CAR-TR-730 DACA76-92-C-0009
CS-TR-3334 August 1994

# Fast Correlation Matching in Large (Edge) Image Databases

D.M. Gavrila
L.S. Davis

Computer Vision Laboratory
Center for Automation Research
University of Maryland
College Park, MD 20742-3275

**Abstract**

Correlation-based matching methods are known to be very expensive when used on large image databases. In this paper, we will examine ways of speeding up correlation matching by phase-coded filtering. Phase coded filtering is a technique to combine multiple patterns in one filter by assigning complex weights of unit magnitude to the individual patterns and summing them up in a composite filter. Several of the proposed composite filters are based on this idea, such as the Circular Harmonic Component (CHC) filters and the Linear Phase Coefficient Composite (LPCC) filters.

We will consider the LPCC(1) filter in isolation and examine ways to improve its performance by assigning the complex weights to the individual patterns in a non-random manner so as to maximize the SNR of the filter w.r.t. the individual patterns. In experiments on a database of 100 to 1000 edge images from the aerial domain we examine the trade-off between the speed-up (the number of patterns combined in a filter) and unreliability (the number of resulting false matches) of the composite filter. Results indicate that for binary patterns with point densities of about 0.05 we can safely combine more than 20 patterns in the optimized LPCC(1) filter, which represents a speed-up of an order of a magnitude over the brute force approach of matching the individual patterns.

**Keywords**: image matching, correlation, composite filters, phase coding

The support of the Advanced Research Projects Agency (ARPA Order No. 8459) and the U.S. Army Topographic Engineering Center under Contract DACA76-92-C-0009 is gratefully acknowledged, as is the help of Sandy German in preparing this paper.

## 1 Introduction

An important function of current and future image databases is content-based retrieval. This is the capability of retrieving (parts of) images based solely on their image properties, without using any external textual or numeric labels attached to them. Not only does this avoid a time-consuming human involvement in the labeling process, but it also allows queries based on spatial configurations of patterns, which cannot easily be described in words.

Image matching in very large databases poses challenges in addition to the limitations of the particular matching method used. A brute-force method which matches a query pattern exhaustively against all possible image locations can be quite successful on a small number of images, but does not scale well when used on very large image databases. One way to achieve scalability is to apply inexpensive matching methods first, using statistics of different features in the images, in order to select from the vast number of images a subset likely to contain the desired match (see [8], for example). On this subset, more elaborate matching methods can be used, such as graph-matching or correlation-based methods.

Another way to improve on the brute-force approach is to speed up the computations associated with the elaborate matching methods themselves. We will take this approach here and consider ways of speeding up correlation matching by combining multiple patterns into one filter. Matching is then done in "parallel", using the composite filter instead of the individual patterns, resulting in a constant speed-up factor. Such a technique can be used, for example, to detect patterns invariant under a certain transformation (e.g. rotation, scale, projection), by storing different transformations of a reference pattern explicitly in the filter. The target patterns can also be unrelated; in other words, the technique is also applicable in cases where we match a given pattern against a collection of unrelated images.

An interesting technique for building composite filters is phase coding [5, 7]. It assigns to each target pattern a complex weight of unit magnitude (a phase). The composite filter is simply the sum of the patterns weighted by their phases. The underlying principle of phase coding is that if the phases are distributed uniformly over the unit circle, noise will excite all target patterns equally and the complex terms of the filter will cancel out. If one of the target patterns is presented to the filter, the response of the filter will be biased in the direction of the phase assigned to that pattern. This case can thus be detected by inspecting the magnitude of the filter response. The phase of the response vector will be indicative of the target pattern responsible for the match.

In this paper we will consider a particular phase coded filter called LPCC(1) [5]. We derive the SNR for this filter, based on a noise model assuming binary patterns. Unlike previous approaches to phase coding, we assign the phases to the target patterns so as to maximize the SNR of the composite filter. In particular, the phase assignment is such that similar patterns are assigned similar phases in order to avoid cancellation effects in the filter response.

The organization of this paper is as follows. Section 2 covers the LPCC(1) filter and related work. Section 3 deals with the definition and derivation of the SNR. In Section 4 we discuss the phase assignment problem. Section 5 describes the experiments; the conclusions are contained in Section 6.

## 2 The LPCC(1) Filter and Related Work

Several techniques have been proposed to recognize patterns in images under various transformations such as rotation and scale [1, 3, 9]. Here we will restrict ourselves to techniques which are based on phase coding, such as the Circular Harmonic Component (CHC) filters [7] and the Linear Phase Coefficient Composite (LPCC) filters [4–6].

The CHC filters [7] are designed to handle rotational invariance by converting a pattern from Cartesian to polar coordinates. In its polar representation, a pattern $s(r, \phi)$ is periodic in $\phi$ with period $2\pi$. Thus it can be expanded into the Fourier series in $\phi$ by

$$s(r, \phi) \;=\; \sum_{M=-\infty}^{\infty} s_M(r) e^{jM\phi} \tag{1}$$

where

$$s_M(r) = \frac{1}{2\pi} \int_0^{2\pi} s(r, \phi) e^{-jM\phi} d\phi \tag{2}$$

The function $s_M(r)$ is known as the circular harmonic component of order $M$. The CHC filter of order $M$ is given by

$$f_{CHC,M}(r, \phi) = s_M(r) e^{jM\phi} \tag{3}$$

The LPCC filter family [5] is defined by adding target patterns weighted by complex numbers of unit magnitude and equi-distributed phase:

$$f_{LPCC,M}(x, y) = \sum_k \mathbf{s}_k \, \mathbf{e}^{-j2\pi kM/K} \tag{4}$$

If the target patterns are (in-plane) rotations of a reference pattern, then the CHC and LPCC filters are equivalent [5]. For that case and close approximations, Hassebrook et al. [5] derive a general signal-to-noise model which enables them to linearly combine $N$ LPCC filters into a filter bank in an optimal way. Which LPCC filters to choose from and how many remains an open problem.

