

Sparse Representations for Image Decompositions

Davi Geiger¹ Tyng-Luh Liu² Michael J. Donahue³

¹Courant Institute, New York University, New York NY 10012, USA

²Institute of Information Science, Academia Sinica, Nankang 115 Taipei, Taiwan

³IMA, University of Minnesota, Minneapolis MN 55455, USA

Abstract

We are given an image I and a library of templates \mathcal{L} , such that \mathcal{L} is an overcomplete basis for I . The templates can represent objects, faces, features, analytical functions, or be single pixel templates (canonical templates). There are infinitely many ways to decompose I as a linear combination of the library templates. Each decomposition defines a representation for the image I , given \mathcal{L} .

What is an optimal representation for I given \mathcal{L} and how to select it? We are motivated to select a sparse/compact representation for I , and to account for occlusions and noise in the image. We present a concave cost function criterion on the linear decomposition coefficients that satisfies our requirements. More specifically, we study a “weighted L^p norm” with $0 < p < 1$. We prove a result that allows us to generate all local minima for the L^p norm, and the global minimum is obtained by searching through the local ones. Due to the computational complexity, i.e., the large number of local minima, we also study a greedy and iterative “weighted L^p Matching Pursuit” strategy.

1 Introduction

1.1 Image decomposition

In the field of signal processing and computer vision an input signal or image is a function f over some subset of \mathbb{R} or \mathbb{R}^2 . To represent, manipulate and analyze f , it is useful to introduce a linear decomposition into basis elements f_j , i.e., $f = \sum_j c_j f_j$. An example of a well known and useful decomposition of this type is the Fourier series expansion.

Many times the image comes from a class of images where various examples are already known. For instance, when studying images of faces, Kirby and Sirovich [21] as well as Turk and Pentland [33] utilized various exemplary images of faces. They have considered a method to construct basis functions, the *principal component analysis* (PCA), where the functions are chosen adaptively, according to the exemplars. For the case of faces they referred to the adaptive basis functions as eigenfaces. Then, given a new face image, this approach will uniquely decompose the image into the basis. A sparse representation is obtained when very few basis functions for such

decomposition, the principal components of PCA, well approximate the image.

1.2 Overcomplete representations

Let us now consider a library \mathcal{L} , an overcomplete basis, and study how the problem of function decomposition changes. Say, in one dimension, \mathcal{L} consists of two classes of basis functions, sinusoids and functions of the form $1/(k+x)$ ($k \in \mathbb{N}$). Assume that $f(x) = \sin 2x + \frac{4}{(3+x)}$ is our target function (our image) so that \mathcal{L} is an overcomplete basis. It is clear that only two terms from the prototype library are sufficient to represent $f(x)$. However, one could write $f(x)$ using either sinusoids alone or as combinations of $1/(k+x)$ alone, but either representation would require many terms. Indeed, there are infinitely many decompositions (representations) for $f(x)$.

In object recognition the library of templates is overcomplete with respect to the space of all possible images, or to the smaller space of a class of images, e.g., faces (see proof in section 2). Thus, given a library with faces and various other templates, to represent an image of face, we would like to decompose it using a face template, and not other possible decompositions. In this way the recognition problem is a problem of representation, of choosing the optimal decomposition of an image.

1.3 Optimal criteria and solution

What makes a decomposition optimal? For object recognition, this is a cognitive question. Why do we see a face in a face image, even when 50 % of the pixels are occluded, if instead we could simply see the image as a collection of independent gray value pixels (yet another decomposition)? We argue that the optimal representation criterion is to seek a compact representation of an image, to utilize as few templates and yet explain the data as accurate as possible. We are interested in image decompositions that can account for large occlusions. Notice that most work on recognition addressing occlusions, e.g., [17, 22, 23], which we are aware of, does not formulate it as a function decomposition, nor provides a way to incorporate occlusions as in our approach.

We argue that the objective function should be an “ L^p norm” on the utilized coefficients of the decomposition. The issue is that the optimization problem will generally have multiple local minima. We show that it is possible to characterize all local minima and obtain the global one by visiting all of them. Since the number of local minima grows exponentially with the size of the template library we are forced to consider an alternative greedy algorithm.

1.4 Matching pursuit

Inspired by Mallat and Zhang’s work [25] (see also Bergeaud and Mallat [5]), we consider a *matching pursuit* strategy where, at each stage, a best selection is based on minimizing an image residue. In regression statistics, this decomposition method is known as *projection pursuit regression*, a non-parametric method that is concerned with “interesting” projections of high dimensional data (see Friedman and Stuetzle [15], Huber [18]).

Our algorithm is a multi-stage iterative algorithm that at each stage (i) apply the L^p “best-matching” template and, (ii) update the image by removing (subtracting) the object matched by the selected template.

1.5 Background

Let us first remark that our results (not the greedy approach) were first presented in a technical report [11] in 1993. Since then various interesting and related work (obviously independent from us) has arrived [4, 14, 26, 6, 7, 28, 35]. We will try to give a today's (post-facts) view of the field.

Our approach resonates with the following concepts (i) Barlow [2] viewed that the brain is a coding system with redundant representations seeking sparse ones. David Field [14] has followed Barlow's view. (ii) Rissanen's *minimum length encoding* (MDL) [29] description proposed optimal representations for a system, utilizing the least set of parameters (in our case, the parameters are the coefficients of the decomposition). (iii) In mathematics and signal processing, Coifman and Wickerhauser [8] studied the problem of decomposing a signal into overcomplete basis functions. In particular they focused on combining multiple wavelets. An "entropy" criterion was then applied to select the coefficients. Mallat and Zhang [25] have devised a matching pursuit strategy, not committed to any particular form of bases, and to select the coefficients they used an L^2 error norm. At the same time as us, Chen and Donoho [6, 7] studied the overcomplete signal representation problems with L^1 norm optimization. Their optimization method is based on linear programming. Similar techniques were also studied earlier by others [3, 32].

Neural Networks: Feed-forward Neural Networks behave as a function decomposition machinery, e.g., Poggio and Girosi [27]. Thus, one can have overcomplete representations. Poggio and Girosi [28] have recently studied relation between neural networks and sparse representation. Their criterion to the selection of the functions and coefficients is the least squares, i.e., L^2 errors with a smoothness constraint (regularization networks). Olshen and Fields [26] as well as Bell and Sejnowsky [4] did consider the study of sparse representations within the realm of neural networks, at the same time as us. Their approach is to construct error functions to encourage sparse representations and use typical gradient descent techniques to obtain local minima. As we point out again in this paper, the number of local minima of error functions encouraging sparse representations is typically large and so careful study of the landscape of an error function is critical for optimal solutions.

Template Matching: A special case of our approach is when applied to only one object template and various one pixel templates. In this case it is very similar to p -norm minimization (where $p = 2$ is correlation). Correlation methods are known to be optimal for detecting templates in the presence of additive Gaussian noise, but fail miserably when dealing with occlusions or more realistic noise. Minimizing a p -norm, for $0 < p < 1$, is more resistant to occlusions. Ben-Arie and Rao [1] have studied alternative methods for template matching based on non-orthogonal basis functions and minimizing the discriminative signal-to-noise ratio. For one template detection the method is an *expansion matching* and, as pointed out in their paper, it is different from match filtering. Their method gives a better performance than correlation methods for severe occlusion cases since the responses are more peaked than correlation ones. They have also extended the approach to multiple template matching. The comparison with ours is difficult because the formalism is different, not an optimization one, but rooted on linear methods.

2 Template Library and Image Decomposition

In our formulation, a template library, say \mathcal{L} , consists of one canonical template and many non-canonical (object) templates. The canonical template ϵ_1 of size N has zero gray value pixels everywhere except one pixel at the extreme left and top corner where the gray value is 1. More precisely, $\epsilon_1[j] = \delta_{1j}$, where $\delta_{ij} = 1$ for $i = j$ and $\delta_{ij} = 0$ otherwise.

First let us consider a lemma about the canonical template and basis functions followed by a proposition on the overcompleteness of the template library,

Lemma 1 *Canonical template ϵ_1 plus a set of all translations form a basis for the image space.*

Proof: Given an image I of size N , let A_i denote a translation by i pixels (we order the pixels from top to bottom and left to right), i.e., $A_i(\epsilon_1)$ is such that the first template pixel of ϵ_1 is positioned at the i -th image pixel of I . So, the set of all translation is $\{A_i \mid i = 1, \dots, N\}$. Thus, $\{A_i(\epsilon_1) \mid i = 1, \dots, N\}$ is a basis for the image space of size N . This follows directly from

$$I[j] = \sum_{i=1}^N I[i]\delta_{ij} = \sum_{i=1}^N I[i]A_i(\epsilon_1)[j] = \sum_{i=1}^N c_i A_i(\epsilon_1)[j] \quad \text{where } c_i = I[i].$$

□

Proposition 1 *With the set of all translations, a library \mathcal{L} containing non-canonical templates (e.g., edge-templates, eye-templates), as well as canonical template ϵ_1 is overcomplete.*

Proof: From the previous lemma, the canonical template plus all possible translations form a basis. Adding to this basis, or to any basis, a set of different (non-canonical) templates, say eye-templates, face-templates, will create an overcomplete library. □

More generally $\{A_i(\cdot)\}$ denotes some group transformations, possibly accounting for deformations [16], applied to the templates. In our experiments we have just considered the group of translations. Our results here do not offer efficiency when using larger group of transformations, i.e., we would simply increase the search space by the number of new degrees of freedom introduced by larger group of transformations. Further investigation to reduce the search is necessary. Moreover, the space of possible object poses and articulations is enormous and the study of these transformations is beyond the scope of this present work.

