

A Characterization of Visual Feature Recognition

Binu Mathew, Al Davis, Robert Evans
School of Computing, University of Utah
{mbinu | ald | revans}@cs.utah.edu

Abstract

Natural human interfaces are a key to realizing the dream of ubiquitous computing. This implies that embedded systems must be capable of sophisticated perception tasks. This paper analyzes the nature of a visual feature recognition workload. Visual feature recognition is a key component of a number of important applications, e.g. gesture based interfaces, lip tracking to augment speech recognition, smart cameras, automated surveillance systems, robotic vision, etc. Given the power sensitive nature of the embedded space and the natural conflict between low-power and high-performance implementations, a precise understanding of these algorithms is an important step in developing efficient visual feature recognition applications for the embedded space. In particular, this work analyzes the performance characteristics of flesh toning, face detection and face recognition codes based on well known algorithms. We show that the problem can be decomposed into a pipeline of filters which could lead to efficient implementations as stream processors. With better than 92% hit rate for a modest 16KB L1 data cache, the algorithms have memory system behavior commensurate with embedded processors. However, our results indicate that their execution requirements strain the performance available on current embedded systems.

1. Introduction

The thrust of embedded computing has both diversified and intensified in recent years as the focus on mobile computing, ubiquitous computing, and traditional embedded applications has begun to converge. A side effect of this intensity is the desire to support sophisticated applications such as speech recognition, visual feature recognition, and secure wireless networking in a mobile, battery-powered platform. Unfortunately these applications are currently intractable for the embedded space. Running these applications on a low-power embedded processor cannot keep up with the inherent real-time processing requirements of speech recogni-

tion for example. The problem is that low-power processors do not have sufficient compute power. Using mainstream high-performance microprocessors comes close to meeting the performance requirements but the energy requirements are not commensurate with mobile and embedded processing domains. The first step in developing new architectures and systems which can adequately support these applications is a precise understanding of the algorithms.

Our focus has been on perception algorithms which form the basis for natural human interfaces to the embedded computing space. This is motivated by a belief that natural human interfaces are essential to realizing the dream of ubiquitous computing. In earlier work we have studied speech recognition and the results of that study have allowed us to create new architectures which support real-time, speaker independent, large vocabulary speech to text applications at a power level that is commensurate with embedded space energy budgets [8]. In this paper, we report on a similar study of visual feature recognition systems.

Visual feature recognition systems vary significantly based on the type of feature that is being recognized. Relatively simple recognizers are regularly employed in industrial visual inspection systems. On the other hand, human face recognition is an extremely complex task given the huge possibility space of facial features and skin tones. Facial recognition systems clearly have utility in security and surveillance domains, and other visual recognizers play key roles in gesture interfaces, lip reading to support speech recognition, and robotics. Our interest in face recognition however is primarily motivated by the difficulty of the problem which cannot be currently supported by embedded systems. Furthermore the structure of our face recognizer appears to be easily adapted to address other visual feature recognition tasks. The main differences for these other tasks is a different training regimen and different frame rate requirements. For example, the Rowley method of face detection described in this paper has been applied to license plate detection [11].

The particular application studied here can be viewed as a pipeline of 3 major functional components. A flesh tone detector is used to isolate areas of a frame where a face is

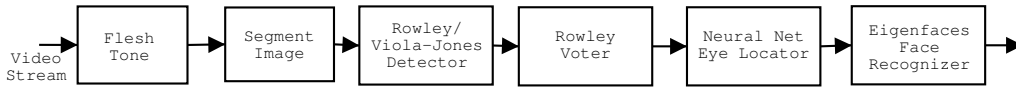


Figure 1. Algorithmic stages of a face recognizer

likely to be present. The next stage is a face detector which determines whether there is a face present or not. The final phase is a face recognizer. Each of these components is based on well known algorithms which have been adapted to fit into our framework. Some algorithmic optimization and restructuring has been done to suit our purposes but the basic approach has been developed by other researchers.

Interestingly the face recognition system when viewed from a structural perspective comprises a series of increasingly discriminating filters. Early stages of the sequence must inherently filter the entire image. As the process proceeds downstream, each stage needs to examine less image data since previous stages have eliminated certain areas from the probable candidate list. The result is an interesting balance of simple algorithms which analyze lots of data early in the sequence and more sophisticated algorithms which only need to analyze limited amounts of data late in the process. The simple sequence of filters structure makes face recognition amenable for implementation as an embedded system.

The following section provides a conceptual description of the recognition process and the algorithms involved in it. Section 3 describes the performance characteristics of the recognizer both by phase and for the whole application. Of particular interest is the percentage execution time of each phase, cache and function unit usage as well as CPU performance relative to real time requirements. Section 4 describes possible avenues for software and hardware optimizations followed by conclusions.

