~, ~, ~, J] =
        lsqcurvefit(T1vTRfunc, x0, TR,
            Signal(q, :)', lb, ub, options);
    T1vTR.Map(q) = beta(2);
end

```

Julia

```

for q in 1:size(indices,1)
    Signal = T1vTR_Images[indices[q,:]]
    Signal ./= maximum(Signal)

    # Create solver and set objective function
    model = Model(solver=NLOptSolver(
        algorithm=:LN_COBYLA))
    @defVar(model, beta[1:2])
    @setNLObjective(model, Min, sum{(Signal[i]-
        beta[1]*(1-exp(-TR[i]/beta[2])))^2, i=1:N})

    # add constraints and initial guess
    @addNLConstraint(model, beta[1] >= lb[1])
    @addNLConstraint(model, beta[1] <= ub[1])
    @addNLConstraint(model, beta[2] >= lb[2])
    @addNLConstraint(model, beta[2] <= ub[2])
    setValue(beta[1], x0[1])
    setValue(beta[2], x0[2])

    status = solve(model)
    T1vTR_Map[indices[q]] = getValue(beta[2])
end

```

Fig. 2. Curve-fitting function in MATLAB and Julia

have a built-in k -means function. However, because Julia features a package manager, typing the command `Pkg.add("Clustering.jl")` at the Julia command prompt immediately downloaded the latest version of the library `Clustering.jl` from a github repository (<https://github.com/JuliaStats/Clustering.jl>), and provided the use of its `kmeans` function, which had the same syntax as MATLAB's implementation. For the task involving least-squares curve fitting, a function from the JuliaOpt package `LsqFit.jl` that appeared to mirror MATLAB's `lsqcurvefit` was inadequate, as it did not allow for bounds on the solution. An interface allowing Julia to call the NLOpt library, `NLOpt.jl`, was used instead (Fig. 2).

A few MATLAB functions involved in the permeability estimation step, such as `cumtrapz` and `nans`, were not present in Julia, but it was straightforward to simply implement them.

4.3 Parallelization of Julia Implementation

Julia allows for a variety of approaches to parallel programming, but for the embarrassingly parallel computations required in this application, a simple shared memory model was sufficient. We utilized a special Julia datatype designed for this purpose, called a `SharedArray`, which is accessible by multiple processors. The demonstration code for `SharedArrays` in the Julia documentation provided a clear blueprint for our implementation. Following this example, the key steps were to copy data from several arrays into `SharedArrays`, and create two kernels that: (1) determined which voxels to process on each processor, and (2) ran the appropriate subset of the least-squares curve-fitting calculations on each processor.

Index determination kernel

```

@everywhere function my_range(shared_indices)
    idx = indexpids(shared_indices)
    if idx == 0
        return 0:1
    end
    n_chunks = length(procs(shared_indices))
    splits = [round(Int, s) for s in linspace(0, size(shared_indices, 1), n_chunks+1)]
    return splits[idx]+1 : splits[idx+1]
end

```

(a) Kernel to divide indices in SharedArray to be processed on different processors

Serial curve-fitting in Julia

```

for q in 1:size(indices,1)
    Signal = T1vTR_Images[indices[q],:]
    Signal ./= maximum(Signal)
    T1vTR_Map[indices[q]] =
        calculate_T1(Signal, TR)
end

```

Parallel curve-fitting kernel run on each worker

