MULTIRATE SEPARABLE IMPLEMENTATION OF STEERABLE FILTER BANKS

Roberto Manduchi and Pietro Perona

California Institute of Technology 116–81
Pasadena, CA 91125
{manduchi,perona}@vision.caltech.edu

ABSTRACT

This paper describes efficient schemes for the computation of a large number of differentely scaled/oriented filtered versions of an image. Previous work has generalized the well-known steerable/scalable ("deformable") filter bank structure by imposing X-Y separability on the basis filters. This systems, designed by an iterative projections technique, was shown to achieve substantial reduction of the computational cost.

To reduce the memory requirement, we adopt a multirate implementation. The resulting structure, however, is not shift-invariant - it gives raise to "aliasing". We introduce a design criterion for multirate deformable structures that jointly controls the approximation error and the shift-variance.

1. INTRODUCTION

Elementary visual structures such as lines, edges, texture, motion, are powerful "cues" to understand the structure of the outside world from its visual appearance (the image), and their identification is instrumental for almost any visual task. Classical image processing problems (enhancement, denoising) may also be approached successfully using elementary descriptors such as edges and textures. Image compression schemes using sub–bands coders, oriented along the preferred texture orientation in the image, proved advantageous in terms of visual rendition. Velocity may be interpreted as orientation in spatio–temporal domain, and motion-compensated spatio-temporal filters may be used succesfully for prediction, interpolation and smoothing as well as for coding.

Regardless of the specific descriptor of interest, most techniques start processing the image (or image sequence) with a family of linear filters tuned at a wide range of orientations and scales of resolution. The multiscale/multiorientation image decomposition is then analyzed to detect features (usually, via a non-linear stage) and to measure their attributes (orientation, dominant scale, velocity).

While orthogonal structures have been intensively studied in the context of wavelet theory, several algorithms are designed to operate on redundant (or overcomplete) image decompositions. A drawback of this approach is that the computational cost to realize the analysis filter bank may easily become too high for practical use. In order to meet prescribed implementation constraints, thus, the use of fast filtering is mandatory.

This paper describes efficient schemes for the computation of a large number of multiscale/multioriented filtered versions of an image, pushing forward previous results by Freeman and Adelson [1] and Perona [2].

To reduce the computational weight of multiscale/ multiorientation filter banks, one may exploit the correlation among the filters in a two-stage scheme usually named "steerable" or "scalable" [1, 3] (or more generally, "deformable" [2]). Such systems may be shown [4] to be a particular instance of the "multistage separable" structure introduced by Treitel and Shanks [5]. In a previous paper [6], we generalized the deformable filter bank structure by imposing X-Y separability on the basis filters. Such an implementation reduces effectively the computational weight associated with deformable filter banks. An iterative projections technique was used for the least-squares design of such structures.

While deformable filter banks are efficient in terms of computational cost, they require substantial extra memory to store intermediate frames. To reduce the overall memory requirements, we propose to use a multirate implementation: for those filters in the filter bank that are narrow-banded, only a subsampled version of the intermediate filtered frames needs to be stored.

In our implementation, we embed the filter bank in an analysis/system separable pyramidal scheme. The system is designed using a novel least-squares procedure for multirate FIR filters [7] introduced by Manduchi and Perona [8]. This technique, based on the time domain, relies on the

We gratefully acknowledge support from NSF Grant IRI 9306155 on "Geometry driven diffusions", NSF Research Initiation grant IRI 9211651, ONR grant NOOO14-93-1-0990, a NSF National Young Investigator Award to P.P.. This work is supported in part by the Center for Neuromorphic Systems Engineering as a part of the National Science Foundation Engineering Research Center Program; and by the California Trade and Commerce Agency, Office of Strategic Technology.

0-7803-4455-3/98/$10.00 © 1998 IEEE

definition of a suitable "multirate" approximation criterion, which jointly controls the approximation error and the shift-variance. The resulting system achieves the goal of joint reduction of computational weight and of memory requirement.

2. MULTIRATE DEFORMABLE FILTER BANKS

Imposing X-Y separability on the basis filters of a deformable filter bank, as described in [6], can reduce the overall computational weight, at the price of supplementary memory for storing the intermediate filtered images. Frame memories are expensive, and it is desirable to reduce the memory requirement while enjoying reduced computational cost.

A solution is derived observing that some of the basis filters may be narrow-banded. The filtered signals are thus highly correlated, and may be subsampled before being stored (therefore using less memory). These reduced rate versions are then interpolated and linearly recombined for each scale and orientation. Although this may seem computationally expensive at a first sight, we can show that with a suitable choice of the interpolator filters, the overall implementation results efficient both in terms of memory and of computational cost.

The idea of multirate implementation of digital filters dates back to the 1975 paper of Rabiner and Crochiere [7]. In the context of computer vision, Burt [9] designed multirate filters shaped as gaussians or as laplacian of gaussians. More recently, Manduchi and Perona [8] proposed a least–squares design procedure for multirate FIR filters. This last approach differs from the previous ones in the sense that the shift-variance (consequent to the sampling rate alteration) is explicitly kept into account.