2. Overview of Visual Feature Recognition

Figure 1 shows the major steps in face recognition. The input is a low-resolution video stream such as 320x200 pixels at 10 frames per second. The stream is processed one frame at a time and sufficient state is maintained to perform history sensitive tasks like motion tracking. The process is essentially a pipeline of filters which reduce the data and attach attributes to frames for the use of down stream components. Typically each filter is invoked at the frame rate. This underlines the soft real time nature of this application. Additional data is required since filters may access large databases or internal tables. These additional data streams add to the aggregate bandwidth requirement of the system. The periodic nature of the application domain often makes it possible to easily estimate the worst case requirements.

Object recognition typically proceeds in two steps: object detection and the actual object identification. Most approaches to object identification require a clearly marked area, normalized to a particular size, and the location of key features. Object detectors find the area where the desired feature is likely to reside, scale the area to meet the normalization requirement, and then creates a location and boundary description for that area. False positives and negatives occur but the algorithms try and minimize their occurrence.

Object detectors also often work at a fixed scale. The detector is swept across the image recording all positions at which a detection was reported. The image is then subsampled or scaled down by a small factor (typically 0.8) and the process is repeated until the frame is below the size of the detector. A decision procedure is then applied to all the predicted hits to decide which ones are the most likely. Detectors often have much lower compute cost per sub-window than their corresponding identifying routines. Since they are swept across the entire image, a significant portion of the application's execution time might be spent in the detector. In contrast, even though identifying filters are more compute intensive, they are applied only to the high probability regions of the frame, so their contribution to the overall execution time might be low. Though object detectors are less compute intensive, they are much more difficult to design due to their generality. For example a face identifier chooses from one of N known faces, but a face detector has to distinguish between the infinite sets of faces and non-faces.

Since detection is time consuming, it is common to structure an object detector as a cascade of filters with cheaper heuristics upstream identifying potential regions for more expensive heuristics downstream. An extreme case of this is the Viola/Jones method which trains a sequence of about 200 increasingly discriminate filters [15]. A more common approach when dealing with faces and gestures is to identify the flesh colored regions of an image and apply a more sophisticated detector to those regions.

The identifier receives candidate regions from the detector along with other information like probability, scale and feature locations. It typically uses some sort of distance metric from known references to provide a positive identification.

In our face recognizer, the first level of detection is provided by identifying skin colors. This process known as flesh toning is further described in Section 2.1. This is followed by an image segmenting algorithm which groups skin

colored pixels into regions of interest where a face could likely be found. Segmenting is described in Section 2.2. The regions of interest are then applied to face detection algorithms. We evaluate two approaches known as the Rowley face detector and the Viola/Jones detector. The details appear in Sections 2.3 and 2.4 respectively. The detectors flag the regions that have a high probability of containing faces. However, they are likely to flag a detection at multiple nearby locations and image scales. So a voting algorithm is used to aggregate nearby detections into a single probable location. Finally a face identifier based on the Eigenfaces method is applied to the face locations to identify the faces from a database of known faces. The Eigenfaces approach is explained in Section 2.5. Since the face identifier requires the pixel coordinates of the eyes, we trained a neural network to discriminate between images that are eyes or not eyes. This network corresponds to the Eye Locator stage in Figure 1. It is swept across the pixel locations in a face to detect the location of the eyes.

2.1. Flesh Toning

Flesh toning identifies flesh colored pixels in an image. The commonly used RGB color space is not well suited for flesh toning because skin color occupies a wide range in the primary color space. Variations due to lighting and ethnicity are hard to deal with and skin-like colors on walls and clothing are difficult to discriminate. However, skin colors are tightly clustered in color spaces like HSV and YCbCr. Flesh toning can be done by converting pixels from sample images into the chosen color space and making a scatter plot with two tags, one for flesh pixels and one for non-flesh pixels. A boundary is then drawn around flesh tone clusters. This boundary is then approximated by curves which can be described by simple geometric equations. In the image under test, any pixel that lies inside this new approximated but easily described boundary is considered to be a flesh pixel. For improved accuracy we employ the consensus of two separate flesh toning algorithms based on the NCC (Normalized Color Coordinates) and HSV color spaces respectively [13, 2]. The output of this phase is a bit mask of the same size as the image where bits are set if the corresponding pixel is flesh colored.