```

for q in my_range
    Signal = T1vTR_Images[indices[q],:]
    Signal ./= maximum(Signal)
    T1vTR_Map[indices[q]] =
        calculate_T1(Signal, TR)
end

```

(b) Serial vs. Parallel curve-fitting function in Julia

Fig. 3. Functions involved in serial and parallel curve fitting in Julia

The first of these (Fig. 3a), which was essentially copied from the Julia documentation [32], partitions a collection of indices into equal-sized groups for each processor. The second (Fig. 3b) merely substitutes an iteration over all voxel indices for an iteration over the indices assigned to the worker running the kernel.

As shown in Table 2, the parallelization produces gains of up to 3.24x over the serial Julia implementation, with reasonable efficiency.

Table 2. Summary of parallel Julia implementations

# Cores	Execution time (seconds)	Speedup relative to serial Julia	Total efficiency
1	2.63 ± 0.37	0.79x	78.56%
2	1.47 ± 0.07	1.41x	70.53%
4	1.03 ± 0.09	2.01x	50.37%
8	0.71 ± 0.01	2.92x	36.51%
12	0.64 ± 0.06	3.24x	27.00%

5 Results

In this section, we evaluate our implementations of the researcher’s algorithm in terms of performance and reliability, and report on differences in programmability between MATLAB and Julia. We find that rewriting the algorithm in Julia resulted in performance gains exceeding those of using the MATLAB PCT.

5.1 Performance

Quantitative performance metrics were generated for the parallel MATLAB, serial Julia, and parallel Julia implementations for the T_1 calculation. The most performant version of parallel MATLAB achieved a speedup of 7.66x with 63.85% efficiency. As shown in Fig. 4, the serial version of Julia achieved a 22.32x speedup compared to the serial MATLAB version, a 72.32x speedup with parallelism, and a 9.44x speedup when comparing the fastest parallel Julia implementation with the fastest parallel MATLAB implementation. However, MATLAB exhibits substantially more efficient parallelism compared to Julia (Tables 1 and 2).

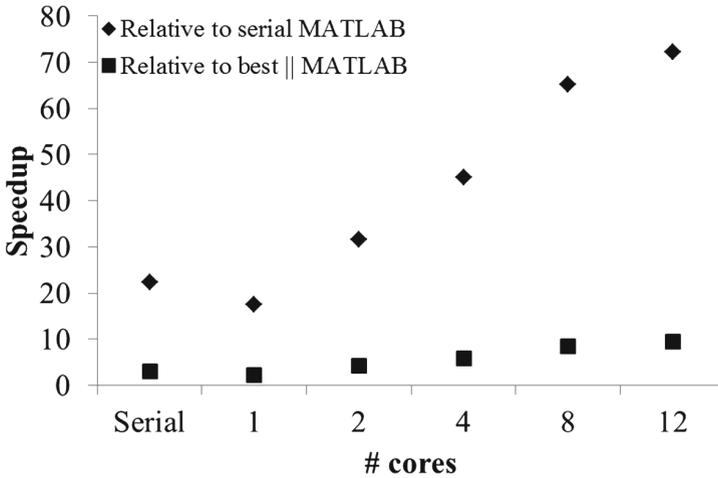


Fig. 4. Speedup of Julia implementations relative to MATLAB implementations

5.2 Reliability

The results obtained with Julia were comparable to those obtained with MATLAB. The T_1 calculation results obtained from Julia had insignificant numerical differences compared to the original MATLAB implementation: within the 1,083 signal pixels in the leg, the root mean square error (RMSE) was $7.90e-4$ and the maximum absolute difference (MAD) was 0.0108; within the 3,199 signal pixels of the tumor, the RMSE was $9.07e-4$ and the MAD was 0.0347. For the final $R^{K_{trans}}$ calculation, the RMSE was $2.31e-3$ and the MAD was 0.0744.

5.3 Programmability

MATLAB and Julia are similar in terms of syntax (Fig. 5) and expressiveness. The serial MATLAB and Julia implementations contain 292 and 244 source lines of code (SLOC) respectively. The fastest parallel implementation in MATLAB contains 239 SLOC compared to Julia's 284 SLOC. Productivity

was also comparable after familiarity was gained with Julia — in both languages it took approximately 1.5 h to parallelize the algorithm. This was somewhat surprising given the differences in familiarity with the languages and parallelism. The graduate student who learned MATLAB and Julia in tandem did not find either language particularly more difficult to understand than the other — however, he was simply translating MATLAB code to Julia code, as opposed to prototyping purely in Julia.

Mask extraction code in MATLAB

```
[x,y,z] = size(data);
data_mtx = reshape(data, x*y, z);
k = 2
kmeans_group_index = kmeans(data_mtx, k);
group_index_noise = kmeans_group_index(1);
mask = reshape(kmeans_group_index ~
group_index_noise, x, y);
```

Mask extraction code in Julia

```
(x,y,z) = size(data)
data_mtx = reshape(data, x*y, z)
k = 2
kmeans_group_index = kmeans(data_mtx', 2)
group_index_noise = kmeans_group_index.assignments[1]
mask = reshape([d == group_index_noise ? 0 : 1
for d in kmeans_group_index.assignments], x, y)
```

Fig. 5. k -means mask extraction code in MATLAB and Julia

One difference between the languages is that Julia’s error messages initially appear cryptic, often including information about types that the programmer did not specify, but were inferred during JIT compilation. In Fig. 6, an error arises because `size(out)` returns a `Tuple` type, while the colon operator expects the right-hand operand to be an integer. In MATLAB, comparable code runs with no error, since the colon simply uses the first element of the 2×1 array returned by `size`. MATLAB was generally found to be more “forgiving” than Julia.