Unlike the CHC filters, which were specially designed to handle rotational invariance, the LPCC filters can be used in a more flexible way. They can be used to combine patterns under different transformations, such as scale, or to combine unrelated patterns together. In this paper, we examine the performance of one of the LPCC filters in isolation, the LPCC(1) filter, following the work of Carlotto [2]. The difference between [2] and our work is twofold. First, we will assign the phases to patterns in a non-random way to maximize the SNR of the filter. Secondly, the focus of this paper will be on extensive experiments to examine the performance of the filter on a large data set.

The LPCC(1) filter is built by assigning target patterns $\mathbf{s}_i, i = 1, \ldots, K$ complex weights (**phases**) of unit magnitude $\phi_i, i = 1, \ldots, K$. The phases are assumed **equi-distributed**, i.e. there exists a permutation $\Pi$ over the indices $i$ such that

$$\phi_i \;=\; \Pi(i)\,\frac{2\pi}{K} \tag{5}$$

Each permutation $\Pi$ represents a valid **phase assignment**. We will call the phase assignment *successive* if $\Pi$ is the identity mapping. The LPCC(1) filter is given by the sum of the weighted patterns

$$\mathbf{f} \;=\; \sum_k \mathbf{s}_k\, \mathbf{e}^{j\phi_k} \tag{6}$$

The response of the filter $\mathbf{f}$ to a pattern $\mathbf{q}$ is

$$\mathbf{r} = \mathbf{q}^T\,\mathbf{f} = \sum_k \mathbf{q}^T\,\mathbf{s}_k\,\mathbf{e}^{j\phi_k} \tag{7}$$

To see how the filter works, assume that $\mathbf{q} = \mathbf{n}$, a noise vector. Then

$$\mathbf{r} \;=\; \mathbf{n}^T\,\mathbf{f} \;=\; \sum_k \mathbf{n}^T\,\mathbf{s}_k\,\mathbf{e}^{j\phi_k} \tag{8}$$

Under the assumption that noise will excite all target patterns equally strong, the terms $\mathbf{n}^T\,\mathbf{s}_k$ will be constant. Their weighted sum will cancel out or be near $\mathbf{0}$, given equi-distributed phases.

Now assume on the other hand that $\mathbf{q} = \mathbf{s}_i$. Then

$$\mathbf{r} \;=\; \mathbf{s}_i^T\,\mathbf{f} \;=\; \mathbf{s}_i^T\,\mathbf{s}_i\,\mathbf{e}^{j\phi_k} + \sum_{k \neq i} \mathbf{s}_i^T\,\mathbf{s}_k\,\mathbf{e}^{j\phi_k} \tag{9}$$

The response is the sum of a vector of length $||\mathbf{s}_i||^2$ and phase $\phi_k$, and an interference term. Again, under reasonable assumptions, the interference term is small compared to the signal term. For example, given $\mathbf{s}_i^T\,\mathbf{s}_i = p$ and $\mathbf{s}_i^T\,\mathbf{s}_k = \alpha p$, $i \neq k$, $0 \leq \alpha \leq 1$ then

$$\mathbf{r} \;=\; p\mathbf{e}^{j\phi_k} - \alpha p\,\mathbf{e}^{j\phi_k} + \sum_{k,i} \alpha p\,\mathbf{e}^{j\phi_k} \;=\; (1-\alpha)p\,\mathbf{e}^{j\phi_k} \tag{10}$$

This means that matching a pattern $\mathbf{q}$ with one "stored" in the composite filter can be detected by examining the radius and phase of the response vector $\mathbf{r}$.

Compare the above phase-coded matching approach with a simple correlation match approach. There we match pattern $\mathbf{q}$ with each target pattern $\mathbf{s}_i, i = 1, \ldots, K$ separately

$$r_k \;=\; \mathbf{q}^T\,\mathbf{s}_k \tag{11}$$

As a general matching procedure, phase-coded matching has the same limitations as simple correlation matching. Its advantage over simple correlation is speed. Given unit cost for performing a simple correlation (whether it is performed in the spatial or in the frequency domain), the cost of a simple correlation match is $K$, while the cost is 2 for the phase-coded approach (one correlation for the real part and one for the imaginary part of eq. (7)). Thus the potential speed-up factor is $K/2$.

In practice, the phase coded filter is best used as an "information filter" to select promising solutions to the correlation match. The technique can lead to false positive and false negative matches. A false positive match arises when pattern $\mathbf{q}$ matches some of the target patterns only partially, but these target patterns happen to have similar phases, so their contributions all add up, leading to a large response vector $\mathbf{r}$. See Figure 1a. Shown is the simple correlation between a non-target query pattern and the target patterns of a filter. The horizontal axis corresponds to the different target patterns (denoted by their assigned phases) and the vertical axis gives the result of the simple correlation $r_k$. Note that although the query pattern doesn't match any of the target patterns in particular, it matches the target patterns around phase $-90$ degrees slightly more than the others. This asymmetry leads to a relative large response vector, shown by the vertical line. This case of a false positive match can be dealt with by verifying the match locations proposed by phase-coded matching with simple matching. There is no need to match all $K$ target patterns at candidate locations; the phase of the response vector (eq. (7)) can be used to limit their number. If at the candidate location there exists a good match with one of the target patterns, the phase of the response vector will not deviate much from the phase assigned to the matching target pattern.

Thus, one can seek matching target patterns among the patterns with phases within a range of the phase of the response vector, rejecting the match if the result of simple correlations is below a threshold. This will eliminate matches such as in Figure 1a.

A false negative match is more serious and arises when pattern $\mathbf{q}$ matches one of the target patterns, say $\mathbf{s}_k$, but it happens that there are one or more target patterns similar to $\mathbf{s}_k$ which have been assigned opposite phases. Pattern $\mathbf{q}$ will then match target patterns with opposing phases, and their contributions in the response vector $\mathbf{r}$ will be attenuated, or even cancelled. See Figure 1b. Shown is the simple correlation between a query pattern and the target patterns of the same filter as in 1a. In this case, the query pattern is identical to the target pattern which was assigned phase 0 degrees. Although Figure 1b shows a relative high correlation match at phase 0 degrees, the same applies for phase 180 degrees. This results in the described cancellation effect which has the effect that the "true" solution in 1b is ranked lower than the "false" solution in 1a, because of its smaller response vector.