2.2. Segmentation

Segmentation is the process of clumping together individual pixels into regions where an object might be found. A common approach is to do a connected component analysis which typically forms irregular regions. Since the Viola and Rowley algorithms need rectangular regions, we use a slightly different approach. An *erosion* operator examines each pixel and blacks it out unless all its neighbors in a

3x3 pixel window are set [5]. This is followed by a *dilation* operator which examines each pixel and lights up that pixel if any of its neighbors in a 4x4 window are set. The net result is that small connections like loose clothing or hair that cause false connections are cut away and the image is smoothed by filling small holes. The outputs of this stage are: a) A list of coordinates that specify the rectangular bounding boxes of pixel clusters b) A gray scale version of the image.

2.3. Rowley Face Detector

Henry Rowley's neural net based face detector is well known as a pioneer in this field [11]. Its implementation was provided to us by the Robotics Institute at CMU. This detector is designed to determine if a 20x20 pixel image contains a face or not. Face detection is done by sweeping the detector over a larger image and computing the decision at each pixel location. Then the image is scaled and reduced in size by a factor of 0.8 and the procedure is repeated. The resulting series of images and detection locations is called an image pyramid. In the case of real faces, a detection will be reported at several nearby pixel locations at one scale and at corresponding locations in nearby scales. False positives do not usually happen with this regularity. Hence a voting algorithm (our terminology) can be applied to the image pyramid to decide the site of any true detections. To further reduce false positives, multiple separately trained neural nets may be applied to the image and their consensus can represent a more accurate detection.

2.4. Viola and Jones' Detector

Viola and Jones present a new and radically faster approach to face detection based on the AdaBoost algorithm from machine learning [15, 12]. They claim a 15x speedup over the Rowley detector for their implementation without using flesh toning. Since their source code is proprietary, we re-implemented their algorithm based on example code obtained from the University of British Columbia. Our re-implementation uses flesh-toning for both the Rowley and the Viola/Jones detectors and the algorithms are close in speed due to factors mentioned later.

The AdaBoost algorithm can combine several weak heuristics to generate a strong heuristic, while providing statistical bounds on the training and generalization error. The weak heuristics used by the Viola/Jones algorithm are based on three types of *features* which are essentially the sum or difference of all the pixels in adjacent rectangular regions. Face detection is reduced to computing the weighted sum of chosen rectangle-features and applying a threshold. Similar to the case of the Rowley algorithm, a 24x24 detector is swept over every pixel

location in the image and the image is rescaled. We apply Rowley’s voting algorithm to decide the final detection locations.

The original slow approach described in the Viola/Jones paper uses 200 features. They then go on to describe a faster approach where they cascade many detectors with more complex detectors following simpler ones. A window is passed to a detector only if it was not rejected by the preceding detector. Since training this cascade is a laborious process, we model the workload characteristics of this algorithm with a single 100 feature detector. Hence, the Viola/Jones and the Rowley approaches are similar in performance in our implementation.

2.5. Eigen Faces

Eigenfaces is a well known Principle Component Analysis (PCA) based face recognition algorithm developed by researchers at MIT [14]. In this paper, we use a re-implementation of the Eigenfaces algorithm from researchers at Colorado State University [4]. Face images are projected onto a feature space called *face space* defined by the eigen vectors of a set of faces. This captures the variation between the set of faces without emphasis on any one facial region like the eyes or nose. The mathematical treatment of Eigenfaces is too involved to discuss here. The approach works by computing and storing the face space corresponding to each face in a training set. A test image is projected on to each saved eigenface and a set of weights is computed based on closeness to the known eigenfaces. The weights are then used to label the test image as one of the known persons or an unknown one. In our evaluation the eigenfaces database contains 10 personalities from television.

3. Characterization

In this section we provide a detailed characterization of visual feature recognition by native execution and profiling using processor performance counters as well as via simulation. The native profiling results that appear later in Figures 2 and 3 were obtained using SGI SpeedShop on a 666 MHz R14K processor. Simulation studies are based on MLRSIM, an out of order processor simulator derived from the Rice University RSIM simulator. It is detailed enough to run a derivative of the NetBSD operating system and can run SPARC binaries compiled for Sun OS without any modification. A multi-GHz processor is required to operate this application in real time. Parameters like L1 cache hit time, memory access time, floating point latencies etc. were measured on a 1.7 GHz AMD Athlon processor using the lmbench hardware performance analysis benchmark [10]. Numbers that could not be directly measured were

| |
|---|
| <p>Default Configuration: Out of order processor at 2 GHz, SPARC V8 ISA 4 ALU, 4 FPU, Max 4 issue Max 4 graduations per cycle 16KB 2 way L1 I and D cache with 2 cycle latency 2MB 2 way L2 Cache, 20 cycle latency 600 MHz, 64 bit DRAM interface.</p> |
| <p>Embedded Configuration: In order processor at 400 MHz, SPARC V8 ISA 1 ALU, 1 FPU, Single issue Single graduation per cycle 32 KB 32 way L1 I and D cache with 1 cycle latency, No L2 cache, 100 MHz 32 bit DRAM interface.</p> |
| <p>Application: gcc 2.95.2 compiled binary for SunOS</p> |

Table 1. Experiment Parameters

obtained from vendor micro-architecture references. MLRSIM was configured to reflect these parameters reported in Table 1. Unless mentioned otherwise, the remainder of this paper uses the default configuration.