The multirate systems we consider here are not perfect reconstruction, and therefore suffer from aliasing, meaning that they are not shift-invariant. Shift-invariance is considered a fundamental property of signal processing systems; on the other side, a system that deviates "slightly" from shift-invariance may still be suitable to most applications. Typical effects of aliasing are staircase patterns in correspondence of brightness edges, and low–frequency fringes (Moirè patterns) in correspondence of textured areas. Both phenomena are visually quite noticeable. However, if our task is the analysis of the visual structures in an image, aliasing may be tolerated as a signal–dependent noise, qualitatively not too different from other forms of "noise" that are traditionally considered.

How do we quantify "shift-variance", and how can we jointly control the approximation error and the shift-invariance of our filters? Manduchi and Perona [8] regarded shift-variance as data-dependent noise, which can therefore be quantified. By studying the behavior of the multirate system in the time domain, one can derive an error functional that keeps into account both the approximation error and the "aliasing" noise. The multirate filter may be optimized using an iterative procedure.

In the case of separable deformable filter banks, we adopt the multirate implementation for those 1-D basis filters that are suitably narrow-banded. A simple design strategy is the *a posteriori* multirate approximation: first, a separable deformable filter bank is designed; then, each narrow-banded 1-D filter is transformed into a multirate structure using the technique of [8].

However, a more clever approach optimizes for the whole multirate system having *a priori* constrained the multirate structure of a number of kernels. The algorithm described in this section achieves such a goal, jointly optimizing for the multirate 1-D filters and for the recombination functions.

We use a pyramidal octave–band separable structure for the steerable/scalable filter bank, represented in Figure 2. Although this is not the most general scheme for a multirate realization, we claim a number of design and implementation advantages deriving from the pyramidal decomposition. The basis filters may be embedded in (and optimized for) *any* separable pyramidal structure, although the interpolation (synthesis) filters need to be "short" in order to achieve computational savings. The proposed implementation uses simple raised-cosine analysis and synthesis filters.

It is important to realize that this structure is different from the "steerable pyramid" of Simoncelli *et al.* [3]. In fact, the input and the output of our system are at the same rate; the pyramid is merely used as an efficient computational scheme.

2.1. Pyramidal implementation of deformable filter banks

As mentioned earlier, we use multirate implementation for those 1-D filters of the X-Y separable deformable filter bank which are suitably band-limited. The computational burden of the basis filtering is reduced, although it is likely that, in order to compensate for the approximation error introduced by the multirate implementation, the number of branches in the filter bank will have to be increased. In fact, important advantages of the multirate implementation are:

1. Reduction of memory requirements (because only the decimated version of the basis filters' outputs are stored);
2. Reduction of the computational burden for the recombination (because multiplications by the recombination coefficients are performed on the lowest rate signals).

The basic multirate structure is represented in Figure 1. The overall pyramidal structure is shown in Figure 2.

In particular, we set the decimation filter equal to the interpolator, and we keep them constant for all branches in the overall deformable structure. This last constraint could be easily removed, however we believe that our choice is justified by the following reasons:

1. We have noticed in our experiments that optimizing for the decimator and the interpolator independently (imposing the same kernel's length), we obtain two identical filters (up to a scale factor). Although we cannot provide a rigorous theoretical explanation of such this phenomenon, we have found no counterexample so far;

2. The proposed structure is typical of pyramidal implementations like the "gaussian pyramid" [11];

3. The optimization procedure is much faster this way, as only one kernel per branch has to be optimized.

Any filter suitable for pyramidal decomposition is a candidate for $h(x)$, so long as it can be implemented efficientely. This is a fundamental requirement because we store the basis filters' outputs in their decimated versions, for memory parsimony. For each scale and orientation, we interpolate back these signals before recombination[1].

In our experiments, we have used a simple 3-taps raised-cosine filter $h(x) = [a, 1, a]$ for the decimator/interpolator (the optimization of $h(x)$ is discussed in [4].) To interpolate by a factor 2 using polyphase implementation, we only need 0.5 multiplications and 0.5 sums per output sample.

To appreciate the reduction of the computational burden for the recombination from the basis filters' outputs, consider a 1-D deformable filter bank where all R basis filters are implemented in a 2-multirate scheme. For each scale/orientation, we need 0.5 (for interpolation) plus $0.5R$ (for recombination) multiplications per input sample. A comparison with the non-multirate case (where the reconstruction requires R multiplications sums per input sample per scale/orientation) shows that, even if the multirate implementation may require to increase the filter bank's rank to compensate for larger approximation error, it makes for the substantial reduction of the computational cost associated with the recombination.

Our approach requires to first select those 1-D filters that will be realized in a multirate fashion, and then to set the corresponding multirate indices $\{M(r)\}$. A simple design strategy would optimize the filter bank for a non-multirate implementation, and then approximate each basis kernel in a multirate scheme under the guidelines of [8]. Clearly, optimizing the overall system for given multirate structure and for given kernels' sizes represents a better solution. In [4], we describe the extension to the algorithm of [6] to achieve such a goal.