This raises the issue of which types of patterns can be used in the phase-coded approach. For applications involving non-binary target patterns with values near an average (non-zero) value, the interference effects will be major. Although the phase-coded approach still can be applied in that case (possibly with an SNR model, as in [5]), it will not produce experimental results as good as when used with binary patterns with relative low point-density (say 0.05). All previous experiments [2, 5, 7] use low point-density binary data for this reason, and so will we. In the next section, we will derive a SNR for the phase-coded filter based on binary patterns. Interference effects are minimized by a phase assignment which maximizes the SNR; see Section 4.

The domain which we will consider in this paper is image matching. One way to encode the information present in images is with a pattern vector. In this paper, we represent an image by a vector of pixel values; the patterns are binary and obtained from the edge map of an image. Of course, pattern vectors can also be used to encode features in an image, such as the presence of straight lines at various orientations. In any case, the operation of interest here is matching templates with images by correlation. There are three different ways to use the phase-coded approach (i.e. LPCC(1) filter) [2]:

- Single Image Multiple Template (SIMT) case

- Multiple Image Single Template (MIST) case

- Multiple Image Multiple Template (MIMT) case

The SIMT case can be used to deal with the important case of detecting transformation-invariant patterns (w.r.t. rotation, scale, projection) in an image. Different transformed versions of a reference template are generated explicitly and combined into one or more complex filters. Matching is then done with the transformation-invariant filter. The MIST case is useful in the context of image databases, where we want to take advantage of the fact that we know the images in advance; we can combine images together off-line to speed up the matching with a query template on-line. The MIMT case is a mixture of the SIMT and MIST case. Since it does not introduce any new issues compared to the previous cases, we do not consider it in this paper.

## 3 SNR

Let $D$ be the dimensionality of the pattern space. Consider first the case in which the binary patterns $\mathbf{s}_i$ $(i = 1, \ldots, K)$ are deterministic and the noise $\mathbf{n}$ is uniformly distributed. More precisely, $\mathbf{n}$ consists of $D$ independent binary random variables $X$ such that $p(X = 1) = \rho_n$ and $p(X = 0) = 1 - \rho_n$. Let $\mathbf{S}_i$ and $\mathbf{N}$ be random variables denoting the filter response to signal $\mathbf{s}_i$ and noise $\mathbf{n}$, resp. The following statistics can be derived:

$$E[\mathbf{S}_i] = \mathbf{S}_i = \sum_k \mathbf{s}_i^T \, \mathbf{s}_k \, \mathbf{e}^{j\,\phi_k} \tag{12}$$

$$\begin{aligned} E[\mathbf{N}] &= \sum_k E[\mathbf{n}^T \mathbf{s}_k] \, \mathbf{e}^{j\,\phi_k} \\ &= \rho_n \sum_k \mathbf{s}_k^T \mathbf{s}_k \mathbf{e}^{j\phi_k} \end{aligned} \tag{13}$$

$$\begin{aligned} E[\mathbf{S}_i \mathbf{S}_i^*] &= E[\mathbf{S}_i] \, E[\mathbf{S}_i]^* \\ &= \sum_k \sum_l \mathbf{s}_i^T \mathbf{s}_k \, \mathbf{s}_i^T \mathbf{s}_l \, \mathbf{e}^{j(\phi_k - \phi_l)} \end{aligned} \tag{14}$$

$$\begin{aligned} E[\mathbf{N}\mathbf{N}^*] &= \sum_k \sum_l E[\mathbf{n}^T \mathbf{s}_k \, \mathbf{n}^T \mathbf{s}_l] \, \mathbf{e}^{j(\phi_k - \phi_l)} \\ &= \sum_k \sum_l \Big( \sum_i \sum_j s_k^i s_l^j E[n^i n^j] \Big) \, \mathbf{e}^{j(\phi_k - \phi_l)} \\ &= \sum_k \sum_l \Big( \rho_n \sum_i s_k^i s_l^i + \rho_n^2 \sum_i \sum_{j \neq i} s_k^i s_l^j \Big) \, \mathbf{e}^{j(\phi_k - \phi_l)} \\ &= \sum_k \sum_l \big( \rho_n \, \mathbf{s}_k^T \mathbf{s}_l + \rho_n^2 (\mathbf{s}_k^T \mathbf{s}_k \, \mathbf{s}_l^T \mathbf{s}_l - \mathbf{s}_k^T \mathbf{s}_l) \big) \, \mathbf{e}^{j(\phi_k - \phi_l)} \\ &= \sum_k \sum_l \big( (\rho_n - \rho_n^2) \, \mathbf{s}_k^T \mathbf{s}_l + \rho_n^2 \, \mathbf{s}_k^T \mathbf{s}_k \, \mathbf{s}_l^T \mathbf{s}_l \big) \, \mathbf{e}^{j(\phi_k - \phi_l)} \end{aligned} \tag{15}$$

$$\begin{aligned} Var[\mathbf{S}_i] &= E[\mathbf{S}_i \mathbf{S}_i^*] - E[\mathbf{S}_i] \, E[\mathbf{S}_i]^* \\ &= 0 \end{aligned} \tag{16}$$

$$
\begin{aligned}
Var[\mathbf{N}] &= E[\mathbf{N}\mathbf{N}^*] - E[\mathbf{N}]E[\mathbf{N}]^* \\
&= (\rho_n - \rho_n^2) \sum_k \sum_l \mathbf{s}_k^T \mathbf{s}_l \, \mathbf{e}^{j(\phi_k - \phi_l)} \qquad (17)
\end{aligned}
$$