Though the embedded configuration in Table 1 simulates a SPARC V8 ISA, its architectural parameters were chosen to model an Intel XScale (StrongARM) development system with the exception that the processor has a floating point unit. Though the XScale does not have an L2 cache, since MLRSIM cannot be configured without an L2 cache, this configuration has a 64KB *inclusive* L2 cache. Since the cache is inclusive and the same size as the sum of the L1 caches, this configuration behaves similar to a machine with no L2 cache.

We discuss characteristics of the application in 5 configurations: a) full pipeline using the Rowley face detector, b) full pipeline using the Viola/Jones face detector, c) only the Rowley face detector with flesh toning and image segmentation, d) only the Viola/Jones face detector with flesh toning and image segmentation, e) only the Eigenfaces recognizer. The last three configurations are important from an energy savings perspective since running the individual algorithms on separate low frequency processors or hardware accelerators, can lead to significant energy savings.

Figures 2 and 3 show the relative execution times of each algorithm using the Rowley detector and the Viola/Jones detector. Figures 4 and 5 show the L1 Dcache miss rate and the L2 cache hit rates for all 5 application configurations. Since the caches are inclusive, the L2 hit rate is defined as the L1 misses that hit in the L2 cache divided by the total number of accesses made by the application. Since this application achieves 99.8% ICache hit rate with a 16KB ICache, no other ICache configurations were studied.

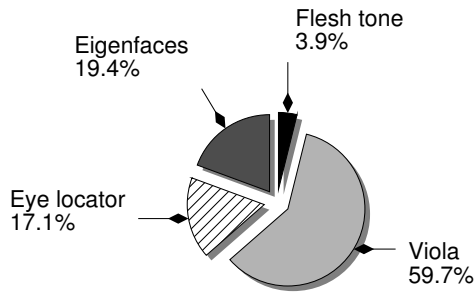


Figure 2. Execution time break down of Viola/Jones detector based face recognizer

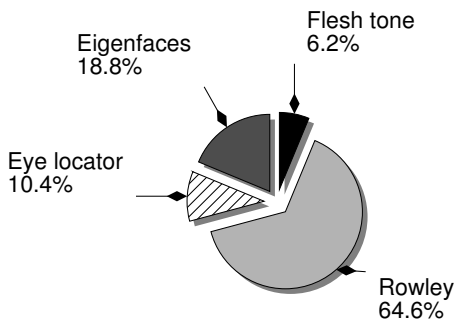


Figure 3. Execution time break down of Rowley detector based face recognizer

For the entire application there is consistently greater than 92% L1 cache hit rate for D-caches of 16KB and above. This indicates that the streaming pipelined model we use for composing the algorithms is a good fit for the problem. Each 320x200 pixel color image is 187.5 KBytes long and the corresponding gray scale versions are about 64 KBytes. The images clearly will not fit in the L1 cache. The explanation is that the color image is accessed in streaming mode, i.e. each pixel is touched exactly once for flesh toning. Image segmentation works on the flesh tone bitmap (approximately 64KB) making at most two passes over it. Since these accesses touch at most two image rows at a time, good cache utilization is ensured. Subsequently, only small windows into the image are used. Since objects in these images are typically smaller than 50x50 pixels, each object is only about 2.5KB in size. The downstream algorithms make several passes over each object, but only a small part of each object needs to be cache resident at each time. For exam-