2.1 Coordinate transformations

Assume template library $\mathcal{L} = \{\tau_j : j = 1, \dots, M\}$, where we have used $\tau_1 \equiv \epsilon_1$ to represent the canonical template. Let the image data be I of dimension N and each template τ_j be of dimension N_j (we assume all N_j to be perfect square numbers). Note that $N_1 = N$, i.e., the canonical template is of image size. Furthermore, let $Q_j = \{1, 2, \dots, N_j\}$ be the pixel set of τ_j . Again, a translation $A_i(\tau_j)$ is such that the first template pixel of τ_j is positioned at the i -th image pixel. We can explicitly describe such relations as follows:

$$\begin{aligned} Q_j &\xrightarrow{A_i} \Gamma_{ij} \subset \{1, 2, \dots, N\} \quad (\text{denoted as } A_i(Q_j) = \Gamma_{ij}) \\ q \in Q_j &\xrightarrow{A_i} \gamma \in \Gamma_{ij} \quad (\text{denoted as } A_i(q) = \gamma) \end{aligned} \tag{1}$$

where Γ_{ij} represents the set of pixels in the image that the template τ_j is occupying after applying the translation A_i . The mapping formula between $\gamma \in \Gamma_{ij}$ and $q \in Q_j$ is ¹

$$\gamma = i + (\lfloor \frac{q-1}{\sqrt{N_j}} \rfloor \times \sqrt{N}) + (q-1 - \lfloor \frac{q-1}{\sqrt{N_j}} \rfloor \times \sqrt{N_j}).$$

Denote $T_{ij} = A_i(\tau_j)$ and $T_{i1} = A_i(\tau_1) = A_i(\epsilon_1) = \epsilon_{i1} = \epsilon_i$ where $\epsilon_i[j] = \delta_{ij}$. Notice that $T_{ij}[\gamma] = A_i(\tau_j)[\gamma] = \tau_j[q]$.

2.2 Image decomposition equations

We can finally write the image decomposition equations where, for $1 \leq x \leq N$,

$$\begin{aligned} I[x] &= \sum_{i=1}^N \sum_{j=1}^M c_{ij} T_{ij}[x] \\ &= \sum_{i=1}^N c_{i1} \epsilon_{i1}[x] + \sum_{i=1}^N \sum_{j=2}^M c_{ij} T_{ij}[x] \quad (\text{denote } c_{i1} \text{ as } c_k) \\ &= \sum_{k=1}^N c_k \epsilon_k[x] + \sum_{\lambda=N+1}^{MN} c_\lambda T_\lambda[x] \end{aligned} \quad (2)$$

where $\lambda = \lambda(i, j) = i + (j - 1) \times N$. For the remainder of this paper, we will use c_k and c_λ to distinguish the coefficients associated with canonical template and non-canonical ones, respectively.

3 Optimization Criteria and Solution

Our approach is to construct an objective function $F(\mathbf{c})$ (\mathbf{c} is a vector of decomposition coefficients) that when minimized selects a best representation, \mathbf{c}^* , among all solutions \mathbf{c} that satisfy (2). Of course, we have also to be able to solve for \mathbf{c}^* .

3.1 Criteria

The modeling focuses on/requires

- 1. Sparse Representation:** to represent (decompose) an image using as few templates as possible in order to have an economical (minimal) representation.
- 2. Oclusions:** to allow for partial oclusions, i.e., the cost of fitting a template must take into account that portions of the template may have a “bad match”.
- 3. Noise:** to model noise via “noise templates” or canonical template, accounting for the difference between the template fit and the image.

¹The expression $\lfloor x \rfloor$ denotes the greatest integer less than or equal to x .

Let us start by focusing the attention on the coefficients related to the canonical template, c_k for $k = 1, \dots, N$ in equation (2). The canonical template plus translations play a special role of “repair”, i.e., they complete the image from the other templates’ work. More precisely, focusing on one pixel, say $k \in \{1, \dots, N\}$, and using that $\epsilon_k[k] = \delta_{kk} = 1$, equation (2) becomes

$$c_k = I[k] - \sum_{\lambda=N+1}^{MN} c_\lambda T_\lambda[k]. \quad (3)$$

Therefore, a cost on the norm of the canonical template coefficients is equivalent to a cost on the fidelity of fitting all the object templates to the data.

The lack of fidelity can be due to occlusion or image noise. The noise model requires *the cost function to escalate with the magnitude of c_k* . Modeling occlusion suggests equal penalties regardless the magnitude of c_k . Both, noise and occlusion, yield the compromise that *the rate of increase in cost as a function of $|c_k|$ should decrease*.

The above consideration leads us naturally to adopt concave objective functions. In particular, we will primarily study the objective function

$$F_p(\mathbf{c}) = \sum_{i=1}^{MN} \omega_i |c_i|^p \left(= \sum_{k=1}^N \omega_k |c_k|^p + \sum_{\lambda=N+1}^{MN} \omega_\lambda |c_\lambda|^p \right), \quad (4)$$

where, again, N is the number of possible (translations) transformations and M is the size of the template library. The scalars ω_i ’s are positive weights and setting their values is the topic of the next section.

Our previous consideration can now be synthesized. The sparse representation suggests $p = 0$ to count the number of utilized templates (weighted by ω_i). Selections related to the canonical template receive special attention, since their roles are to model occlusions and noise. Occlusions suggest the penalty not to depend on c_k , i.e., $p = 0$. However, noise modeling of any c_k requires a cost penalty depending on how large the error $|c_k| = |I[k] - \sum_{\lambda=N+1}^{MN} c_\lambda T_\lambda[k]|$ is. The balance between both processes, occlusions and noise modeling, lead to $0 < p < 1$.

Comments: Penalties on $|c_\lambda|$, for the non-canonical templates, will bias templates to fit in dark regions since the coefficients get smaller and noise gets smaller. Thus, everything else being equal, the errors get smaller in dark regions.

From equation (3), penalties on $|c_k|$, for the canonical template, become penalties on $|I[k]|$ wherever a template is not present. These penalties encourage (i) non-canonical templates to be present at bright regions (and leave out dark regions instead, i.e., small $I[k]$), and (ii) brighter images to request larger size templates so as not to leave image regions left out (cluttered regions). The resulting balance does not let non-canonical templates to either strongly bias dark or bright regions (still biasing towards larger size templates for brighter images).

3.2 Setting ω_λ

There are many (translated) canonical templates used on an image decomposition (used often to “repair” the decomposition), and we do not want them to dominate the cost. We then normalize the costs by setting the

weights proportional to the corresponding template size, i.e., $\omega_\lambda \propto N_j$, except for the canonical templates where we set their weights to 1 (the number of nonzero pixels). It is also reasonable to encourage templates that exhibit more contrast to be preferred over plain constant templates. The rationale is that a high contrast template that appears in an image is a more significant fact (i.e., less likely) than a plain one (most images have background with plain regions and various regions of various images are smooth). A way to measure this complexity is to use the template variances, possibly with a $\frac{p}{2}$ power ($p < 1$) to shorten the effect, i.e., we set $\omega_\lambda \propto Var(\tau_j)^{\frac{p}{2}}$. Note that the canonical template has $Var(\epsilon_1)^{\frac{p}{2}} \approx 1$. We then propose

$$\omega_\lambda = \omega_{\lambda(i,j)} = N_j \cdot Var(\tau_j)^{\frac{p}{2}}.$$

One can think both criteria, the template size and variance, as a measure of the template complexity, the higher the complexity the more favorable to select it.

3.3 Optimization procedure

Equation (2) can be written in matrix notation as

$$\mathbf{T}\mathbf{c} = \mathbf{I}, \tag{5}$$

where

$$\mathbf{T} = \begin{pmatrix} \epsilon_1[1] & \cdots & \epsilon_N[1] & T_{N+1}[1] & \cdots & T_{MN}[1] \\ \epsilon_1[2] & \cdots & \epsilon_N[2] & T_{N+1}[2] & \cdots & T_{MN}[2] \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \epsilon_1[N] & \cdots & \epsilon_N[N] & T_{N+1}[N] & \cdots & T_{MN}[N] \end{pmatrix},$$

$$\mathbf{c} = (c_1, c_2, \dots, c_{MN})^t \quad \text{and} \quad \mathbf{I} = (I[1], I[2], \dots, I[N])^t.$$

Note that if the prototype library forms a basis (linearly independent), then $M = 1$, and there is no freedom in choosing the coefficients (c_i); the coefficients are uniquely determined by the constraint. If $M > 1$, then there are linear dependencies in the prototype library (the prototype library over-spans) and the set of all solutions (c_i) to the constraint forms an $(M - 1)N$ dimensional affine subspace in the MN -dimensional coefficient space. Let S denote this solution space, i.e., $\dim(S) = (M - 1)N$. Using the above matrix notations, our optimization problem can be written as:

$$\min_{\mathbf{c}} F_p(\mathbf{c}) = \min_{\mathbf{c}} \sum_{i=1}^{MN} \omega_i |c_i|^p \quad \text{subject to } \mathbf{T}\mathbf{c} = \mathbf{I}, \tag{6}$$

where $\mathbf{T} \in \mathbb{R}^{N \times MN}$, $\mathbf{c} \in \mathbb{R}^{MN}$, $\mathbf{I} \in \mathbb{R}^N$, $M > 1$. The constraint space, S , is the set of all \mathbf{c} satisfying $\mathbf{T}\mathbf{c} = \mathbf{I}$, and is an affine subspace of dimension $(M - 1)N$.