Consider now the case in which both the signal patterns and the noise are uniformly distributed. Assume all signals have the same distribution ($\mathbf{S}_i = \mathbf{S}$). Let $\rho_s$ and $\rho_n$ denote the corresponding parameters of the $D$ independent binary random variables making up the signal $\mathbf{s}$ and the noise $\mathbf{n}$. Assume that the phase assignment is equi-distributed. Then the following can be derived, similarly to [2]:

$$
\begin{aligned}
E[\mathbf{S}] &= D(\rho_s - \rho_s^2)\, \mathbf{e}^{j\phi_i} && (18) \\
E[\mathbf{N}] &= \mathbf{0} && (19) \\
E[\mathbf{S}\mathbf{S}^*] &= D(\rho_s - \rho_s^2)\, (\rho_s^2(1-D) + \rho_s(D+K-3) + 1) && (20) \\
E[\mathbf{N}\mathbf{N}^*] &= D(\rho_s - \rho_s^2)\, K \, \rho_n && (21) \\
Var[\mathbf{S}] &= D(\rho_s - \rho_s^2)\, (\rho_s^2 + \rho_s(K-3) + 1) && (22) \\
Var[\mathbf{N}] &= D(\rho_s - \rho_s^2)\, K \, \rho_n && (23)
\end{aligned}
$$

We would like to have a measure of the performance of a certain filter $\mathbf{f}$. Define the signal to noise ratio, $SNR(i)$, of filter $\mathbf{f}$ w.r.t. component signal $\mathbf{s}_i$ as

$$
SNR(i) = \frac{E[\mathbf{S}_i\mathbf{S}_i^*]}{E[\mathbf{N}\mathbf{N}^*]} \qquad (24)
$$

The overall SNR of filter $\mathbf{f}$ can be based on different statistics of the distribution of $SNR(i)$ for $i = 1, \ldots, K$, such as the minimum value or average value. We chose

$$
SNR = \left( \frac{1}{K} * \sum_i SNR(i)^{1/2} \right)^2 \qquad (25)
$$

For the uniform noise model, using eqs. (20), (21) and (24), we obtain a simplified expression for the signal-to-noise ratio

$$
SNR = \frac{\rho_s^2(1-D) + \rho_s(D+K-3) + 1}{K \, \rho_n} \qquad (26)
$$

which for typical values of $D$, $K$ and $\rho_s$ (i.e. $\rho_s^2 \ll \rho_s$, $K \ll D$), leads to the approximation

$$
SNR = \frac{\rho_s D}{K \, \rho_n} \qquad (27)
$$

## 4 Phase Assignment

In the previous section we defined the SNR for the phase-coded filter. In the uniform signal/noise model the SNR does not depend on the phase assignment of the patterns, as long as the phases are equi-distributed. This is because the similarity between two different target patterns $\mathbf{s}_i$ and $\mathbf{s}_j$, i.e. their inner-product $\mathbf{s}_i^T \mathbf{s}_j$, is taken to be constant. In practice, the similarity between two target patterns in the filter can vary and the chosen phase assignment does influence the SNR, as can be seen from eqs. (14), (15) and (24). Consider the following worst case scenario: we have $K$ target patterns which are rotations of a reference pattern which is point-symmetric around the origin, using the phase assignment such that the rotated versions are assigned successive phases. A pattern and its 180 degree rotated version will then be identical, and will be assigned opposite phase. This will result in pairwise cancellation of the terms of the filter response.

Of course, the case where the filter consists of rotations of a reference pattern around its point of symmetry occurs rarely, and could be detected beforehand. But more generally, cancellation effects occur in the filter when two similar target patterns are assigned opposite phases. [5] deals with this case by using multiple LPCC filters with different phase assignments, increasing the matching cost. In this paper, we examine ways to improve the phase assignment for one filter, the LPCC(1) filter.

In the absence of an analytical method to find the optimal phase assignment of a set of $K$ target patterns in terms of the maximum SNF, we turn to heuristics to develop a sub-optimal phase assignment. Note that an exhaustive search through the space of all permutations is prohibitive for typical values of $K > 10$, given its size of $(K-1)!$ excluding cyclical symmetry. One could use well-known AI search techniques to prune this space. A more efficient approach is a hierarchical one:

1. group the $K$ patterns into $M$ groups ($4 \leq M \leq K$), each group containing patterns which are assigned successive phases.

2. assign the $M$ groups over the unit circle

In step 1 we use similarity between patterns as the criterion for grouping. Similarity between two patterns $\mathbf{s}_i$ and $\mathbf{s}_j$ is defined in terms of their scaled inner-product:

$$\text{Dist}(\mathbf{s}_i, \mathbf{s}_j) \;=\; \frac{\mathbf{s}_i^T \mathbf{s}_j}{\min(||\mathbf{s}_i||, ||\mathbf{s}_j||)} \tag{28}$$

The idea behind this "similar patterns, similar phases" heuristic is to avoid cancellation effects in the filter and to increase the denominator of the $SNR(i)$ (eq. (24)). An implementation for step 1 is based on a simple iterative clustering algorithm which partitions the $K$ patterns initially into $M$ groups, and shifts patterns back and forth between groups in order to minimize the sum of the intra-cluster distances between pairs of patterns. Step 2 can now be done exhaustively over all $(M-1)!$ different distributions of the $M$ groups over the unit circle, for small values of $M = 4, 5, 6$. We retain the phase assignment with the maximum SNR.

So far we have dealt with the problem of how to distribute $K$ patterns within a complex filter. In typical applications we are given $N$ patterns and some decisions have to be made concerning the number of complex filters to represent the pattern set and how to perform the partitioning. One simple way to partition the pattern set, without computing the $N^2$ cross-correlations, is to rank the patterns according to decreasing point density and assign successive patterns to the same filter, until the SNR (based on the model of uniform distribution of signal and noise) falls below a certain threshold. At that point a new filter is assigned for the remaining pattern set. This has the benefit that it avoids grouping a pattern having small point density with one having large point density, where the latter can dominate the former. If there is some a-priori information available about the patterns, for example that they are rotated versions of one another, other heuristics can be used. In general, more effort can be put into finding a good phase assignment for target patterns if this process can be done off-line, prior to matching.