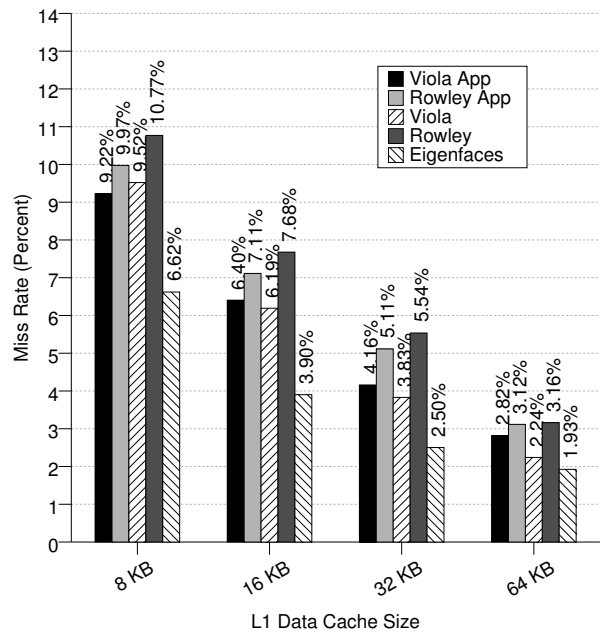


Figure 4. L1 Dcache miss rate

ple, the integral image computation which is a part of evaluating the rectangular features in the Viola/Jones algorithm is based on a recurrence that involves two adjacent image rows and an additional row for intermediate storage and has an L1 cache footprint of about 4.4KB. The Rowley algorithm touches at most 20 rows of the object at the same time. However, as it sweeps across the image left to right and top to bottom only a 20x20 pixel window needs to be cache resident at a time. Since it shifts its position one pixel at a time, a 19x19 region of this window will be reused by the next iteration contributing to high L1 cache hit rate. A similar pattern occurs in the later phase of the Viola/Jones algorithm on a 24x24 region. The Eigenfaces algorithm uses a projected image of the object to be recognized as well as basis, mean and projected image matrices corresponding to each reference object. The target object is reused while it is compared against each candidate. Each candidate however is accessed only once per target object.

The objects and their attributes from each stage are typically touched again by the next stage. The auxiliary information used by the algorithms is somewhat small. Both detector algorithms use fixed size data structures. The worst case is the Viola/Jones algorithm which needs a weight and a type for each feature corresponding to $100 * 2 * 4 = 800$ bytes of L1 cache usage. The data set for the Eigenfaces algorithm on the other hand is linear in the number of the reference faces. But since these could potentially be streamed into the L1 Dcache once per target object (or once per frame) its footprint is small. Only the projected tar-

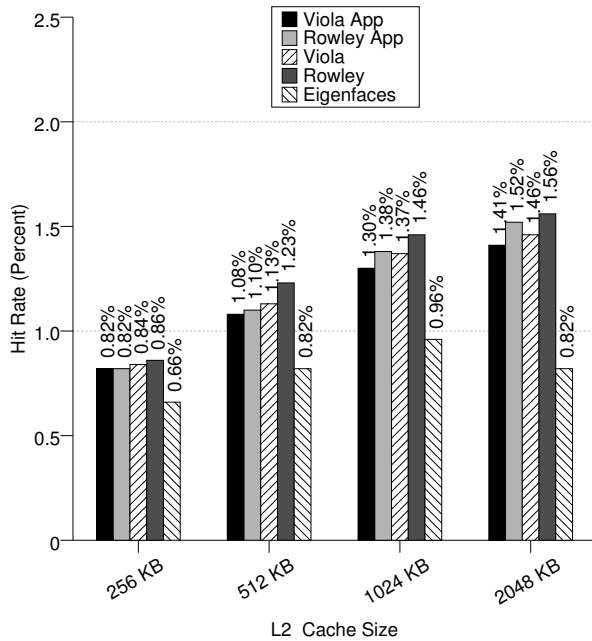


Figure 5. L2 Cache hit rate

get object and a small part of the basis/mean/projected reference images need to be resident in the L1 Dcache. From Figure 5 it is seen that the L2 cache is largely ineffective since it is accessed infrequently due to the low L1 miss rate.

From a cache footprint perspective, both the detector algorithms and the entire application appear to be a good match for embedded processors with limited cache resources. Since images are accessed left to right, multiple rows at a time, sequential prefetch (or strided prefetch) would hide memory access latencies even when the L1 Dcache is small. However, quite a different view unfolds on examining the IPC and speedup characteristics. Figure 6 shows IPC for a variety of execution unit configurations. IPC is seen to saturate early on for two main reasons. The first is caused by dependences in the loop bodies. For example, neural net evaluation involves computing $\sum_{i=0..n} Weight[i] * Input[i]$. In addition to the loop carried dependence on the sum, each of the inputs is accessed indirectly via a pointer since an input to one neuron could be the output of another neuron. Secondly the high ratio of array variable accesses to arithmetic operations causes saturation of the Dcache ports.