It is natural when analyzing F_p in (6) as a function in the coefficient space (c_i) to decompose the domain into octants, where each coefficient is of constant sign. This allows the removal of the absolute values in (6), so we may treat F_p as a smooth function inside each octant. For example, if we consider the restriction of F_p to the

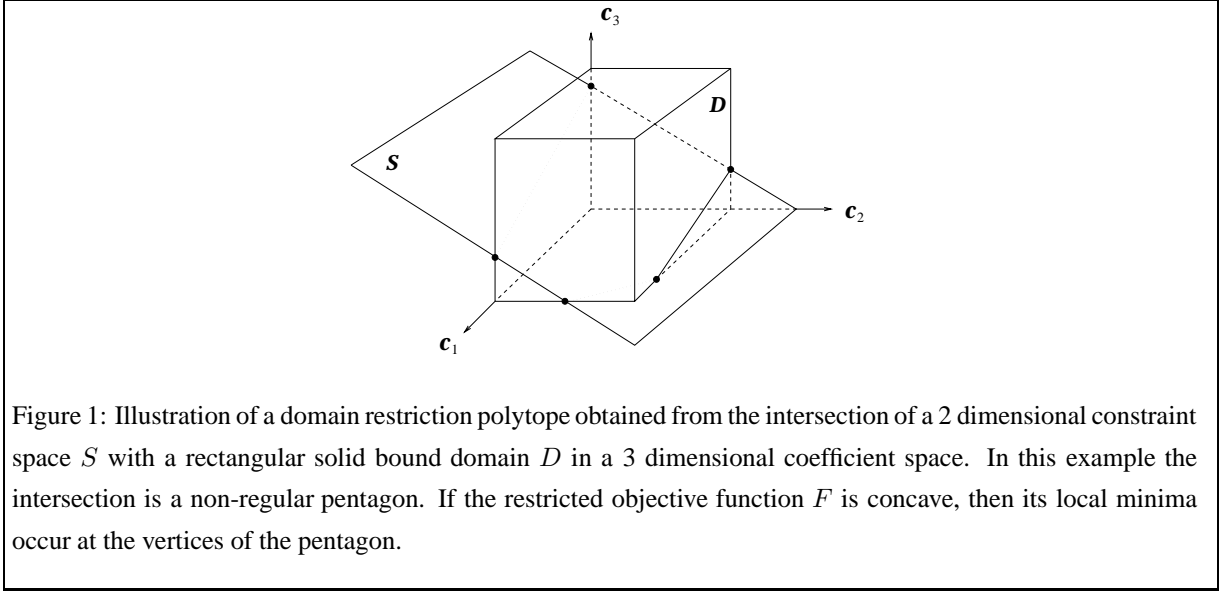


Figure 1: Illustration of a domain restriction polytope obtained from the intersection of a 2 dimensional constraint space S with a rectangular solid bound domain D in a 3 dimensional coefficient space. In this example the intersection is a non-regular pentagon. If the restricted objective function F is concave, then its local minima occur at the vertices of the pentagon.

octant consisting of all points \mathbf{c} such that $c_1 < 0$, $c_2 < 0$, and $c_i > 0$ for $i \geq 3$, then the cost function in (6) becomes

$$F_p(\mathbf{c}) = \omega_1(-c_1)^p + \omega_2(-c_2)^p + \sum_{i=3}^{MN} \omega_i(c_i)^p.$$

Moreover, it is clear that $F_p(\mathbf{c}) \rightarrow \infty$ as $\|\mathbf{c}\| \rightarrow \infty$, so for minimization purposes it suffices to consider bounded \mathbf{c} . The bound will depend upon the constraint equation (5), but, for example, if \mathbf{c}_0 is any solution to (5), then it suffices to consider only those \mathbf{c} satisfying $|c_i| \leq (F_p(\mathbf{c}_0)/\omega_i)^{1/p}$ for all i . Recall that each ω_i is a positive scalar and can be computed in advance. When combined with the restriction to octants, we have a decomposition of the pertinent domain of F_p into MN -dimensional cubes of edge length $(F_p(\mathbf{c}_0)/\omega_i)^{1/p}$.

The intersection of the constraint space S with these domain cubes gives rise to convex polytopes, as illustrated in Figure 1. The system of domain restrictions can be written out explicitly. For the first (positive) octant they are

$$\begin{aligned} \mathbf{T}\mathbf{c} &= \mathbf{I} \\ c_i &\leq d_i, \quad 1 \leq i \leq MN \end{aligned} \tag{7}$$

$$-c_i \leq 0, \quad 1 \leq i \leq MN, \tag{8}$$

where previously we considered the case that each d_i is at least as large as $(F_p(\mathbf{c}_0)/\omega_i)^{1/p}$. The relation $c_1 = (1, 0, \dots, 0)^t \cdot \mathbf{c} \leq d_1$ describes a half-space in the coefficient space \mathbf{c} , and the entire collection (7) and (8) together describe the intersection of $2MN$ halfspaces, i.e., a polytope with at most $2MN$ faces. The general inequality defining a half-space is $\mathbf{v} \cdot \mathbf{c} \leq d_i$, where \mathbf{v} is a vector normal to the bounding hyperplane, and d_i determines an offset from the origin. So an arbitrary convex polytope having N' faces can be described in the form $\mathbf{B}\mathbf{c} \leq \mathbf{d}$, where $\mathbf{B} \in \mathbb{R}^{N' \times MN}$, $\mathbf{d} \in \mathbb{R}^{N'}$, and the inequality is interpreted coordinatewise. So the

generalized constraint relations can be written:

$$\begin{aligned} \mathbf{Tc} &= \mathbf{I} \\ \mathbf{Bc} &\leq d. \end{aligned}$$

These relations can be viewed as defining a polytope inside the affine space S . If we were to perform a bases transformation to obtain coordinates conducive to representations inside S , then F_p under the same transformation would lose its simple form. Even without this consideration, it is useful to study more general objective functions. The specific property of F_p of interest to us is *concavity*. A function F mapping from a convex domain Ω of a vector space X to \mathbb{R} is concave if

$$F(\beta x + (1 - \beta)y) \geq \beta F(x) + (1 - \beta)F(y)$$

for all x and y in Ω and $\beta \in [0, 1]$. The result we desire (Proposition 2) actually requires only a weaker property, which we call *pseudo-concave*. A function $F : \Omega \rightarrow \mathbb{R}$ as above is pseudo-concave if

$$F(\beta x + (1 - \beta)y) \geq \min\{F(x), F(y)\}$$

for all x and y in Ω and $\beta \in [0, 1]$. Clearly any concave function is also pseudo-concave.

Proposition 2 *Let Ω be a closed, bounded, convex polytope in a vector space X , and let $F : \Omega \rightarrow \mathbb{R}$ be pseudo-concave. Then the global minimum of F on Ω occurs at a vertex of Ω .*

Proof: Let $x \in \Omega$. We show that there is a vertex v of Ω such that $F(x) \geq F(v)$, from which the proposition follows.

It is well known that x can be represented as a convex combination of vertices of Ω , i.e., there exists vertices v_1, v_2, \dots, v_n of Ω and corresponding strictly positive coefficients a_1, a_2, \dots, a_n , with $\sum_{i=1}^n a_i = 1$, such that

$$x = \sum_{i=1}^n a_i v_i.$$

If $n = 1$, then $a_1 = 1$, and so $x = v_1$ is a vertex and there is nothing to prove. Otherwise we have by pseudo-concavity that

$$\begin{aligned} F(x) &= F\left(a_1 v_1 + (1 - a_1) \sum_{i=2}^n \frac{a_i}{1 - a_1} v_i\right) \\ &\geq \min\{F(v_1), F(\sum_{i=2}^n a_i v_i / (1 - a_1))\}. \end{aligned} \tag{9}$$

If $n = 2$, then $a_2 = 1 - a_1$, so $F(x) \geq \min\{F(v_1), F(v_2)\}$, and we are finished. Otherwise simply iterate the decomposition on the last term in (9) until

$$F(x) \geq \min\{F(v_1), F(v_2), \dots, F(v_n)\}$$

is obtained. □

4 One Template Matching

In many cases we do not want the representation (2) to use more than one non-canonical template. For example, if we are to find a specific face in an image, then it suffices to use only one face template. Also, when various object templates are present in an image, but no overlapping, it can be reduced to a combination of one template matching problems. The non-canonical template represents a key feature and the canonical template $\epsilon_1 \equiv \tau_1$ (of size N) plus translations represent non-interest elements, e.g., noise.

Let the non-canonical template be τ_2 (of size N_2) and assume A_i is the translation applied to τ_2 . Since $A_i(\tau_2) = T_{i2} = T_{N+i}$, this implies that we look for a decomposition of the form:

$$I[x] = c_{N+i}T_{N+i}[x] + \sum_{k=1}^N c_k \epsilon_k[x]. \quad (10)$$

It is clear that, from (1) and (2), $c_k = I[k]$ if k is outside the template domain, i.e., $k \notin \Gamma_{i2}$. Thus, the equation (5) can be restricted to

$$\begin{pmatrix} \epsilon_{A_i(1)}[A_i(1)] & \cdots & \epsilon_{A_i(N_2)}[A_i(1)] & T_{N+i}[A_i(1)] \\ \epsilon_{A_i(1)}[A_i(2)] & \cdots & \epsilon_{A_i(N_2)}[A_i(2)] & T_{N+i}[A_i(2)] \\ \vdots & \ddots & \vdots & \vdots \\ \epsilon_{A_i(1)}[A_i(N_2)] & \cdots & \epsilon_{A_i(N_2)}[A_i(N_2)] & T_{N+i}[A_i(N_2)] \end{pmatrix} \begin{pmatrix} c_{A_i(1)} \\ \vdots \\ c_{A_i(N_2)} \\ c_{N+i} \end{pmatrix} = \begin{pmatrix} I[A_i(1)] \\ \vdots \\ I[A_i(N_2)] \end{pmatrix}.$$

Again, using $\epsilon_i[j] = \delta_{ij}$, we can rewrite the equation above as

$$\begin{pmatrix} 1 & 0 & \cdots & 0 & \tau_2[1] \\ 0 & 1 & \cdots & 0 & \tau_2[2] \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & \tau_2[N_2] \end{pmatrix} \begin{pmatrix} c_{A_i(1)} \\ \vdots \\ c_{A_i(N_2)} \\ c_{N+i} \end{pmatrix} = \begin{pmatrix} I[A_i(1)] \\ \vdots \\ I[A_i(N_2)] \end{pmatrix},$$

where we have used $T_{N+i}[A_i(q)] = \tau_2[q]$ (note that $A_i(1) = i$). We can also assume that $\tau_2[q] \neq 0$ for $q = 1, \dots, N_2$ (otherwise, we can redefine either τ_2 or the pixel ordering to get a smaller value for N_2).