## 5 Experiments

We examined the performance of the phase-coded filter with and without the phase assignment algorithm. Two cases were considered: the SIMT (Single Image Multiple Template) and the MIST (Multiple Image Single Template) case. For the SIMT case, we considered the case of combining different rotations of a reference template into the phase coded filter and then matching it with the image database. The MIST case involved combining different (unrelated) images from the database together and matching them with a reference template. Ten subimages, containing salient patterns P0–P9, were extracted from the database to represent the reference templates. See Figure 4. For both SIMT and MIST experiments, we examined the trade-off between the speed-up (the number of target patterns) and unreliability (the number of resulting false matches) of the composite filter. To do this, we ranked the match locations produced by the phase-coded filter in decreasing order of size of the response vector. Then it was determined where the correct location of the reference template was ranked. This was done for varying number of target patterns $K$ in a filter.

The database consisted of edge images of size $128 \times 128$ from the aerial domain. For the MIST case the size of the database was 1000, while for the SIMT case it was 100. The average point density of the edge images in the database was about 0.05. See Figures 2 and 3 for 16 images of the database and the corresponding edge maps. Correlation between template and image was implemented in the frequency domain using the well known equivalence between correlation in the spatial domain and multiplication in the frequency domain. That is, if $f(x, y)$, $g(x, y)$ and $z(x, y)$ denote functions in the spatial domain where

$$z(x, y) = f(x, y) \oplus g(x, y) \tag{29}$$

and where $\oplus$ denotes the correlation operator, then

$$Z(u, v) = F(u, v) \times G(x, y)^* \tag{30}$$

where $F(u, v)$, $G(u, v)$ and $Z(u, v)$ are the Fourier transforms of $f(x, y)$, $g(x, y)$, and $z(x, y)$. Correlation thus involves transforming the template and the image to the frequency domain, performing a multiplication, and transforming the result back to the spatial domain. There are three Fourier transforms to be performed, each with complexity $O(n^2 \log n)$.

### 5.1 Single Image Multiple Template Experiments

The effects of the proposed phase assignment on the SNR of the phase-coded filter can be seen in Figures 5 and 6. In Figure 5, 45 rotations of the reference pattern P5 were used to construct the filter. Figure 5a shows the response of the filter (eq. (12)) for the 45 target patterns and for 45 patterns at intermediate rotations between the target patterns. Figure 5b shows the expected response of the filter to uniform noise with densitity 0.1 ($E[\mathbf{NN}^*]^{1/2}$, see eq. (15)). Figure 5c shows the SNR w.r.t. the target patterns. The patterns are represented by their assigned phases. The gray line indicates the un-optimized phase assignment case, where successive rotations of a pattern are assigned successive phases. The black line indicates the optimized phase assignment. Figure 6 shows the SNR for the patterns P0–P9. As can be expected, the SNR is higher for the target patterns of the filter than for patterns of intermediate rotations. This explains the oscillating nature of the curves. For patterns P7 and P8 the phase assignment was identical to the successive phase assignment, resulting in the same SNR. For the other cases, clear improvements can be seen when using the optimized phase assignment.

Now we turn to the results obtained on the dataset of 100 images. Figures 7a and 7b list the rank of the correct location in the ranking of all potential match locations by the phase-coded approach. The results are given for different values of $K$ and for both the successive and optimized phase assignment. Given an average template size of $30 \times 30$, there are roughly $100 \times 98^2 = 9.6 \times 10^5$ potential match locations; in the experiments we retained the top $1 \times 10^4$ (i.e. top 1%). From Figure 7 it can be seen that matching with the optimized filter produces some improvement w.r.t. un-optimized version, but both versions of the phase-coded filter actually perform very well, ranking the correct solution consistently in the top 1%.

To evaluate the performance of the phase-coded filter in the case that the desired pattern only partially matches one of the target patterns, we performed the following experiment. For a fixed $K = 45$, we matched all $K$ individual target patterns with the database by simple correlation and retained the overall best 10 solutions. These solutions all matched more than 50% of some target pattern. Then it was determined where these solutions rank in the phase-coded ranking. See Figure 8 for the results for the successive and optimized phase assignment. A clear improvement is achieved by the optimized phase-coded filter. If, for example, we regard all solutions not within the top 50000 (i.e. top 5%) as missed, with the un-optimized filter we miss 24 out of 100 solutions compared to only 4 out of 100 for the optimized filter.

### 5.2 Multiple Image Single Template Experiments

The database of 1000 images was compressed into $M$ filters using the procedure described at the end of Section 4. The parameters used were $D = 30^2$ and SNR = 150, 100, 50 resulting in M = 74, 50, 26, or equivalently $K_{\text{avg}} = 14, 20, 38$ resp. There are roughly $M \times 98^2$ locations to consider with the phase-coded filter. Figure 9 gives the results for the un-optimized and optimized phase coded filters. In this case, the gain in using the optimized phase assignment is not evident. This is because the images are unrelated to each other so the inner product $\mathbf{s}_i^T \mathbf{s}_j$ does not vary much among the images assigned to a filter. Thus, unlike the situation in the previous section, we cannot expect much improvement over random phase assignment. Observe that this is different from the case in which the images represent a video sequence, where images grouped in a filter are related. The results for the optimized phase assignment turn out to be even slightly worse for some patterns. This can occur because the phase assignment is based on the global similarity of the two images.

It is possible that two images are assigned opposite phases, but are locally similar to each other. This is what happened in the case of P4 with $K = 38$. But both versions of the filter perform very well in retrieving the correct location of the reference template. With one exception (for $K = 38$), all correct locations of the patterns P0–P9 are found in the top 5%.

## 6 Conclusion

We have performed experiments with the LPCC(1) filter and an optimized version where the phase assignment maximizes a defined overall SNR. Improvements are clear for the case in which the filter consists of rotated versions of a reference pattern. For unrelated patterns, the proposed phase assignment can still be useful to avoid occasional cancellation effects in the filter. The phase coded technique proved successful in the case of binary patterns with low point densities (0.05). For this case, more than 20 patterns can be safely combined into a filter, resulting in a speed-up of an order of magnitude over the brute force approach of matching the individual patterns.