Figure 7 shows the run times normalized to real-time. Here, 1.0 represents minimum real time performance corresponding to 5 frames per second. For example, in Figure 7 in the 1 ALU + 1 FPU configuration, the Rowley algorithm is 1.13 times slower than real-time while the Eigenfaces algorithm processes 5 frames in 0.69 seconds. Clearly, embedded processors are inadequate to handle the work load

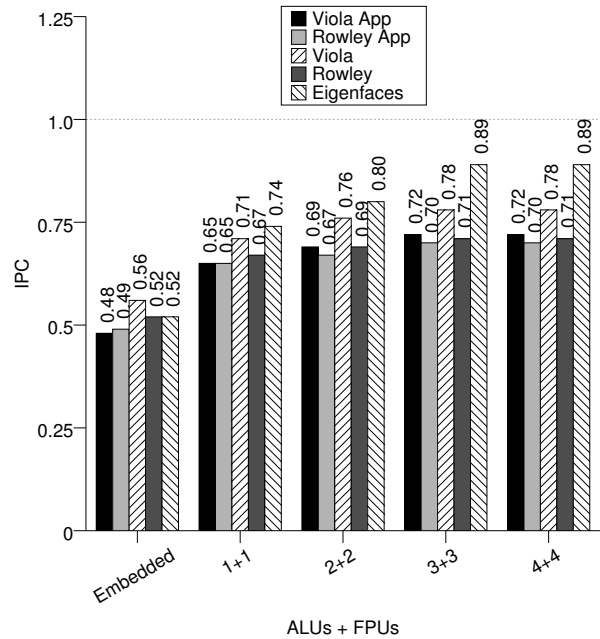


Figure 6. IPC

in real time. In this case instruction throughput is the culprit. Even when function units are available, dependences and contention for the Dcache ports causes low IPC. The power budgets required for real-time performance are beyond what is available on normal low power embedded platforms. Thermal dissipation is a problem even on high performance processors and energy saving solutions are important for real time workloads like visual feature recognition. Hardware accelerators that use specialized data paths and stream array operands out of multiple SRAM buffers stand a good chance of accelerating these algorithms at embedded power budgets.

4. Opportunities for Optimization

One recurring theme in image processing is computing a kernel that operates on an $M \times N$ sub-window of a larger $W \times H$ image. The kernel is re-computed for every possible pixel location within the larger image. This resembles sliding the $M \times N$ sub-window over the $W \times H$ image. There is significant scope for compiler based re-ordering of computations in such kernels. We provide two concrete examples.

As described in Section 2.4, the heuristics used by the Viola/Jones algorithm are based on the sum/difference of pixels in adjacent rectangular regions. Recomputing these sums for each pixel location is very expensive. A major contribution of their approach is an intermediate image representation called the integral image. The sum of the pixels in a rectangular window can be computed easily using the in-

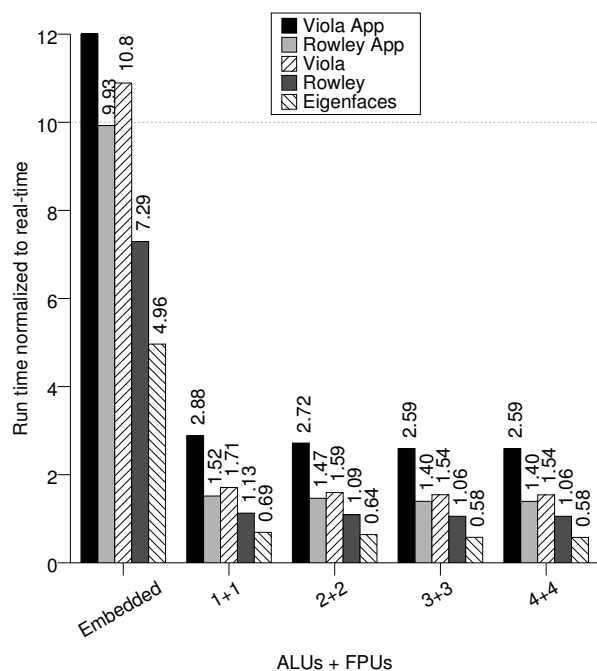


Figure 7. Speedup or Slow down over real time

intermediate representation. The integral image value at pixel location (x,y) in an image is defined as the sum of all pixels to the left and above the pixel (x,y) . This is computationally prohibitive. By expressing the same relationship as a pair of recurrences, it is possible to compute the integral image with just one pass over the image. This transformation required careful study and insight from the originators of the algorithm. But given the fact that the sums of rectangular sub-windows of the larger image are recomputed at each pixel location, a compiler based tool aware of the access pattern and rules of arithmetic could be made to deduce the recurrences.