```
julia> function seq(a, b)
    out = zeros(b-a)
    for i in 1:size(out)
        out[i] = a+i
    end
    out
end
seq (generic function with 1 method)

julia> seq(2,4)
ERROR: MethodError: `colon` has no method matching colon(::Int64, ::Tuple{Int64})
Closest candidates are:
  colon{T<:Real}(::T<:Real, ::Any, ::T<:Real)
  colon{A<:Real,C<:Real}(::A<:Real, ::Any, ::C<:Real)
  colon{T}(::T, ::Any, ::T)
  ...
in seq at none:3
```

Fig. 6. A typical Julia error message

Moreover, MATLAB’s Code Analyzer will infer and warn the programmer if, for instance, there is code that would distribute a large array to many parallel workers, while Julia does not provide warnings of this type. More broadly, the

MATLAB IDE was felt to be more convenient and powerful than using a Julia the Juno IDE (<http://junolab.org/>) or a Jupyter Notebook (<http://jupyter.org/>).

Another difficulty that arose when implementing the algorithm in Julia was locating a library that would provide the same functionality as MATLAB's `lsqcurvefit` function. The fact that a Julia library function with the same name had been written, but which used a different algorithm that did not allow bounds constraints, was troubling.

In spite of these hurdles, our experience suggests that learning Julia is comparable to learning a dynamically-typed programming language.

5.4 Limitations

Our implementations in the two languages used different algorithms for performing the fitting for the T_1 calculation: the MATLAB `lsqcurvefit` function used the Trust-Region-Reflective Least Squares Algorithm from the Optimization Toolbox, while the Julia version used a modified version of Powell's implementation of the COBYLA algorithm in the NLOpt library. In order to ensure that the execution times were not different due to this distinction, the original MATLAB implementation was altered to solve the non-linear fitting using the NLOpt library with the same parameters and objective function that the Julia implementation did. The T_1 calculation took 80.53 ± 0.41 seconds using NLOpt (compared to 52.40 ± 0.65 seconds using `lsqcurvefit`). The numerical differences between the two implementations were insignificant: The voxels in the leg had an RMSE of $3.63\text{e-}04$ and MAD of $4.25\text{e-}03$, and the tumor had an RMSE of $3.75\text{e-}04$ and MAD of $1.06\text{e-}02$. Since MATLAB is proprietary software, it is difficult to determine the source of the performance differences observed in a detailed fashion.

6 Related Work

To the best of our knowledge this is the first comparison of the performance of two implementations of the linear reference region model (LRRM) for the analysis of DCE MRI data. Smith et al. [25] recently described DCEMRI.jl, a Julia implementation of the commonly used Tofts model for DCE MRI, and compared it briefly against implementations in IDL [22] and R language [28]. DCEMRI.jl was reported to be 24X faster than DCE@urLAB; both implementations used the Levenberg—Marquardt algorithm (LMA). Additionally, DCEMRI.jl was reported to be 10X faster than dcmriS4, but this is not a straight comparison because dcmriS4 uses a Bayesian hierarchical approach for curve fitting that is more demanding than the LMA.

Much research has been done on compiling and automatically parallelizing MATLAB [6, 7, 12, 19]. We did not compare the Julia and MATLAB parallel implementations against what the most active MATLAB compiler project, McLab at McGill, can perform and this is future work. One important consideration in the selection of Julia was that it is a programming environment

that is gaining significant community support. For the medical imaging research community to switch to a new programming platform, the platform will need to show signs of significant community support and longevity.

A review of the literature found no references comparing Julia’s performance to that of other languages for the same algorithm in an authentic parallel programming application (i.e. beyond microbenchmarks).

7 Conclusion