It follows from Proposition 2 that the local minima of $F_p(\mathbf{c})$ can be found by setting c_{N+i} , $c_{A_i(1)}$, \dots , $c_{A_i(N_2)}$ to zero one at a time. If we set $c_{N+i} = 0$ then we get $c_k = I[k]$ for all k . This is the ‘‘pure noise’’ solution. The first nontrivial (τ_2 -template using) solution sets $c_{A_i(1)} = 0$. This forces the template coefficient $c_{N+i} = I[A_i(1)]/\tau_2[1]$, from which it follows $c_{A_i(q)} = I[A_i(q)] - c_{N+i} \cdot \tau_2[q]$, for $q = 2, \dots, N_2$. The F_p cost of (10), when setting $c_{A_i(1)} = 0$, can then be explicitly calculated as

$$F_p(\mathbf{c}) = \omega_{N+i}|c_{N+i}|^p + \sum_{q=1}^{N_2} \omega_{A_i(q)}|c_{A_i(q)}|^p + \sum_k' \omega_k|c_k|^p, \quad (11)$$

where \sum_k' denotes the summation over $k \in \{1, \dots, N\} - \{A_i(q) \mid q = 1, \dots, N_2\}$. The solution determined by setting $c_{A_i(q)} = 0$ ($2 \leq q \leq N_2$) can be calculated in an analogous fashion.

Finally, the corresponding value for $F_p(\mathbf{c})$, when the solution is determined by setting $c_{A_i(j)} = 0$, for $j = 1, \dots, N_2$, is then

$$\begin{aligned}
F_p(\mathbf{c}) &= \omega_{N+i} \left| \frac{I[A_i(j)]}{\tau_2[j]} \right|^p + \sum_{q=1}^{N_2} \omega_{A_i(q)} \left(\left| I[A_i(q)] - \frac{I[A_i(j)]\tau_2[q]}{\tau_2[j]} \right|^p - |I[A_i(q)]|^p \right) \\
&\quad + \sum_{k=1}^N \omega_k |I[k]|^p \\
&= N_2 \cdot \text{Var}(\tau_2)^{\frac{p}{2}} \left| \frac{I[A_i(j)]}{\tau_2[j]} \right|^p + \sum_{q=1}^{N_2} \left(\left| I[A_i(q)] - \frac{I[A_i(j)]\tau_2[q]}{\tau_2[j]} \right|^p - |I[A_i(q)]|^p \right) \\
&\quad + \sum_{k=1}^N |I[k]|^p,
\end{aligned}$$

where we used $\omega_{N+i} = N_2 \cdot \text{Var}(\tau_2)^{\frac{p}{2}}$ and $\omega_k = 1$ for weights associated with the canonical template. The optimal cost of the match of the template in the (translation) position i is the smallest of the values of $F_p(\mathbf{c})$ across all $N_2 + 1$ solutions \mathbf{c} , i.e., $c_{N+i} = 0$ or $c_{A_i(q)} = 0$ for $q = 1, \dots, N_2$. One can perform a similar analysis for all template translations, and define the matching position of the template to be the position which generated the smallest match cost.

4.1 Simulations

We have designed a sequence of experiments focused on the effects of noise and occlusions to demonstrate both the weighted and unweighted (all weights are set to 1) L^p decomposition methods are superior to the conventional correlation techniques.

The experiments consist of numerous trials on random images with fixed occlusion size and fixed noise variance. The latter determines the signal-to-noise ratio (SNR) for the experiment, defined here as the ratio of the standard deviation of the image to the standard deviation of the noise.

Each trial has four components: an image, a template, an occlusion, and noise. The image is 64 pixels wide by 64 pixels high, randomly generated using an uncorrelated uniform distribution across the range $(-256, 256)$. The template is a 4 pixel by 4 pixel subimage of the image. After selecting the template, a portion of the image from which the template is drawn is ‘‘occluded’’ by redrawing from the same distribution that formed the image, i.e., from an uncorrelated uniform distribution with range $(-256, 256)$. (Occlusion sizes range from 0–14 pixels, from a total subimage size of 16 pixels.) Finally, noise is added to the (occluded) image, drawn from an uncorrelated Gaussian mean-zero random variable.

Translations of the template are compared against the noisy, occluded image, using both weighted and unweighted L^p -norm decomposition method. (Because both the template and the image are drawn from zero-mean random variables, there is little difference between 2-norm error minimization and standard correlation.) For each method the translation position yielding the best score is compared with the position of the original subimage from which the template was formed. If the two agree then the match is considered successful, otherwise the match fails for the trial in question. The first experiment, displayed in Figure 2-(a), displays the percentage of successful match trials at various occlusion sizes and no noise. Dashed curves there show the results from unweighted L^p -norm optimization for $p = 0.125, 0.25, 0.5$. Solid curves correspond to results obtained from our

proposed technique, the weighted L^p -norm optimization, for the same p -values and in addition for $p = 1.0, 2.0,$ and 4.0 . (We do not currently have an appropriate formulation of our method for $p > 1$.) Note that smaller values of p outperform larger values, and that for a given p value our method performs slightly better than minimizing with respect to the corresponding unweighted L^p -norm, providing nearly 100% correct results with $p = 0.125$ for occlusions as large as 11 (out of 16) pixels.

This result is somewhat artificial, however, since noise is generally present in real images. Figure 2-(b) presents results for when noise is present at a SNR of 37. Here we note that $p = 0.125$ still performs very well, although good results can not be obtained if the occlusion is larger than half the template size. Notice that the results using larger values of p are less affected by noise, especially those with $p > 1$.

Figure 2-(c) displays the results for varying noise levels with a constant occlusion size of 5 pixels. Note again that larger values of p produce results which are less sensitive to noise. For example, the results for $p = 0.125$, which are best for large SNR, are poorest for SNR of less than about 3.

The final graph, Figure 2-(d), shows the results obtained by varying the p -value for fixed occlusion and noise levels. We see there that for high noise levels (SNR=2) with no occlusions the best p value is 2. For an occlusion size of 4/16 and a SNR of 9.2, the best p value for our weighted L^p method is somewhere between 0.25 and 0.5, whereas the optimal p value for the unweighted method is somewhat smaller. The remaining curve corresponds to no noise and an occlusion size of 8/16. Also note that for small p (generally best for large occlusions), the performance difference between the weighted and unweighted L^p decomposition method is smallest.

To conclude, these experiments show that both our proposed weighted and unweighted L^p decomposition method with $p < 1$ are superior to standard correlation for template matching in the presence of occlusions and low levels of (Gaussian) noise. Smaller values of p tend to be more robust against occlusions at the cost of greater sensitivity to noise.

4.2 Real image results

To properly gauge the effectiveness of the proposed decomposition approach it is necessary to apply it to natural images. We performed a simple study on one (non-canonical) template matching with real images (see Figure 3).

The image in Figure 3-(a) is a cropped unobstructed view of the subject's face. It is used as a template for matching against an image with the subject's face partially occluded, which we refer to as the **match image**, Figure 3-(b). We then experimented our proposed L^p decomposition method against the standard L^2 correlation approach.

To conduct the match, the template is overlaid at each position in the match image, and a match score is obtained (at each position). This is displayed visually as **match result images** in Figures 3-(c), (d), (f) and (g). Each pixel in these images denotes the match score for the template centered at that position in the corresponding match image, with darker pixels denoting a better match. We constrained the template overlay to remain completely inside the match image, so there is a zero border around each match result image with width equal to half the template size.

Correct recognition results with the L^p decomposition, e.g., see Figures 3-(e), while the L^2 correlation failed to identify the subject's face due to occlusion (Figures 3-(h)). Notice that the **match result images**, Figures 3-(c) and (f), were derived with a *similarity ratio* threshold (to be defined later) set to 0.75, where $Simil = 0$ means