While computing the Viola/Jones heuristics we also need the standard deviation of pixel values within a 24×24 pixel window starting at each pixel location within the probable face rectangle. Our initial implementation that simply recomputed the standard deviation function at each pixel location was seen to occupy between 10-15% of the compute time of the whole application. When going from one pixel to the next, the windows overlap by 23×24 pixels and the mean and sum of squares for one pixel can be easily calculated from its predecessors values by adjusting for the non-overlapping pixels alone. By defining a set of recurrences for the mean and mean square for 24×24 sub windows over a wider region, we were able to compute the standard deviations in one pass over the image thereby reducing the ex-

ecution time of this component to less than 1%. Currently, such transformations require a lot of attention from the programmer and insight into the algorithm and are error prone because of corner cases. This bolsters the argument in favor of compiler based loop re-structuring that can apply axioms of algebra to deduce the right set of recurrences.

Another possible optimization is to reorder the computation so that data may be streamed through a set of execution units and results computed in the minimum number of passes while observing limits on the amount of intermediate storage used. Compiler based tools that can accept a kernel, the size of the image and sub-window as well as the size of the intermediate storage available on the processor and effect the transformation automatically have an opportunity to enhance performance and reduce resource utilization significantly.

As seen in Figures 6 and 7, wide issue clearly helps performance. But wide issue comes at the cost of increased hardware complexity and could potentially limit the clock frequency as well as exceed a limited energy budget. This application is embarrassingly parallel in most sections due to the intrinsic data parallelism in the pixel processing. One way of achieving good performance at low power is to use a cluster of simple processors or function units operating in parallel with a very small quantity of SRAM for local storage and no cache.

In our previous research with speech recognition systems we have observed that an order of magnitude or better improvement was possible in the energy delay product using such an approach. We expect similar results for visual feature recognition. As a concrete example, it has been demonstrated that using such a statically scheduled cluster approach for neural network evaluation improved the process normalized energy consumption by a factor of 17 and the energy delay product by a factor of 8 over a 2 GHz Pentium 4 [9]. Recall that the Rowley detector is based on neural networks.

5. Related Work

Perception processing which encompasses a wide range of topics like computer vision, speech recognition and gesture recognition is currently the focus of vigorous research. While it is common in the literature to see the relative merits and performance of algorithms compared, architecture level analysis of whole perception applications is extremely rare. Wang, Bhat and Prasanna discussed methods for implementing scalable computer vision algorithms on commercial parallel computers [17]. The Image Understanding Architecture, a parallel processor for real time machine vision and its software environment developed jointly by the University of Massachusetts and Hughes research is described in [19]. Wawrzynek et al described the performance of the

SPERT II, a vector microprocessor based hardware accelerator for neural network algorithms [18]. Agaram, Keckler and Burger presented a detailed architecture level analysis of the CMU Sphinx II speech recognizer [1].

The next step in this research is to investigate special purpose architectures for machine vision. This approach has shown significant advantages in both performance and energy consumption for speech recognition [8]. Existing efforts on vision architectures can be partitioned into analog and digital approaches. The analog approaches have primarily been based on modeling the human neural system. This neuromorphic approach has produced a variety of efficient special purpose devices for early-vision functions [6, 7, 20]. Digital approaches have covered a broader spectrum of vision functions. Representative examples are commercial offerings by companies such as Cognex and Coreco (www.cognex.com, www.coreco.com) which provide application specific software for industrial applications such as visual inspection, security monitoring, motion detection, etc. These commercial approaches employ standard processor and memory architectures but may or may not dispense with a general operating system environment. Others have experimented with directly mapping algorithm flows onto FPGA based systems, and the utility of new highly parallel architectures such as the MIT RAW machine for vision applications [3, 16]. ASIC implementation of vision applications are rare, perhaps due to the significant diversity of computer vision algorithms and the high cost of ASIC development. Each approach offers distinct advantages but is also limited by intrinsic disadvantages. General purpose processors are slow when compared to ASIC approaches. Analog circuits are both fast and energy efficient but they lack generality and are costly to produce. FPGA based systems are a compromise between generality and efficiency. We feel that there is a significant opportunity for a much better compromise which involves a customizable compute cluster which retains most of the generality of the generic processor approach while achieving performance/energy efficiency levels close to that of special purpose ASIC approaches.

6. Conclusions

We have presented a detailed analysis of the performance characteristics of a face recognition system based on well-known algorithms. Existing face recognition systems are inadequate to support real-time operation on embedded systems. We have shown that by taking advantage of the stream oriented nature of the application it is possible to solve this problem. It is also evident that other visual feature recognizers can benefit from similar tactics since our face recognizer needs few modifications in order to recognize other visual features. This effort and our previous work on speech

recognition systems leads to a strong belief that recasting perception algorithms into a stream oriented style is the key to improving performance, reducing power consumption, and supporting these sophisticated applications on embedded devices with limited cache, processing, and energy resources.