Although the initial MATLAB code processed the benchmark data set on the order of minutes, a typical experiment for a medical imaging researcher often includes more and larger images, taken from dozens of patients, and may take weeks to run. The 72x speedup achieved using Julia would reduce weeks to hours, removing a significant constraint on researchers. Based on our experience while writing this paper, Julia appears to be a very attractive, emergent programming language for scientific computing.

As a result of our work on the DCE MRI algorithm, and his subsequent investigations into Julia, the research scientist has adopted the following model: (1) prototype and validate in MATLAB, (2) use MATLAB to identify bottlenecks, (3) port performance-critical portions of code to Julia. Taking into account the results of this case study, we feel that it would not be difficult for other research groups (who have already navigated the MATLAB learning curve) to similarly utilize Julia. Furthermore, since none of the performance gains found in this study resulted directly from the use of MRI data, it seems likely that other scientific computing applications could be prototyped rapidly in MATLAB, then efficiently ported to Julia for high-throughput applications. The DCE MRI code in both languages can be found at github.com/fdanford/LCPC2016_MATLAB_Julia/.

Acknowledgments. An allocation of computer time from the UA Research Computing High Performance Computing (HPC) and High Throughput Computing (HTC) at the University of Arizona is gratefully acknowledged by the authors.

References

1. Almási, G., Padua, D.: MaJIC: compiling matlab for speed and responsiveness. In: ACM SIGPLAN Notices, vol. 37, pp. 294–303. ACM (2002)
2. Basili, V.R., Carver, J.C., Cruzes, D., Hochstein, L.M., Hollingsworth, J.K., Shull, F., Zekowitz, M.V.: Understanding the high-performance-computing community: a software engineer’s perspective. *IEEE Softw.* **25**(4), 29 (2008)
3. Bezanson, J., Edelman, A., Karpinski, S., Shah, V.B.: Julia: a fresh approach to numerical computing. arXiv preprint [arXiv:1411.1607](https://arxiv.org/abs/1411.1607) (2014)
4. Bezanson, J., Karpinski, S., Shah, V.B., Edelman, A.: Julia: a fast dynamic language for technical computing. arXiv preprint [arXiv:1209.5145](https://arxiv.org/abs/1209.5145) (2012)

5. Cárdenas-Rodríguez, J., Li, X., Whisenant, J.G., Barnes, S., Stollberger, R., Gore, J.C., Yankeelov, T.E.: The basic principles of dynamic contrast-enhanced magnetic resonance imaging. In: Bammer, R. (ed.) *MR & CT Perfusion Imaging: Clinical Applications and Theoretical Principles*, chapter 31. Lippincott Williams & Wilkins, Philadelphia (2016)
6. Casey, A., Li, J., Doherty, J., Chevalier-Boisvert, M., Aslam, T., Dubrau, A., Lameed, N., Aslam, A., Garg, R., Radpour, S., Belanger, O.S., Hendren, L., Verbrugge, C.: McLab: an extensible compiler toolkit for MATLAB and related languages. In: *Proceedings of the Third C* Conference on Computer Science and Software Engineering, C3S2E 2010*, pp. 114–117. ACM, New York (2010)
7. Chauhan, A., Kennedy, K.: Optimizing strategies for telescoping languages: procedure strength reduction and procedure vectorization. In: *Proceedings of the 15th ACM International Conference on Supercomputing*, New York, pp. 92–102 (2001)
8. Chevalier-Boisvert, M., Hendren, L., Verbrugge, C.: Optimizing MATLAB through just-in-time specialization. In: *International Conference on Compiler Construction*, pp. 46–65. Springer, Heidelberg (2010)
9. De Rose, L., Padua, D.: Techniques for the translation of MATLAB programs into FORTRAN 90. *ACM Trans. Program. Lang. Syst. (TOPLAS)* **21**(2), 286–323 (1999)
10. DeGrandchamp, J.B., Whisenant, J.G., Arlinghaus, L.R., Abramson, V.G., Yankeelov, T.E., Cardenas-Rodrguez, J.: Predicting response before initiation of neoadjuvant chemotherapy in breast cancer using new methods for the analysis of dynamic contrast enhanced MRI (DCE MRI) data. In: *International Society for Optics and Photonics, SPIE Medical Imaging*, pp. 978811–978811, March 2016