## References


[1] J. Ben-Arie and K.R. Rao: "A novel approach for template matching by nonorthogonal image expansion," *IEEE Trans. on Circuits and Systems for Video Technology*, Vol. 3, pp. 71–83, 1993.

[2] M.J. Carlotto: "Parallel pattern recognition using phase-coded filters," submitted to *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 1993.

[3] D. Casasent and D. Psaltis, "New optical transforms for pattern recognition," *Proc. IEEE*, Vol. 65, No. 1, 1977.

[4] L. Hassebrook, B.V.K. Vijaya Kumar and L. Hostetler: "Linear phase coefficient composite filter banks for optical pattern recognition," *Proc. SPIE*, Vol. 1053, pp. 218–226, 1989.

[5] L. Hassebrook, B.V.K. Vijaya Kumar and L. Hostetler: "Linear phase coefficient composite filter banks for distortion invariant optical pattern recognition," *Optical Engineering*, Vol. 29, pp. 1033–1043, 1990.

[6] L. Hassebrook, M. Rahmati, and B.V.K. Vijaya Kumar: "Hybrid composite filter banks for distortion invariant optical pattern recognition," *Optical Engineering*, Vol. 31, pp. 923–933, 1992.

[7] Y. Hsu, H.H. Arsenault and G. April: "Rotation-invariant digital pattern recognition using circular harmonic expansion," *Applied Optics*, Vol. 21, pp. 4012–4015, 1982.

[8] W. Niblack et al.: "The QBIC project: Querying images by content using color, texture and shape, " *SPIE Int'l. Symp. on Electronic Imaging Science and Technology, Conf. 1908, Storage and Retrieval for Image and Video Databases*, 1993.

[9] A. Rosenfeld and A. Kak: *Digital Picture Processing*, Academic Press, 1982.