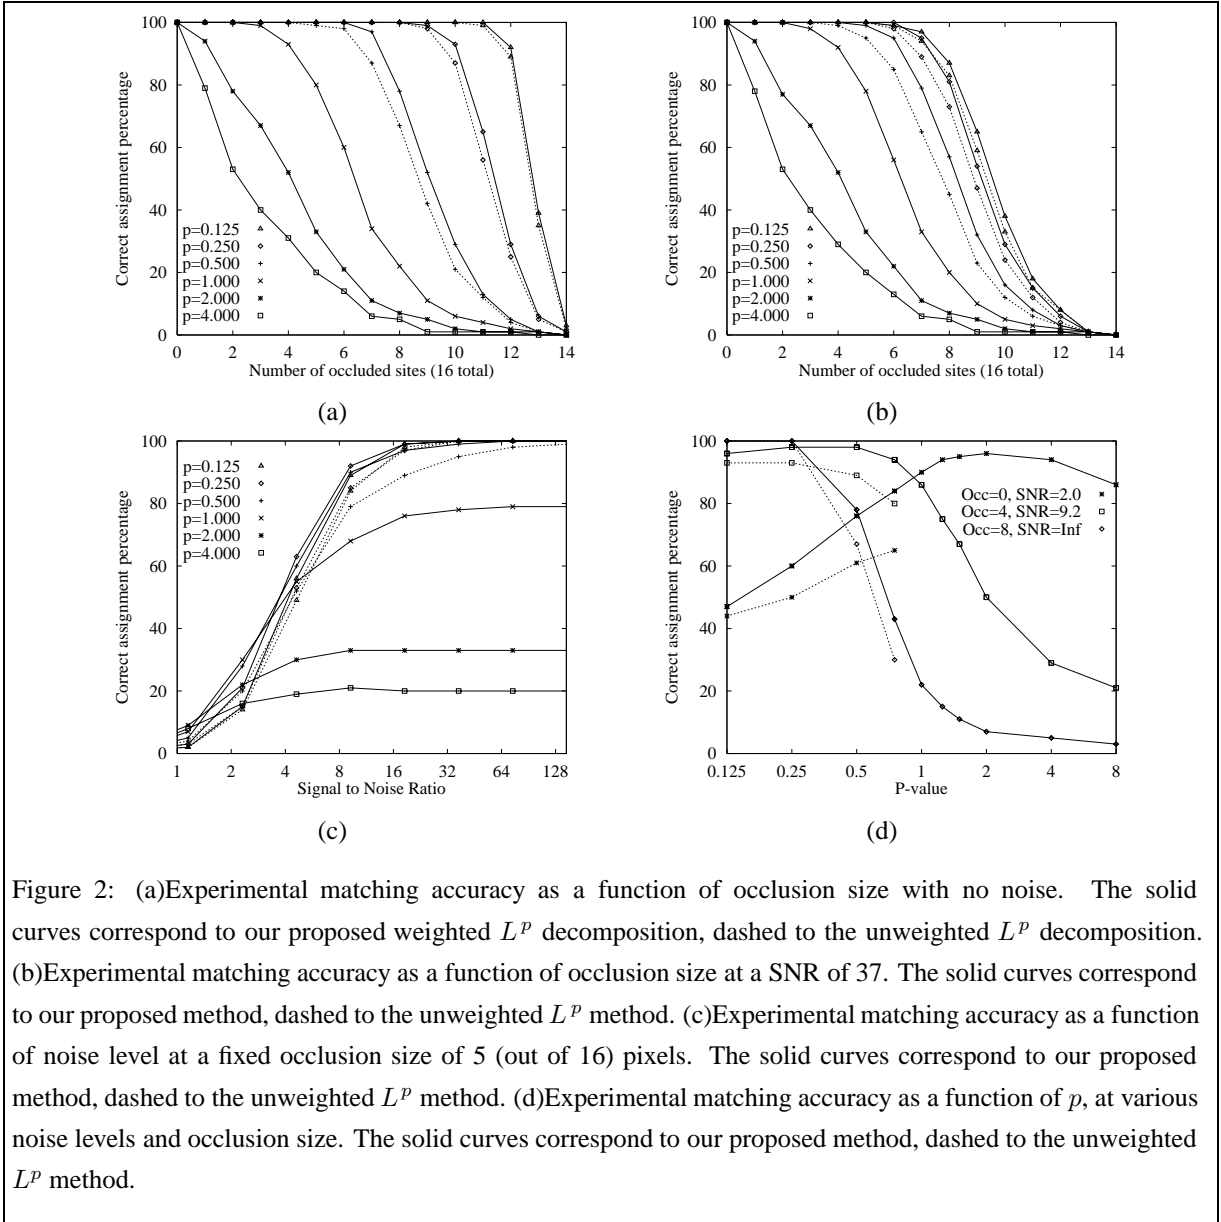


Figure 2: (a)Experimental matching accuracy as a function of occlusion size with no noise. The solid curves correspond to our proposed weighted L^p decomposition, dashed to the unweighted L^p decomposition. (b)Experimental matching accuracy as a function of occlusion size at a SNR of 37. The solid curves correspond to our proposed method, dashed to the unweighted L^p method. (c)Experimental matching accuracy as a function of noise level at a fixed occlusion size of 5 (out of 16) pixels. The solid curves correspond to our proposed method, dashed to the unweighted L^p method. (d)Experimental matching accuracy as a function of p , at various noise levels and occlusion size. The solid curves correspond to our proposed method, dashed to the unweighted L^p method.

similarity threshold is not used.

5 Multiple Templates and Matching Pursuit

The complexity of dealing with many templates is large and the number of local minima to consider grows exponentially with the number of templates. Alternatively, we devise a greedy iterative method where at each stage only one template is selected and thus, we can rely on the result of previous section. This is inspired by Mallat and Zhang’s work on matching pursuit [25] and Bergeaud and Mallat [5].

The original matching pursuit is based on the standard L^2 (Hilbert space) method. To extend to an L^p matching pursuit, $0 < p < 1$, the Hilbert space (inner product) is lost. However, the notion of projection can be recaptured by the optimization criterion for a template to be “best matching” or “closest” to the image. This modification improves robustness of the pursuit scheme but the convergence of L^p pursuit is now not guaranteed. The energy conservation equation and so Jones’ proof [20] of convergence of projection pursuit no longer holds. Despite these limitations, the advantages of efficiency are enormous and we have experimented with this approach.

Let us assume that the residue at the initial stage is the input image, i.e., $R^0 I = I$. Then, at stage n , if a translated template $T_{\lambda_n} (= T_{i_n j_n} = A_{i_n}(\tau_{j_n}))$ and coefficient c_{λ_n} are chosen, the n -th residual image can be updated by “projecting” the $(n - 1)$ -th residue in the direction of T_{λ_n} . More precisely,

$$R^n I[k] = R^{n-1} I[k] - c_{\lambda_n} T_{\lambda_n}[k](1 - O[k]), \quad (12)$$

where $k = 1 \dots N$ and $\lambda_n = (j_n - 1) \times N + i_n$. Note that T_{λ_n} is only of dimension N_{j_n} and we have assumed that $T_{\lambda_n}[k] = 0$ if $k \notin \Gamma_{i_n}$ as defined in (1).

We have introduced the occlusion unit $O[k]$ to avoid additive/repetitive removal of the overlapping pixels. $O[k] = 1$ if pixel k is covered by some template recovered in the previous $(n - 1)$ stages and $O[k] = 0$, otherwise.

The cost function F_p in (6) is the total cost and is a function of a vector \mathbf{c} . The pursuit process is an iteratively greedy algorithm. Let us denote the (iterative) cost function, at stage n , as \mathcal{F}_p^n to be distinguished from the global F_p . Then, at every stage we pursue one template by minimizing \mathcal{F}_p^n , i.e., the cost function \mathcal{F}_p^n is a function of a scalar c .

The cost function \mathcal{F}_p^n , using the L^p norm and taking possible overlappings into consideration, is defined as

$$\mathcal{F}_p^n(c_{\lambda_n}) = \omega_{\lambda_n} |c_{\lambda_n}|^p + \sum_{k \in \Gamma_{i_n}} (|r_k|^p - |I[k]|^p)(1 - O[k]) \quad (13)$$

where $r_k = |R^{n-1} I[k] - c_{\lambda_n} T_{\lambda_n}[k]|$ is the residue at pixel k . The weight is defined as $\omega_{\lambda_n} = (\beta_{\lambda_n} \times Var(\tau_{j_n}))^{\frac{p}{2}}$, where β_{λ_n} is the number of pixels covered by T_{λ_n} .

Clearly, by minimizing the greedy cost \mathcal{F}_p^n at each pursuit stage does not guarantee a global minimization of F_p , the total cost.

An alternative cost we have considered is to divide (13) by the template coefficient $|c_{\lambda_n}|^p$. This cost will guarantee that the template coefficient is not small and so will drive the template away from dark regions. In