7. Acknowledgments

The authors would like to thank the Vision & Autonomous Systems Center (VASC) at the Robotics Institute of the Carnegie Mellon University and Tsuyoshi Moriyama in particular for providing us with the source code of the CMU/Rowley face detector. We would also like to thank Peter Carbonetto of the University of British Columbia for providing the example code which formed the starting point for our AdaBoost face detector.

References

- [1] AGARAM, K., KECKLER, S. W., AND BURGER, D. A characterization of speech recognition on modern computer systems. In *Proceedings of the 4th IEEE Workshop on Workload Characterization* (Dec. 2001).
- [2] BERTRAN, A., YU, H., AND SACCHETTO, P. Face detection project report. <http://ise.stanford.edu/2002/projects/ee368/Project/reports/ee368group17.pdf>, 2002.
- [3] BONDALAPATI, K., AND PRASANNA, V. Reconfigurable computing systems. *Proceedings of the IEEE* 90, 7 (July 2002), 1201–1217.
- [4] D. BOLME, R. BEVERIDGE, M. T., AND DRAPER, B. The CSU Face Identification Evaluation System: Its Purpose, Features and Structure. In *International Conference on Vision Systems* (April 2003), pp. 304–311.
- [5] HARALICK, R., AND SHAPIRO, L. *Computer and Robot Vision*, vol. 1. Addison-Wesley Publishing Company, 1992, ch. 5.
- [6] HIGGINS, C. Multi-chip neuromorphic motion processing. In *Proceedings of the Conference on Advanced Research in VLSI, Atlanta* (March 1999).
- [7] MAHOWALD, M. A., AND MEAD, C. *Analog VLSI and Neural Systems*. Addison-Wesley, 1989, ch. Silicon Retina, pp. 257–277.
- [8] MATHEW, B., DAVIS, A., AND FANG, Z. A Low-Power Accelerator for the SPHINX 3 Speech Recognition System. In *Proceedings of the International Conference on Compilers, Architecture and Synthesis for Embedded Systems (CASES '03)* (October 2003).
- [9] MATHEW, B., DAVIS, A., AND IBRAHIM, A. Perception Coprocessors for Embedded Systems. In *Proceedings of the Workshop on Embedded Systems for Real-Time Multimedia (ESTIMedia)* (October 2003).
- [10] McVOY, L. W., AND STAELIN, C. Imbench: Portable tools for performance analysis. In *USENIX Annual Technical Conference* (1996), pp. 279–294.

- [11] ROWLEY, H. A., BALUJA, S., AND KANADE, T. Neural network-based face detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, 1 (1998), 23–38.
- [12] SCHAPIRE, R. E. The boosting approach to machine learning: An overview. In *In MSRI Workshop on Nonlinear Estimation and Classification* (2002).
- [13] SORIANO, M., MARTINKAUPPI, B., HUOVINEN, S., AND LAAKSONEN, M. Using the skin locus to cope with changing illumination conditions in color-based face tracking. In *Proceedings of the IEEE Nordic Signal Processing Symposium* (2000), pp. 383–386.
- [14] TURK, M., AND PENTLAND, A. Face recognition using Eigenfaces. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (June 1991), pp. 586–591.
- [15] VIOLA, P., AND JONES, M. Rapid object detection using a boosted cascade of simple features. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (Dec. 2001).
- [16] WAINGOLD, E., TAYLOR, M., SRIKRISHNA, D., SARKAR, V., LEE, W., LEE, V., KIM, J., FRANK, M., FINCH, P., BARUA, R., BABB, J., AMARASINGHE, S., AND AGARWAL, A. Baring it all to software: Raw machines. *IEEE Computer* 30, 9 (1997), 86–93.
- [17] WANG, C.-L., BHAT, P. B., AND PRASANNA, V. K. High performance computing for vision. *Proceedings of the IEEE* 84, 7 (July 1996), 931–946.
- [18] WAWRZYNEK, . J., ASANOVIC, K., KINGSBURY, B., BECK, J., JOHNSON, D., AND MORGAN, N. Spert-ii: A vector microprocessor system. *IEEE Computer* 29, 3 (March 1996), 79–86.
- [19] WEEMS, C. C. The second generation image understanding architecture and beyond. In *Proceedings of Computer Architectures for Machine Perception* (Nov. 1993), pp. 276–285.
- [20] WYATT, J. L., KEAST, J. C., SEIDEL, M., STANDLEY, D., HORN, B., KNIGHT, T., SODINI, C., LEE, H., AND POGGIO, T. Analog vlsi systems for early vision processing. In *Proceedings of the 1992 IEEE International Symposium on Circuits and Systems* (May 1992), pp. 1644–1647.