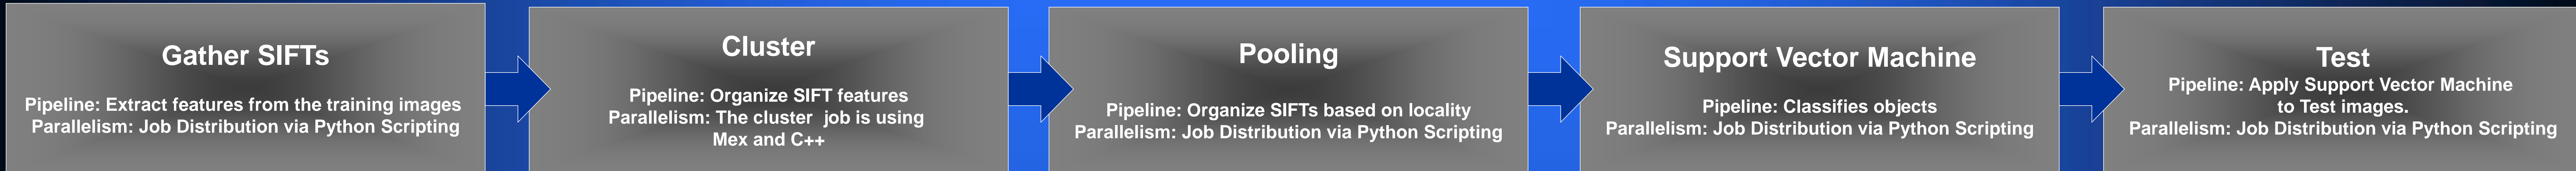




Exploiting Parallelism in Computer Vision

Prof. Joshua Stough, Paul Nguyen '13, Lee Davis '13, Department of Computer Science
Washington and Lee University



Motivation

To improve our object recognition algorithm's speed through parallel programming in order to implement more complicated features that will yield more accurate results.

Improvements

•Parallelism – splitting up work into parts to use multiple processors for faster completion



•Low level programming - "closer to the hardware" programming. Reduces abstraction from binary instructions the hardware executes, which reduces computational overhead. Allows significant code optimization.

Low Level Languages

Machine Code
10010010
10011010
01100110
00101101

```

    #include <omp.h>
    #include "mex.h"
    void mexFunction(int nlhs, mxArray *plhs[], int nrhs, const mxArray *prhs[]) {
        unsigned int i, j;
        D = mxGetM(prhs[0]);
        N = mxGetN(prhs[0]);
        K = mxGetN(prhs[1]);
        plhs[0] = mxCreateDoubleMatrix(K, N, mxREAL);
        double *restrict X = mxGetPr(prhs[0]), *restrict C = mxGetPr(prhs[1]), *restrict DM = mxGetPr(plhs[0]);
        #pragma omp parallel for
        for(j=0; j<N; j++){
            for(i=0; i<K; i++){
                DM[i+j*K]=0;
                for(i2=0; i2<D; i2++){
                    DM[i+j*K]+=(X[D*i2+i]-C[D*i2+i2])*(X[D*i2+i2]-C[D*i2+i2]);
                }
            }
        }
    }
  
```

High Level Languages

```

    function [D] = euclideanDistance(X,C)
    k = size(C,2);
    n = size(X,2);
    D = zeros(k, n);
    for i = 1:k
        D(i,:) = sum((X - repmat(C(:,i),[1 n])).^2);
    end
    end
  
```

•Principal Component Analysis (PCA) - statistical tool used for dimensionality reduction in SIFT features. Reduces the amount of work, which allows for speed increase.