11. DeGrandchamp, J.B., Whisenant, J.G., Arlinghaus, L.R., Abramson, V.G., Yankeelov, T.E., Cárdenas-Rodríguez, J.: Predicting response before initiation of neoadjuvant chemotherapy in breast cancer using new methods for the analysis of dynamic contrast enhanced MRI (DCE MRI) data. In: *Proceedings of SPIE*, pp. 9788:978811–978811-10 (2016)
12. DeRose, L., Gallivan, K., Gallopoulos, E., Marsolf, B., Padua, D.: A MATLAB compiler and restructurer for the development of scientific libraries and applications. In: *Preliminary Proceedings of the 8th International Workshop on Languages and Compilers for Parallel Computing*, pp. 18.1–18.18, May 1995
13. Doherty, J., Hendren, L., Radpour, S.: Kind analysis for MATLAB. *ACM SIGPLAN Not.* **46**(10), 99–118 (2011)
14. Dubrau, A.W., Hendren L.J.: *Taming MATLAB*, vol. 47. ACM (2012)
15. Kumar, V., Hendren, L.: *Compiling MATLAB for high performance computing via x10*. Sable Technical report 03 (2013)
16. Li, X., Hendren, L.: Mc2FOR: a tool for automatically translating MATLAB to FORTRAN 95. In: *2014 Software Evolution Week-IEEE Conference on Software Maintenance, Reengineering and Reverse Engineering (CSMR-WCRE)*, pp. 234–243. IEEE (2014)
17. Markonis, D., Schaer, R., Eggel, I., Müller, H., Depeursinge, A.: Using mapreduce for large-scale medical image analysis. arXiv preprint [arXiv:1510.06937](https://arxiv.org/abs/1510.06937) (2015)
18. MathWorks. Matlab parallel computing toolbox users guide (2016). www.mathworks.com/help/pdf_doc/distcomp/distcomp.pdf. Accessed 20 May 2016
19. Menon, V., Pingali, K.: A case for source-level transformations in MATLAB. In: *Proceedings of the 2nd Conference on Domain-Specific Languages*, pp. 53–66. USENIX Association, Berkeley, 3–5 1999

20. O'Connor, J.P., Jackson, A., Parker, G.J., Jayson, G.C.: DCE-MRI biomarkers in the clinical evaluation of antiangiogenic and vascular disrupting agents. *Br. J. Cancer* **96**(2), 189–195 (2007)
21. O'Connor, J.P.B., Tofts, P.S., Miles, K.A., Parkes, L.M., Thompson, G., Jackson, A.: Dynamic contrast-enhanced imaging techniques: CT and MRI. *Br. J. Radiol.* **84**(2) (2011)
22. Ortuño, J.E., Ledesma-Carbayo, M.J., Simões, R.V., Candiota, A.P., Arús, C., Santos, A.: Dce@urlab: a dynamic contrast-enhanced mri pharmacokinetic analysis tool for preclinical data. *BMC Bioinform.* **14**(1), 1 (2013)
23. Radpour, S., Hendren, L., Schäfer, M.: Refactoring MATLAB. In: Jhala, R., Bosschere, K. (eds.) CC 2013. LNCS, vol. 7791, pp. 224–243. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-37051-9_12](https://doi.org/10.1007/978-3-642-37051-9_12)
24. Sharma, G., Martin, J.: MATLAB: a language for parallel computing. *Int. J. Parallel Prog.* **37**(1), 3–36 (2009)
25. Smith, D.S., Li, X., Arlinghaus, L.R., Yankeelov, T.E., Welch, E.B.: DCEMRI.jl: a fast, validated, open source toolkit for dynamic contrast enhanced MRI analysis. *PeerJ* **3**, e909 (2015)
26. Smith-Bindman, R., Miglioretti, D.L., Johnson, E., Lee, C., Feigelson, H.S., Flynn, M., Greenlee, R.T., Kruger, R.L., Hornbrook, M.C., Roblin, D., Solberg, L.I., Vaneman, N., Weinmann, S., Williams, A.E.: Use of diagnostic imaging studies and associated radiation exposure for patients enrolled in large integrated health care systems, 1996–2010. *JAMA* **307**(22), 2400–2409 (2012)
27. Smith-Bindman, R., Miglioretti, D.L., Larson, E.B.: Rising use of diagnostic medical imaging in a large integrated health system. *Health Aff.* **27**(6), 1491–1502 (2008)
28. Whitcher, B., Schmid, V.J., et al.: Quantitative analysis of dynamic contrast-enhanced and diffusion-weighted magnetic resonance imaging for oncology in r. *J. Stat. Softw.* **44**(5), 1–29 (2011)
29. Matlab execution engine. <http://www.mathworks.com/products/matlab/matlab-execution-engine/>. Accessed 26 Aug 2016
30. Loren on the art of matlab: Run code faster with the new matlab execution engine. <http://blogs.mathworks.com/loren/2016/02/12/run-code-faster-with-the-new-matlab-execution-engine/>. Accessed 26 Aug 2016
31. Techniques to improve performance. http://www.mathworks.com/help/matlab/matlab_prog/techniques-for-improving-performance.html
32. Parallel computing - julia language 0.4.7 predocumentation. <http://docs.julialang.org/en/release-0.4/manual/parallel-computing/>. Accessed 11 July 2016