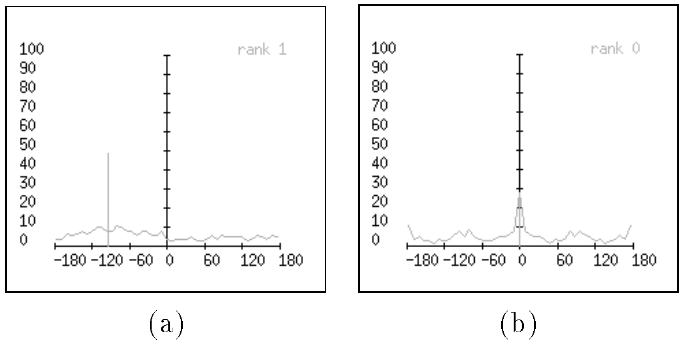


Figure 1: (a) False positive (b) False negative

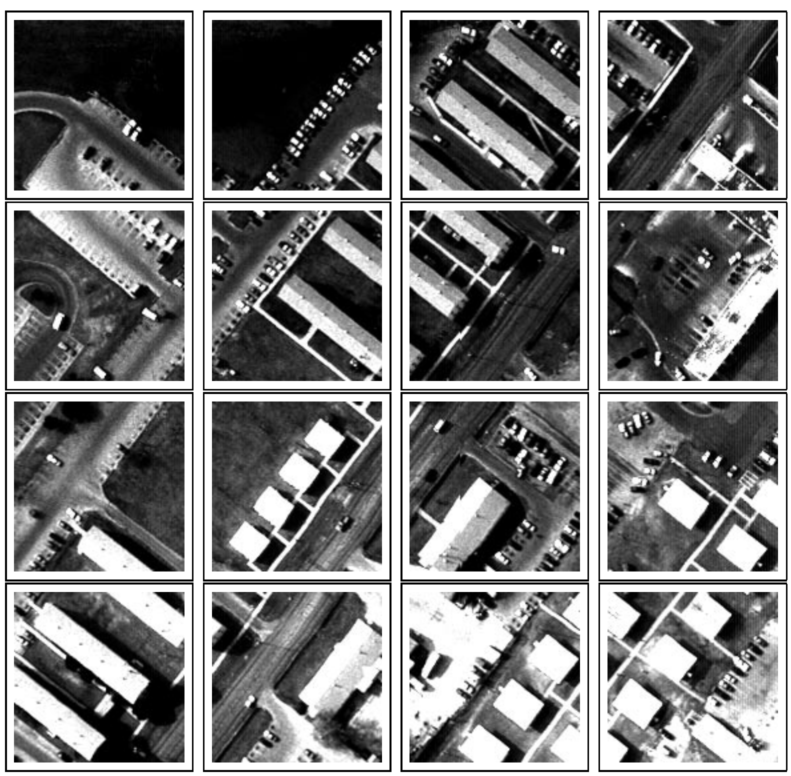

Figure 2: 16 images from the database

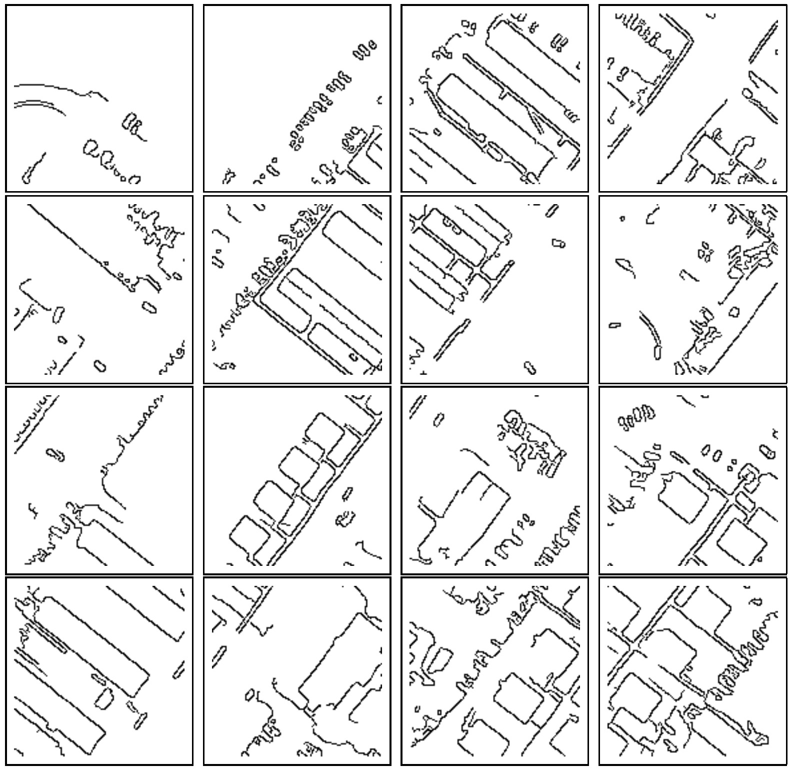

Figure 3: 16 edge images from the database

(a) P0 (b) P1 (c) P2 (d) P3 (e) P4

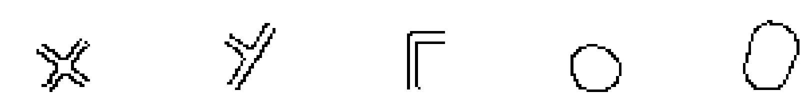

(f) P5 (g) P6 (h) P7 (i) P8 (j) P9

Figure 4: The test patterns P0–P9

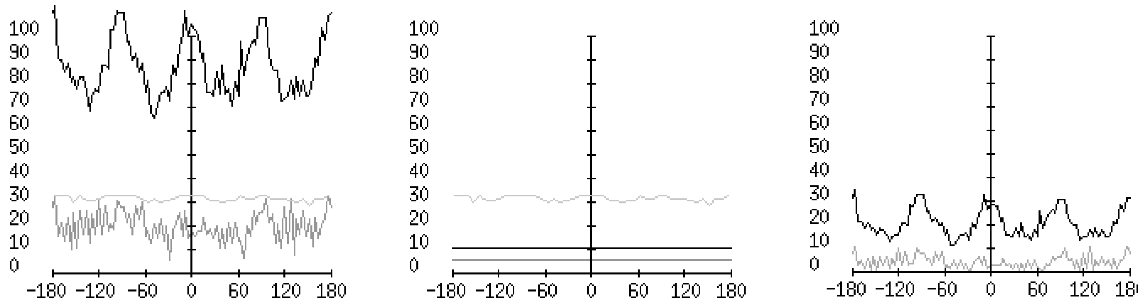


(a) (b) (c)