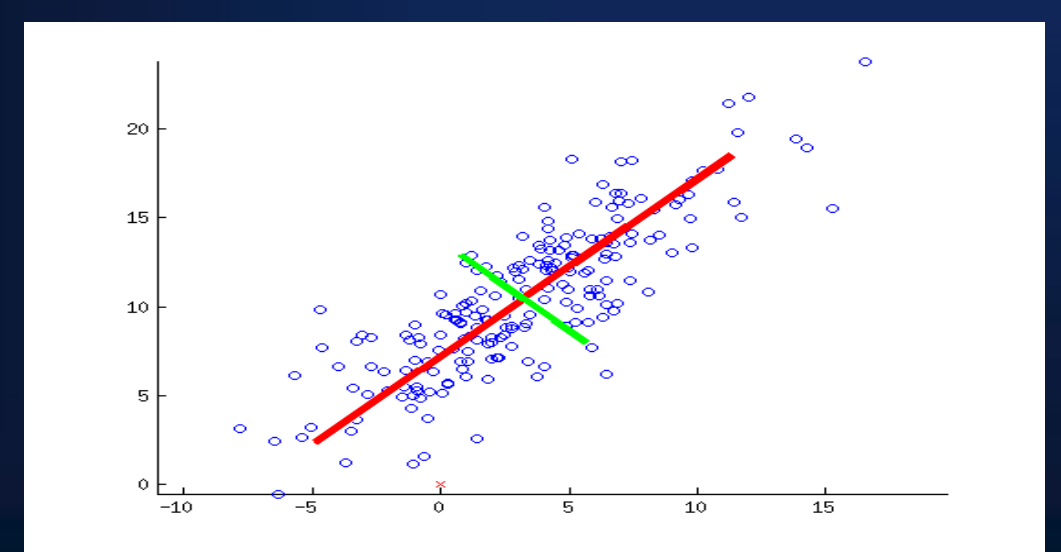


Figure 1 - No PCA. Red/Green shows data variance


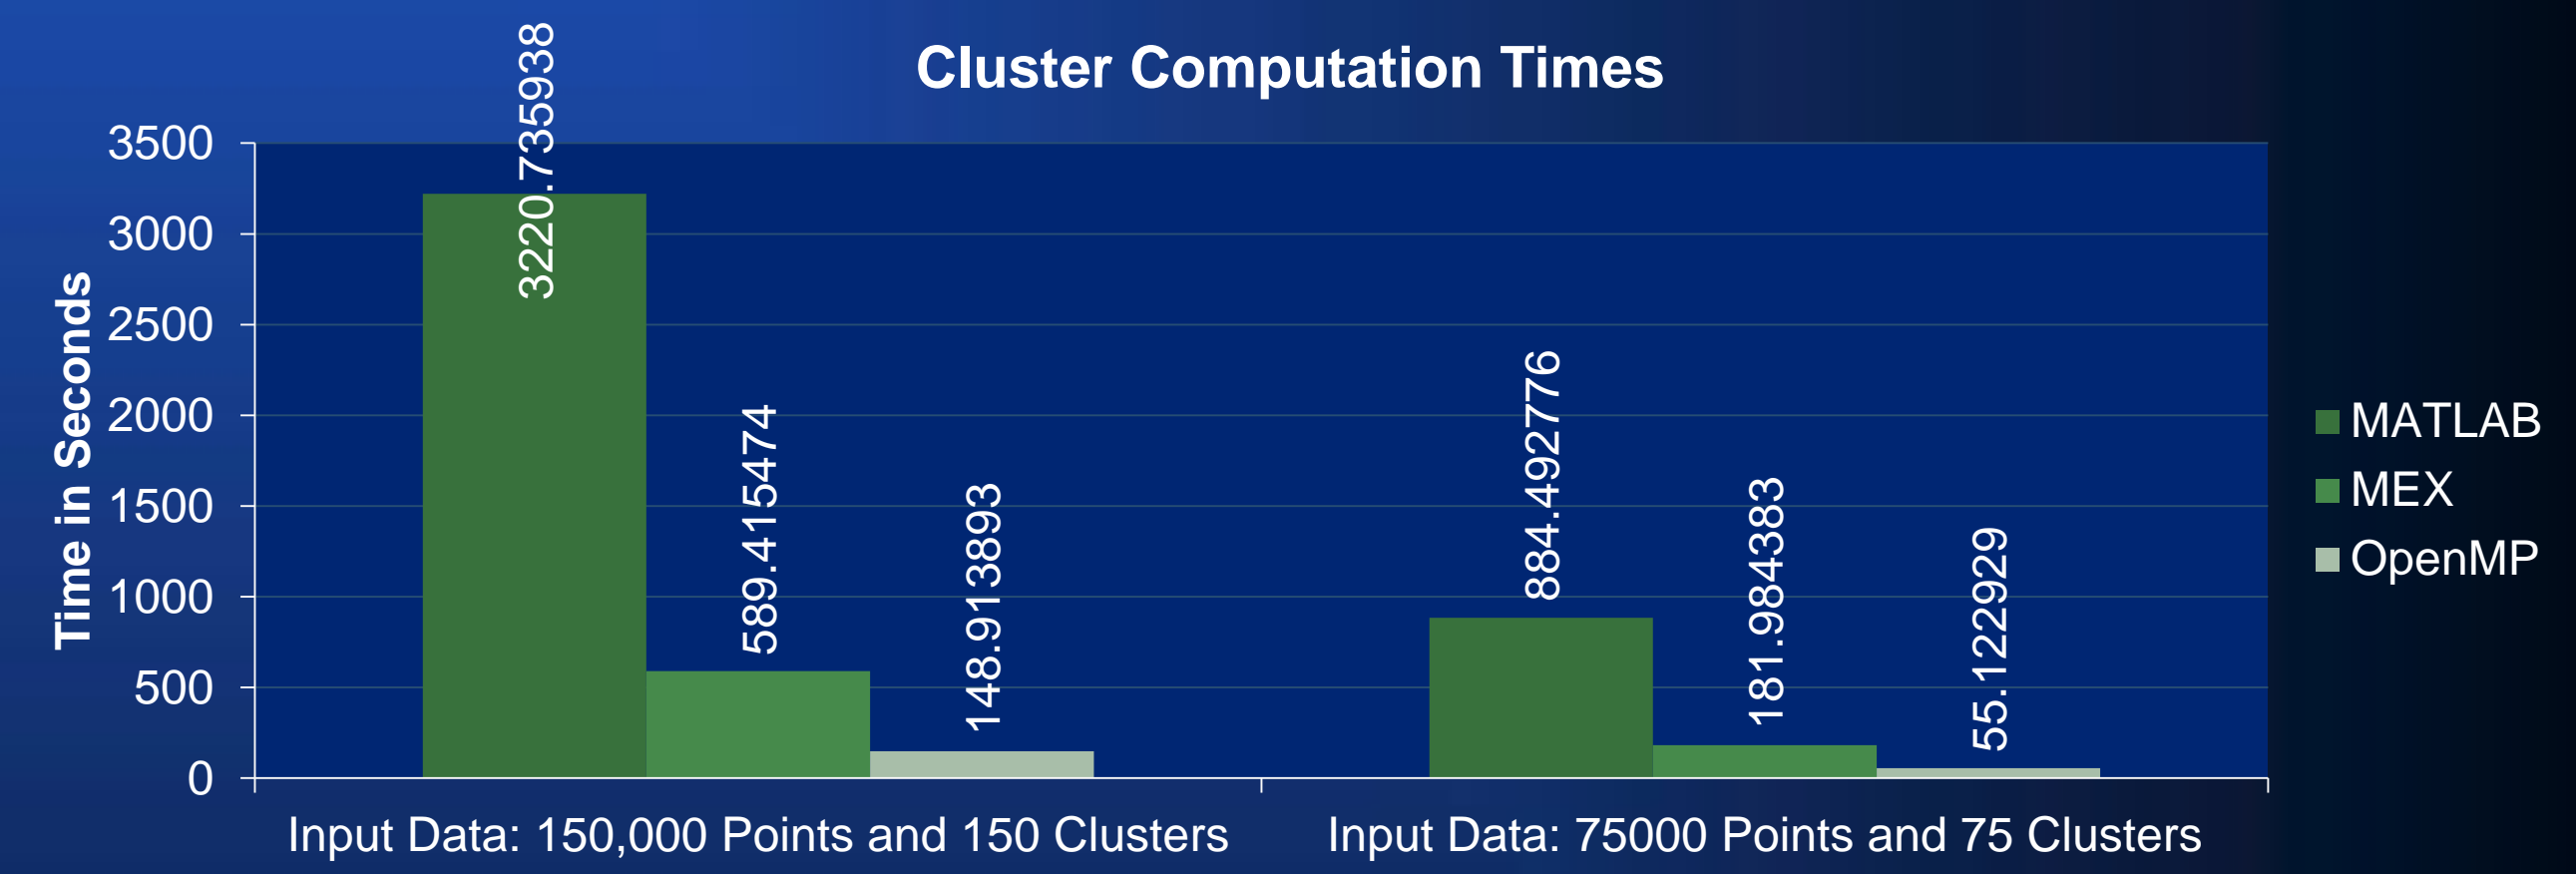




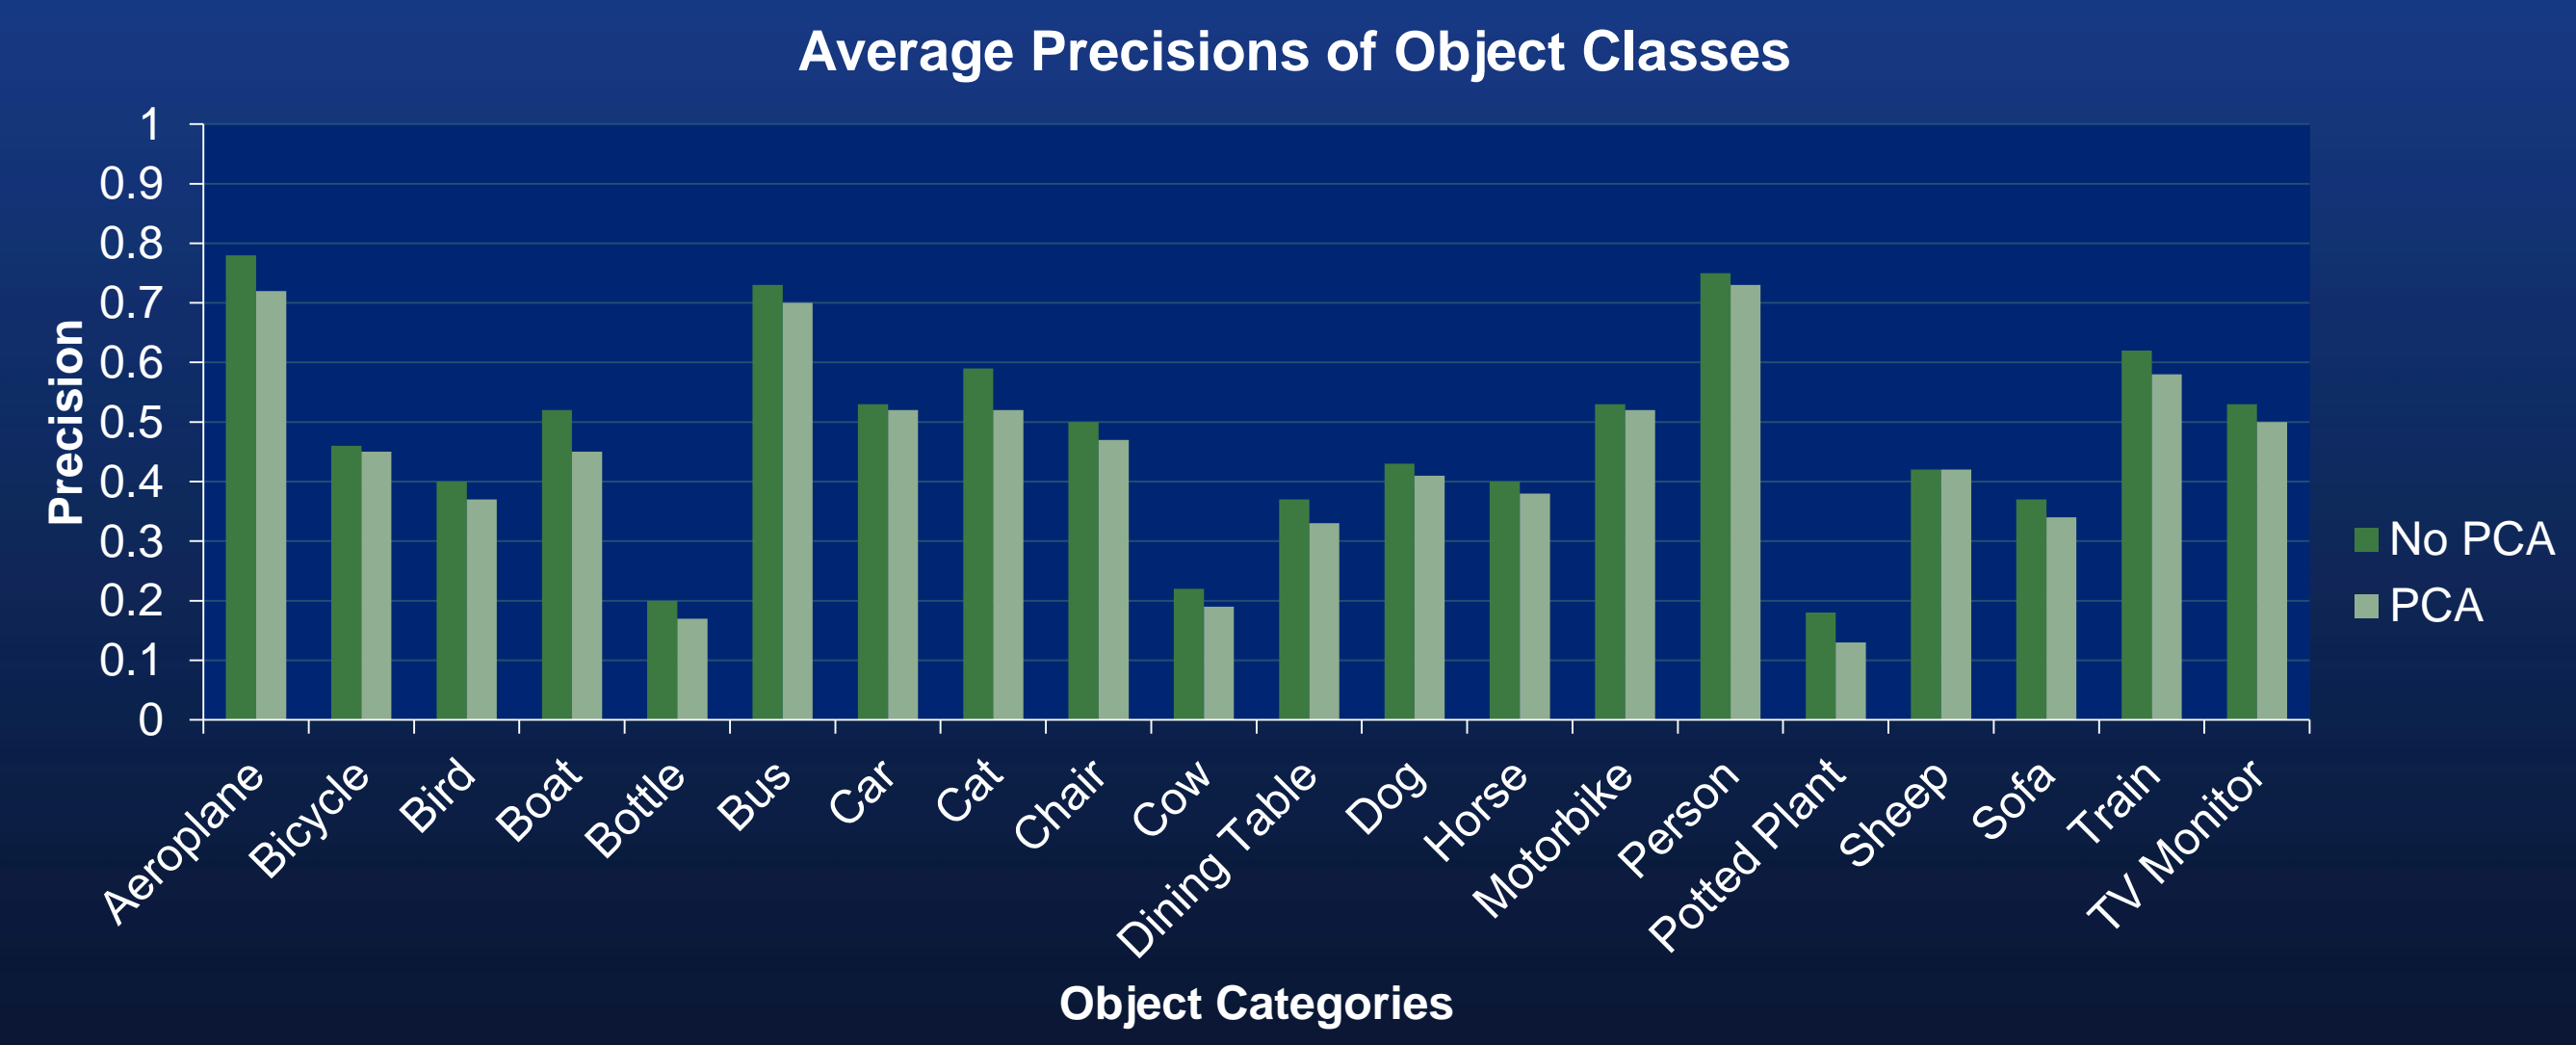
Figure 2 - With PCA. Less data to compute with little loss in accuracy

•Various Metrics – tried various metrics to compare the "similarity" of images.

Results



High Level Parallelism- send parts of each step, except clustering, to different computers to execute in parallel via Python.



PCA applied to SIFTS. Reduced dimensions from 128xN to 60xN with similar precision and accuracy as shown above.

- MATLAB - high level scripting language. Very slow.
- MEX (Matlab Executable) code – MATLAB calling C code. Very fast.
- OpenMP (Open Multi-Processing) and Pthreads (POSIX threads) – runs in parallel on multiple cores of a processor.
- CUDA (Compute Unified Device Architecture) - main computing engine on most NVIDIA graphics cards. Allows general purpose computation.

Acknowledgements
RE Lee Summer Scholar Program, Lenfest Grant, Steve Goryl