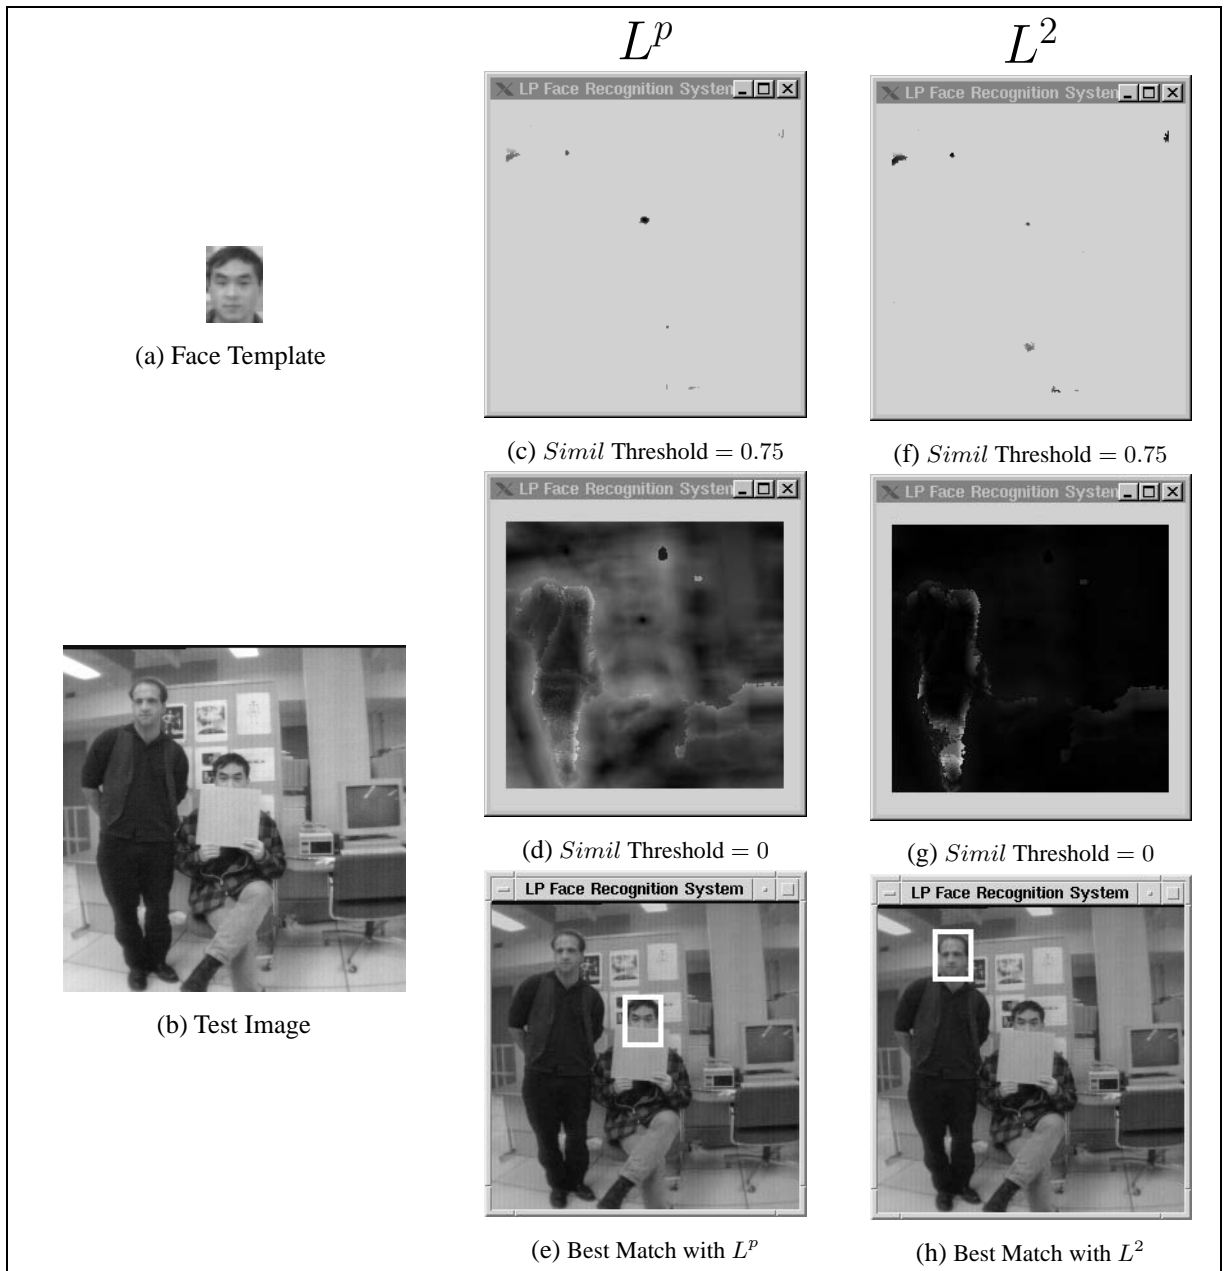


Figure 3: (a) Face template used in the experiments. (b) Test image with the target face partially occluded. (c), (d) and (e) Results using L^p with $p = 0.50$. (c) and (d) Results indicating the matching “response” where the gray level value (darker is better) at each pixel corresponds to the possibility of a potential match. The occluded face is successfully recognized with the L^p method as shown in (e). (f), (g) and (h) Results using standard L^2 correlation. Notice that false recognition has occurred in (h).

this case there is no need for the second term that encourages templates to be placed at bright regions. Thus, an alternative cost is

$$\mathcal{F}_p^n(c_{\lambda_n}) = 1 + \sum_{k \in \Gamma_{i_n}} \frac{|r_k|^p}{\omega_{\lambda_n} |c_{\lambda_n}|^p} (1 - O[k]). \quad (14)$$

5.1 Simplifications

As discussed in section 3 the local minima to $\mathcal{F}_p^n(c)$ are at the vertices of the polytope, each defined by $r_k = 0$, $k \in \Gamma_{i_n}$ and $O[k] = 0$. For each hyperplane $r_k = 0$ the solution of the constraint equation (5) is $c_k = I[k]/T_{\lambda_n}(k)$. We now construct a definition for the solution space associated with each template T_{λ_n} .

Definition 1 *At pursuit stage n , we define the coarse solution space S_{λ_n} generated by T_{λ_n} (with respect to I) as:*

$$S_{\lambda_n} = \{c_k : c_k = [R^{n-1}I(k)/T_{\lambda_n}(k)]^*, \quad k \in \Gamma_{i_n}, \quad O[k] = 0, \quad I[k] \cdot T_{\lambda_n}(k) \neq 0\}.$$

The notation $[x]^*$ is defined as

$$[x]^* = \begin{cases} [x], & \text{when } |x| \geq 1 \\ \frac{1}{[1/x]}, & \text{otherwise} \end{cases}$$

where $[x]$ denotes the closest integer to x . The coarse solution space S_{λ_n} can be classified into subsets by grouping same value of c_k 's.

Definition 2 *Assume that, after grouping, $S_{\lambda_n} = S_{\lambda_n,1} \dot{\cup} S_{\lambda_n,2} \dot{\cup} \dots \dot{\cup} S_{\lambda_n,l}$ of which $m_{\lambda_n} = |S_{\lambda_n,1}| \geq |S_{\lambda_n,2}| \geq \dots \geq |S_{\lambda_n,l}|$ and l, m_{λ_n} are some positive integers. Then a sub-solution set $S_{\lambda_n,k}$, $k \in \{1, 2, \dots, l\}$, is said to be a maximal sub-solution set if $|S_{\lambda_n,k}| = m_{\lambda_n}$.*

As an example of illustration, let $S_{\lambda_n} = \{2, 2, 2\} \dot{\cup} \{4, 4\}$ then $|S_{\lambda_n,1}| = |\{2, 2, 2\}| = 3$ and that the representative (scalar) of S_{λ_n} , in this example, is 2.

With the above definition, we make the following approximation. At pursuit stage n , if T_{λ_n} is indeed embedded in image I and its principal contrast scalar is c_{λ_n} , then

$$\mathcal{F}_p^n(c_{\lambda_n}) = \min_{c \in S_{\lambda_n}} \mathcal{F}_p^n(c) \simeq \min_{c \in S_{\lambda_n,1}} \mathcal{F}_p^n(c).$$

That is, the optimal cost when matching T_{λ_n} can be approximated by only considering for those maximal sub-solution sets of T_{λ_n} . For $p = 0$ and the unweighted L^p pursuit, the above approximation becomes exact, since in this case, to minimize the cost \mathcal{F}_p^n is equivalent to minimize the number of nonzero residues.

Similarity Ratio: We further reduces the search via an ‘‘interesting’’ operator to further filter possible solutions. improve our algorithm.

Definition 3 *At pursuit stage n , for each candidate template T_{λ_n} , we define the following interesting operator,*

$$\text{Simil}_{T_{\lambda_n}} = \frac{|S_{\lambda_n,1}|}{|S_{\lambda_n}|}.$$

Clearly, $0 < \text{Simil}_{T_{\lambda_n}} \leq 1$ and it is called the similarity ratio of T_{λ_n} to its corresponding area in image I .

Applied $Simil_{T_{\lambda_n}}$ to, e.g. the tasks of face recognition as a filter threshold, we may set $Simil_{Threshold} = 0.85$ then this suggests any face template T_{λ_n} can match to any interesting object in I only when it resembles the object more than 85%.

5.2 Algorithm

The algorithm is illustrated by the execution of one iterative stage, say stage n , and the stopping criterion. Let $\mathcal{L}^0 = \mathcal{L}$ be the initial template library. We use an auxiliary set *Overlapped* to memorize pixels matched by more than one templates in the pursuit process. Also, we assume that each non-canonical template can appear only once in the image I . The completion of every stage n is done by executing the following three steps.

A. Matching step : For each non-canonical template τ_j in the template library \mathcal{L}^{n-1} and for each possible translation A_i , let $T_\lambda = A_i(\tau_j)$:

(A-1) compute the solution space S_λ and the representative c_λ for T_λ .

(A-2) compute $Simil_{T_\lambda}$. If $Simil_{T_\lambda}$ is less than $Simil_{Threshold}$ this indicates that T_λ is not a candidate for possible matching. Otherwise, proceed to (A-3).

(A-3) compute $COST_\lambda$ and record the quadruple entry $(T_\lambda, c_\lambda, Simil_{T_\lambda}, COST_\lambda)$ into n -th cost catalogue \mathcal{C}^n . In fact, for $n > 2$, $COST_\lambda$ for T_λ can be looked up directly from the previous cost catalogue \mathcal{C}^{n-1} if the position of T_λ is not overlapped with the best matching template selected in the previous stage. So, in most cases, matching step (A) can be reduced to a single look-up operation.

If all costs in catalogue \mathcal{C}^n are greater than $COST_{Threshold}$, this suggests we have recovered all interested objects in the image I then the algorithm jumps to the stopping stage. Otherwise, continue to the next selection step.

B. Selection step : Choose a template which “best” matches the residual image $R^{n-1}I$ by looking up the cost catalogue \mathcal{C}^n . A best matching template $T_{\lambda^*} (= T_{i^*j^*} = A_{i^*}(\tau_{j^*}))$ to image $R^{n-1}I$ is a template that has the minimal $COST_{\lambda^*}$.

C. Updating step : Remove τ_{j^*} from \mathcal{L}^{n-1} , that is, $\mathcal{L}^n \leftarrow \mathcal{L}^{n-1} - \{\tau_{j^*}\}$. We then update the residual image according (12) and leave the possible overlapping ambiguities to the end.

Stopping stage : When we are at the stopping stage, it implies that we have recovered the main decomposition for I and the remaining task is to check if there are overlapping regions happened during the whole matching pursuit process.

If the set *Overlapped* is not empty, we then decompose it into one or several overlapping regions. In each region, we compute the matching costs restricted on it for all selected templates covered this region. The template producing the smallest cost will be considered as the top most one. This is like “back projection” of the matching pursuit to assume the best approximation achieved by these selected templates. Once all occlusion ambiguities have been resolved, we can match the residual image in the main decomposition areas by the canonical template ϵ_1 .

6 Matching Pursuit Experiments

A small library of face templates has been established (see Figure 4 (a)-(j)). The dimension of all templates is 64×64 . Numerous experiments have been carried out to test our algorithm. To illustrate, consider the three real images, $I_1 - I_3$, in Figure 4 (a)-(c). We obtained decomposition results R_1, R_2 and R_3 shown in Figure 4, for $p = 0.25$ (Similar results were derived for $p = 0.50, 0.75$). For each experiment, it took approximately 30-40 seconds to complete when executed on a Pentium-II PC. When $p = 2$, it is indeed the L^2 matching pursuit method and the recognition results are R_4, R_5 and R_6 . Our proposed L^p matching pursuit has the robustness advantage over the L^2 one.

We have presented a sparse representation for image decomposition problems. Our method is an L^p functional analysis approach. We have shown the L^p sparse representation can account for occlusions where the conventional L^2 correlation will fail. To overcome the complexity of a concave optimization process, we make use the matching pursuit scheme to approximate the global solution.