Figure 5: Pattern P5: (a) Signal response (b) Noise response (c) SNR

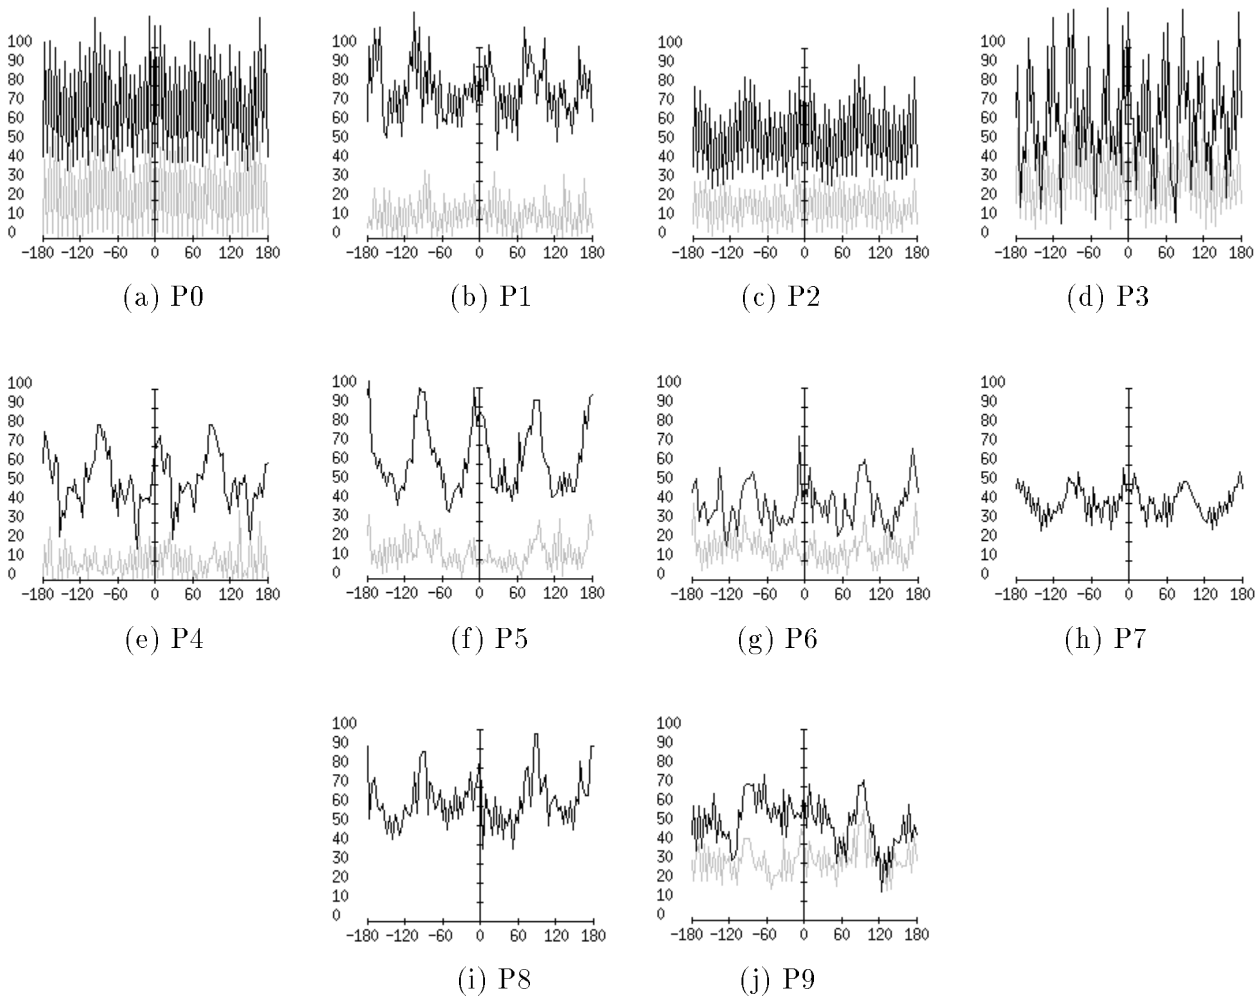


Figure 6: SNR for test patterns P0–P9, successive vs. optimized phase assignment

| | K = 15 | K = 45 | K = 75 |
|---|---|---|---|
| P0 | 1 | 1 | 35 |
| P1 | 1 | 51 | 40 |
| P2 | 1 | 1 | 20 |
| P3 | 1 | 1 | 1 |
| P4 | 1 | 1 | 10000+ |
| P5 | 1 | 3449 | 55 |
| P6 | 1 | 7347 | 4775 |
| P7 | 1 | 43 | 36 |
| P8 | 1 | 1 | 1 |
| P9 | 1 | 2 | 1 |

(a)

| | K = 15 | K = 45 | K = 75 |
|---|---|---|---|
| P0 | 1 | 1 | 1 |
| P1 | 1 | 19 | 45 |
| P2 | 1 | 201 | 1199 |
| P3 | 1 | 1 | 1 |
| P4 | 1 | 2 | 3 |
| P5 | 1 | 4 | 2 |
| P6 | 1 | 8 | 4 |
| P7 | 1 | 43 | 36 |
| P8 | 1 | 1 | 1 |
| P9 | 1 | 1 | 5 |

(b)

Figure 7: SIMT. Effect of $K$ on rank best solution: successive (a) vs. optimized (b) phase assignment

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---|---|---|---|---|---|---|---|---|---|---|
| PO | 1 | 22 | 6429 | 160 | 726 | 173 | 350 | 974 | 5228 | 5884 |
| P1 | 51 | 50000+ | 131 | 50000+ | 11 | 118 | 1811 | 1472 | 57 | 50000+ |
| P2 | 1 | 6 | 70 | 4776 | 5424 | 50000+ | 2235 | 4414 | 291 | 2577 |
| P3 | 1 | 119 | 121 | 26471 | 50000+ | 14 | 50000+ | 695 | 269 | 50000+ |
| P4 | 1 | 3117 | 315 | 155 | 160 | 50000+ | 50000+ | 33737 | 33144 | 1294 |
| P5 | 3449 | 87 | 451 | 43 | 50000+ | 50000+ | 50000+ | 50000+ | 822 | 3830 |
| P6 | 7347 | 50000+ | 50000+ | 50000+ | 50000+ | 50000+ | 50000+ | 29763 | 50000+ | 50000+ |
| P7 | 43 | 2242 | 1090 | 3283 | 12315 | 168 | 27234 | 7058 | 11970 | 5385 |
| P8 | 1 | 65 | 14 | 50000+ | 27 | 9 | 3 | 10 | 17 | 2195 |
| P9 | 1 | 2 | 3 | 11 | 77 | 33 | 50000+ | 50000+ | 2230 | 3550 |

(a)

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---|---|---|---|---|---|---|---|---|---|---|
| P0 | 1 | 6 | 7 | 4 | 2 | 5 | 9 | 8 | 15 | 3 |
| P1 | 19 | 1 | 14 | 50 | 40 | 107 | 438 | 12013 | 440 | 5 |
| P2 | 201 | 472 | 719 | 1064 | 7055 | 13756 | 6590 | 898 | 685 | 2177 |
| P3 | 1 | 8 | 2 | 7 | 3 | 26 | 10 | 12 | 31 | 4 |
| P4 | 2 | 1 | 40 | 19 | 8131 | 159 | 8657 | 6996 | 8903 | 50000+ |
| P5 | 4 | 2 | 3 | 1 | 20 | 6 | 26 | 184 | 44 | 37 |
| P6 | 8 | 2056 | 2936 | 20 | 39 | 42 | 9 | 6 | 83 | 2 |
| P7 | 43 | 2242 | 1090 | 3283 | 12315 | 168 | 27234 | 7058 | 11970 | 5385 |
| P8 | 1 | 65 | 14 | 50000+ | 27 | 9 | 3 | 10 | 17 | 2195 |
| P9 | 1 | 2 | 3 | 11 | 77 | 33 | 50000+ | 50000+ | 2230 | 3550 |

(b)

Figure 8: Phase coded rank of top 10 solutions, $K = 45$: successive (a) vs. optimized (b) phase assignment

| | K = 14 | K = 20 | K = 38 |
|---|---|---|---|
| P0 | 1 | 1 | 1 |
| P1 | 1 | 1 | 1 |
| P2 | 1 | 12 | 1 |
| P3 | 1 | 1 | 1 |
| P4 | 4 | 15 | 1 |
| P5 | 1 | 1 | 1 |
| P6 | 1 | 6 | 1 |
| P7 | 141 | 73 | 3316 |
| P8 | 4 | 235 | 2167 |
| P9 | 2 | 1 | 46 |

(a)

| | K = 14 | K = 20 | K = 38 |
|---|---|---|---|
| P0 | 26 | 26 | 10 |
| P1 | 1 | 1 | 1437 |
| P2 | 1 | 1 | 2 |
| P3 | 1 | 1 | 66 |
| P4 | 1 | 132 | 10000+ |
| P5 | 10 | 26 | 26 |
| P6 | 5 | 15 | 646 |
| P7 | 598 | 316 | 8187 |
| P8 | 1 | 5 | 7 |
| P9 | 1 | 4 | 27 |

(b)

Figure 9: MIST. Effect of $K$ on rank best solution: successive (a) vs. optimized (b) phase assignment