[1] A different strategy consists in storing the interpolated versions of the basis filters' outputs. While no memory reduction is achieved this way, the computational burden is relieved from the repeated interpolation (although now the recombination is performed on the full rate signals). This solution may be suitable when memory constitutes no problem, and larger size interpolators are used.

We are left with the problem of the *a priori* selection of the multirate order and kernel's size for each basis filter. Here is a simple heuristic procedure (we consider the scalable 1-D case for simplicity's sake):

- Design the non-multirate filter bank for given rank R and common kernels' size N;
- For each basis filter $u_r(x)$:
  - For pyramid depth i from 0 to some limit I:
    * Compute the multirate approximations with inner kernel's size equal to N;
    * Compute the conventional filter length $N_c(r, i)$ of the optimized inner kernel;
    * Compute again the multirate approximations with inner kernel's size equal to $N_c(r, i)$;
  - Retain the largest pyramid depth $i(r)$ that gives multirate approximation error below some fixed threshold.

Then, we optimize the overall pyramidal deformable filter bank where we set the depth and the kernel's length of the r-th basis filter equal to $i(r)$ and $N_c(r, i(r))$ respectively. The design algorithm, based on an iterative procedure, is decribed in [4], where we also present experimental results.

3. CONCLUSION

We have described a technique for the efficient implementation of deformable filter banks, based on the use of separable basis filters and multirate implementation. The filter bank is embedded in a pyramidal structure, designed under a novel multirate error criterion that jointly minimizes the approximation error and the aliasing. The computational weight and memory requirement are dramatically reduced with respect to the common steerable/scalable decomposition based on SVD.

We have used the least-squares criterion, which allowed us to attack the design problem in a rigorous formal setting. The choice of other design criteria, which may suit more effectively the image analysis tasks at hand, remains an open problem.

The Matlab software to implement the filter banks described in this paper may be found at http://www.vision.caltech.edu/manduchi/def.tar.Z.

4. REFERENCES

[1] W. Freeman and E. Adelson. The design and use of steerable filters. *IEEE Trans. Pattern Anal. Mach. Intell.*, 13:891–906, 1991.

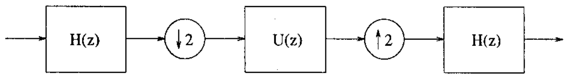

Figure 1: Basic multirate structure used in the pyramidal scheme.

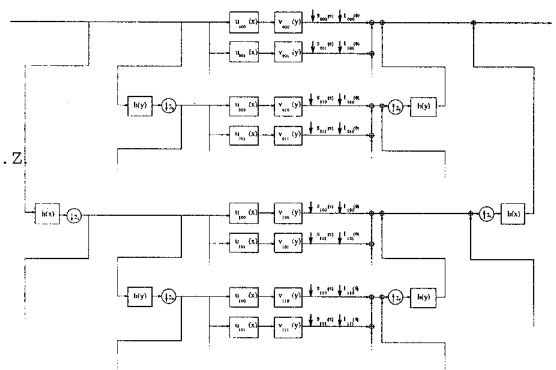

Figure 2: The pyramidal scheme for the 2-D multirate separable scalable/steerable decomposition.

[2] P. Perona. Deformable kernels for early vision. *IEEE Trans. Pattern Anal. Mach. Intell.*, 17(5):488–499, May 1995.

[3] E. Simoncelli, W. Freeman, E. Adelson, and D. Heeger. Shiftable multi–scale transform. *IEEE Trans. Inform. Theory*, 38(2), March 1992.

[4] Perona P. Manduchi, R. and Shy D. Efficient implementation of deformable filter banks. Technical Report CNS-TR-97-04, California Institute of Technology, 1997. Also available at `http://www.vision.caltech.edu/manduchi/deformable.ps.Z`

[5] S. Treitel and J. Shanks. The design of multistage separable planar filters. *IEEE Trans. Geosci. Electron.*, 9(1):10–27, 1971.

[6] R. Manduchi and P. Perona. Pyramidal implementation of deformable kernels. In *Proceedings of the IEEE ICIP'95*, pages 378–381, Washington D.C., 1995.

[7] L.R. Rabiner and R.E. Crochiere. A novel implementation for narrow–band FIR digital filters. *IEEE Trans. Acoust., Speech, Signal Processing*, 23:457–464, October 1975.

[8] R. Manduchi and P. Perona. Least–squares multirate FIR filters. *Electronics Letters*, 32(8):726–727, April 1996.

[9] P.J. Burt. Fast filter transforms for image processing. *Computer Graphics and Image Processing*, 16:20–51, 1981.

[10] Y. Neuvo, D. Cheng-Yu, and S.K. Mitra. Interpolated finite impulse response filters. *IEEE Trans. Acoust., Speech, Signal Processing*, 32:563–570, June 1984.

[11] P.J. Burt and E.H. Adelson. The Laplacian pyramid as a compact image code. *IEEE Trans. Commun.*, 31:532–540, April 1983.