Appendix: Special Cases of Optimization Criteria

This section presents results to the optimization problem (6) for some special values of p . Typically the finite dimensionality of the spaces \mathbb{R}^N and \mathbb{R}^M plays a prominent role.

Case $p = \infty$

Let us define

$$F_\infty(c) = \lim_{p \rightarrow \infty} (F_p(c))^{1/p} \quad (15)$$

$$= \lim_{p \rightarrow \infty} \left(\sum_{j=1}^M |c_j|^p \right)^{1/p} \quad (16)$$

$$= \max_{1 \leq j \leq M} |c_j|. \quad (17)$$

(Note that we take the p^{th} root of F_p to normalize the limit.) So for $p = \infty$ the minimization criterion is

$$\min_{c \in S^*} F_\infty(c) = \min_{c \in S^*} \lim_{p \rightarrow \infty} (F_p(c))^{1/p}. \quad (18)$$

$$= \min_{c \in S^*} \max_{1 \leq j \leq M} |c_j| \quad (19)$$

The minimizing solution c^* is the element of S^* which has smallest maximal component $|c_j^*|$. Furthermore, for all $p > 0$,

$$\min_{c \in S^*} F_\infty(c) \leq \min_{c \in S^*} (F_p(c))^{1/p} \leq M^{1/p} \min_{c \in S^*} F_\infty(c),$$

so

$$\min_{c \in S^*} F_\infty(c) = \lim_{p \rightarrow \infty} \min_{c \in S^*} (F_p(c))^{1/p}.$$

Therefore, for p large enough, the minimizing solution c^* for F_p will have maximal component $|c_j^*|$ close to the smallest possible subject to the constraint in (6).

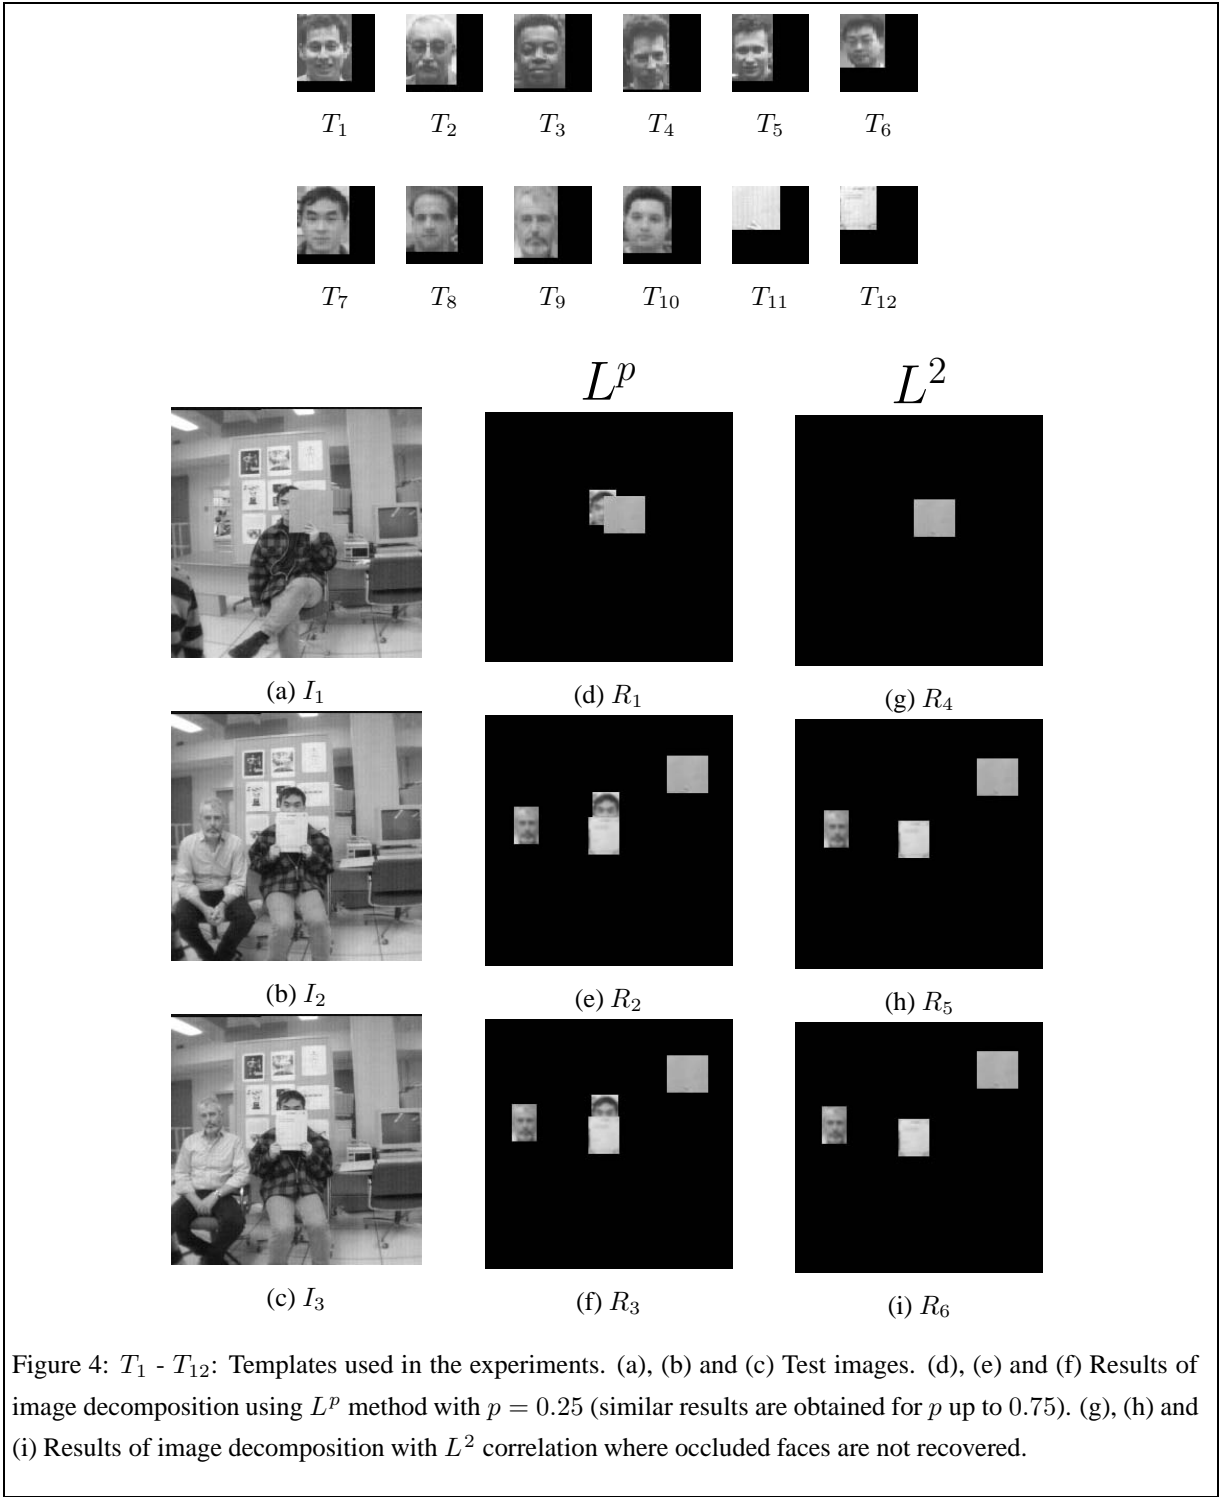


Figure 4: $T_1 - T_{12}$: Templates used in the experiments. (a), (b) and (c) Test images. (d), (e) and (f) Results of image decomposition using L^p method with $p = 0.25$ (similar results are obtained for p up to 0.75). (g), (h) and (i) Results of image decomposition with L^2 correlation where occluded faces are not recovered.

Case $p = 0$

Let us define

$$\begin{aligned} F_0(c) &= \lim_{p \downarrow 0} F_p(c) \\ &= |\{c_j \neq 0 \mid j = 1, 2, \dots, M\}|, \end{aligned} \quad (20)$$

and the minimization criterion for $p = 0$ is

$$\min_{c \in S^*} F_0(c) = \min_{c \in S^*} |\{j \mid c_j \neq 0, j = 1, 2, \dots, M\}|.$$

In other words, the minimizing solution c^* minimizes the number of nonzero components c_j^* . It will be shown later in this document that there exists a finite set $S^{**} \subset S^*$ such that for each $0 < p < 1$, the minimizing point $c^* \in S^*$ of $F_p(c)$ satisfies $c^* \in S^{**}$. It follows from this that

$$\min_{c \in S^*} F_0(c) = \lim_{p \downarrow 0} \min_{c \in S^*} F_p(c),$$

and so for small enough p , the minimizing solution c^* for F_p will be an element of S^* which has the minimal number of nonzero components.

Case $p = 1$

Equation (6) becomes

$$\min_{\mathbf{c}} F_p(\mathbf{c}) = \min_{\mathbf{c}} \sum_{i=1}^{MN} \omega_i |c_i| \quad \text{subject to } \mathbf{T}\mathbf{c} = \mathbf{I}.$$

This problem becomes a linear programming problem as pointed out in [6, 7] and [3][32].

Case $p = 2$

The Hilbert space structure arising when $p = 2$ allows for a closed form solution to the constrained optimization problem. In fact, this is a classical least squares minimization problem.

Let $UDI_{N \times M}V$ be the singular value decomposition of A , where $U \in \mathbb{C}^{N \times N}$ and $V \in \mathbb{C}^{M \times M}$ are unitary, $D \in \mathbb{R}^{N \times N}$ is diagonal, and $I_{N \times M}$ is the matrix with ij -th entry equal to the Kronecker delta δ_{ij} . The diagonal entries of D are the **singular values** of A , and are strictly positive since A is assumed to have full rank. The optimization problem can then be written as

$$\text{Minimize} \quad \sum_{j=1}^M c_j^2 \quad (22)$$

subject to

$$UDI_{N \times M}Vc = b. \quad (23)$$

We can rewrite the constraint as

$$I_{N \times M}Vc = D^{-1}\bar{U}^T b, \quad (24)$$

where \bar{U}^T denotes the conjugate transpose (which is the inverse) of U . Since unitary matrices preserve the L^2 -norm, minimizing the norm of c is equivalent to minimizing the norm of Vc . But clearly the last $M - N$ components of Vc , $(Vc)_{N+1}, (Vc)_{N+2}, \dots, (Vc)_M$ are unconstrained, so the solution with minimal norm has these components set to 0. This means that the minimal solution satisfies

$$Vc = I_{M \times N} D^{-1} \bar{U}^T b, \quad (25)$$

or

$$c = \bar{V}^T I_{M \times N} D^{-1} \bar{U}^T b. \quad (26)$$

The matrix $A^I = \bar{V}^T I_{M \times N} D^{-1} \bar{U}^T$ is called the **pseudo-inverse** of A . (See [10].)

Theorem 1 *The minimizing point $c^* \in S^*$ of $F_2(c)$ is given by*

$$c^* = A^I b.$$

A special case occurs when the singular values D_{ii} are identical:

Theorem 2 *If the singular values of A are identically equal, say $D_{ii} = \sigma$ for all $i = 1, 2, \dots, N$, then the minimizing point $c^* \in S^*$ of $F_2(c)$ is given by*

$$c^* = \sigma^{-2} A^T b.$$

Proof: Using the preceding theorem and $D = \sigma I_{N \times N}$ we get

$$\begin{aligned} c^* &= \bar{V}^T I_{M \times N} D^{-1} \bar{U}^T b \\ &= \sigma^{-1} \bar{V}^T I_{M \times N} \bar{U}^T b \\ &= \sigma^{-2} \bar{V}^T I_{M \times N} D \bar{U}^T b \\ &= \sigma^{-2} A^T b. \end{aligned}$$

□

As an example, if the matrix A is composed of columns forming separate sets of orthonormal bases (for \mathbb{R}^N), then the solution is given after proper renormalization by projecting b onto each column.

Corollary 1 *Let $p = 2$ and assume N divides evenly into M . Suppose, moreover, that the columns $A_{kN+1}, A_{kN+2}, \dots, A_{(k+1)N}$ of A form an orthonormal basis for \mathbb{R}^N for each $k = 0, 1, \dots, (M/N) - 1$. Then the minimizing point $c^* \in S^*$ of $F_2(c)$ is given by*

$$c_j^* = \frac{N}{M} \langle A_j, b \rangle, \quad (27)$$

and

$$F_2(c^*) = \frac{N}{M} \|b\|_2^2. \quad (28)$$

Proof: The rows of A are pairwise orthogonal, and have magnitude (with respect to the 2-norm in \mathbb{R}^M) of $\sqrt{M/N}$. Therefore the N singular values of A are identically equal to $\sqrt{M/N}$, and the result follows immediately from the preceding theorem. □

Unconstrained optimization

The singular value decomposition presented above for $p = 2$ can be used in general to recast the constrained optimization problem as an unconstrained optimization problem. Recall that we want to minimize $F_p(c)$ subject to the constraint $Ac = b$, which can be rewritten

$$I_{N \times M} Vc = D^{-1} \bar{U}^T b.$$

Letting \tilde{c} denote the vector Vc , note that the first N components of \tilde{c} are fixed by the constraint, but the last $M - N$ are unconstrained. So the original optimization problem is equivalent to

$$\min_{\tilde{c}_{N+1}, \tilde{c}_{N+2}, \dots, \tilde{c}_M} F_p(\bar{V}^T \tilde{c}) \quad (29)$$

where

$$\begin{pmatrix} \tilde{c}_1 \\ \tilde{c}_2 \\ \vdots \\ \tilde{c}_N \end{pmatrix} = D^{-1} \bar{U}^T b. \quad (30)$$

This result cannot be used to provide a closed form solution for $p \neq 2$ because the matrix V does not in general preserve the p -norm.

References

- [1] J. Ben-Arie, and K. R. Rao, “*On the Recognition of Occluded Shapes and Generic Faces Using Multiple-Template Expansion Matching*,” Proceedings IEEE International Conference on Pattern Recognition, New York City, 1993.
- [2] H. B. Barlow, *The Coding of Sensory Messages. Current Problems in Animal Behavior*, Cambridge, Cambridge University Press, 1961.
- [3] I. Barrodale and F. Roberts, “*Solution of an Overdetermined System of Equations in the L_1 Norm*,” Math. Dept. Report No. 69, University of Victoria, 1972.
- [4] A. J. Bell and T. J. Sejnowski, “*An information-maximisation approach to blind separation and blind deconvolution*,” *Neural Computation*, 7 (1995), pp. 1129–1159.
- [5] F. Bergeaud and S. Mallat, “*Matching Pursuit of Images*,” SPIE, Orlando, 1995.
- [6] S. Chen and D. Donoho, “*Atomic Decomposition by Basis Pursuit*,” Technical Report, Stanford University, May, 1995.
- [7] S. Chen and D. Donoho, “*Basis Pursuit*,” TR, Stanford University, Nov. 1994.
- [8] R. Coifman and V. Wickerhauser, “*Entropy-based Algorithms for Best Basis Selection*,” *IEEE Transactions on Information Theory*, vol. 38, no. 2, 1992.

- [9] T. H. Cormen, C. E. Leiserson and R. L. Rivest, *Introduction to Algorithms*, McGraw-Hill, 1990.
- [10] G. Dahlquist and Å. Björck, *Numerical Methods*, Prentice-Hall, Englewood Cliffs, N.J., 1974.
- [11] M. J. Donahue and D. Geiger, “*Template Matching and Function Decomposition Using Non-Minimal Spanning Sets*,” Technical Report, Siemens, 1993.
- [12] M. J. Donahue, D. Geiegr, R. Hummel and T. Liu, “*Representations for Image Decomposition with Occlusions*,” Proc. Computer Vision and Pattern Recognition, pp. 7-12, San Francisco, 1996.
- [13] H. Eklblom, “ *L_p -methods for Robust Regression*,” BIT 14, pp. 22-32, 1973.
- [14] D. Field, “*What Is the Goal of Sensory Coding*,” Neural Comp. 6, pp. 559-601, 1994.
- [15] J. H. Friedman and W. Stuetzle, “*Projection Pursuit Regression*,” Journal of the American Statistical Association, vol. 76, pp. 817-823, 1981.
- [16] U. Grenander, “*Advances in Pattern Theory: The 1985 Rietz Lecture*,” Annals Statists., vol. 17, no. 1 pp. 1-30, 1989.
- [17] Huttenlocher, D., G. Klanderman, and W. Rucklidge, 1993, “*Comparing Images Using the Hausdorff Distance*,” IEEE Transactions on Pattern Analysis and Machine Intelligence, **15**(9):850-863.
- [18] P. J. Huber, “*Projection Pursuit*,” The Ann. of Stat., vol. 13, No.2, pp. 435-475, 1985.
- [19] P. J. Huber, *Robust Statistics*, John Wiley & Sons, New York, 1981.
- [20] L. K. Jones, “*On a Conjecture of Huber Concerning the Convergence of Projection Pursuit Regression*,” The Ann. of Stat., vol. 15, No.2, pp. 880-882, 1987.
- [21] M. Kirby and L. Sirovich, “*Application of the Karhunen-Loève Procedure for the Characterization of Human face*,” IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-12(1):103-108, January 1990.
- [22] J. Krumm, “*Eigenfeatures for Planar Pose Measurement of Partially Occluded Objects*,” Proc. Computer Vision and Pattern Recognition, pp. 55-60, San Francisco, 1996.
- [23] A. Leonardis and H. Bischof, “*Dealing with Occlusions in the Eigenspace Approach*,” Proc. Computer Vision and Pattern Recognition, pp. 453-458, San Francisco, 1996.
- [24] T. Liu, M. Donahue, D. Geiger and R. Hummel, “*Image Recognition with Occlusions*,” Proc. 4th European Conf. on Computer Vision, London, 1996.
- [25] S. Mallat and Z. Zhang, “*Matching Pursuit with Time-Frequency Dictionaries*,” IEEE Trans. on Signal Processing, Dec. 1993.
- [26] B. A. Olshausen and D. J. Field, “*Natural image statistics and efficient coding*,” Network: Computation in Neural Systems, 7 (1996), pp. 333–339.

- [27] T. Poggio and F. Girosi, “Regularization algorithms for learning that are equivalent to multilayer network,” *Science*, 247 (1990), pp. 978–982.
- [28] T. Poggio and F. Girosi, “A Sparse Representation for Function Approximation,” *Neural Computation*, in press, 1998.
- [29] J. Rissanen, “Modeling by Shortest Data Description,” *Automatica*, 14:465-471, 1978.
- [30] P. J. Rousseeuw and A. Leroy, *Robust Regression and Outlier Detection*, John Wiley, New York, 1987.
- [31] D. L. Ruderman and W. Bialek, “Statistics of natural images: Scaling in the woods,” *Phys Rev. Letters*, 73 (1994).
- [32] K. Spyropoulos, E. Kountouzis and A. Young, “Discrete Approximation in the L_1 Norm,” *The Computer Journal*, Vol. 16, No. 2, pp. 180–186, 1973.
- [33] M. Turk and A. Pentland, “Eigenfaces for Recognition,” *J. of Cognitive Neuroscience*, vol. 3, pp. 71-86, 1991.
- [34] B. Uhrin, “An Elementary Constructive Approach to Discrete Linear l_p -approximation, $0 < p \leq +\infty$,” *Colloquia Mathematica Societatis János Bolyai*, 58. Approximation Theory, Kecskemét, 1990.
- [35] S. C. Zhu, Y. N. Wu, and D. Mumford, “FRAME (Filters, Random Fields, and Minimax Entropy): Towards a unified theory for texture modeling”. *Proc. Comp. Vision and Patt. Recog.*, San Francisco, 1996.