

Department of Electrical
and
Computer Systems Engineering

Technical Report
MECSE-9-2005

An Analysis of Linear Subspace Approaches for Computer
Vision and Pattern Recognition (supersedes MECSE-6-2003)

P. Chen and D. Suter

MONASH
UNIVERSITY

An Analysis of Linear Subspace Approaches for Computer Vision and Pattern Recognition

Pei Chen and David Suter

*ARC Centre for Perceptive and Intelligent Machines in Complex Environments
Department of Electrical and Computer Systems Engineering
Monash University, Australia, 3800
{pei.chen, d.suter}@eng.monash.edu.au*

Abstract: *Linear subspace analysis (LSA) has become rather ubiquitous in a wide range of problems arising in pattern recognition and computer vision. The essence of these approaches is that certain structures are intrinsically (or approximately) low dimensional: for example, the factorization approach to the problem of structure from motion (SFM) and principal component analysis (PCA) based approach to face recognition. In LSA, the singular value decomposition (SVD) is usually the basic mathematical tool. However, analysis of the performance, in the presence of noise, has been lacking. We present such an analysis here. First, the “denoising capacity” of the SVD is analysed. Specifically, given a rank- r matrix, corrupted by noise – how much noise remains in the rank- r projected version of that corrupted matrix? Second, we study the “learning capacity” of the LSA-based recognition system in a noise-corrupted environment. Specifically, LSA systems that attempt to capture a data class as belonging to a rank- r column space will be affected by noise in both the training samples (measurement noise will mean the learning samples will not produce the “true subspace”) and the test sample (which will also have measurement noise on top of the ideal clean sample belonging to the “true subspace”).*

These results should help one to predict aspects of performance and to design more optimal systems in computer vision, particularly in tasks, such as SFM and face recognition. Our analysis agrees with certain observed phenomenon, and these observations, together with our simulations, verify the correctness of our theory.

Index terms: *Singular value decomposition, Linear subspaces, Principal component analysis, Structure from motion, Factorization method, Homography, Face recognition, Matrix perturbation, First-order perturbation, Multiple eigenvalue/singular value.*

1. Introduction

Linear subspace analysis has found application in many problems in computer vision and pattern recognition, where the high-dimensional representations of certain structures are intrinsically (or approximately) low dimensional. In this paper we focus on several very prominent problems: Structure from Motion (SFM), homography estimation, and PCA-based face recognition, although several other computer vision and pattern recognition tasks fall within the framework of our analysis.

Note: This paper supersedes technical report MECSE-6-2003 P. Chen and D. Suter "An Analysis of Linear Subspace Approaches for Computer Vision and Pattern Recognition". The present paper contains corrections, an extra section on homography refinement (see section 3.1) and the text is considerably re-organised).

1.1. Applications of Linear Subspace Analysis

Homography Estimation

The induced homography matrix over 2 views lies in a dimension-4 subspace (Shashua and Avidan 1996). Similarly, the "relative homographies" of 2 planes over multiple views reside in a dimension-4 subspace (Zelnik-Manor and Irani 1999; Zelnik-Manor and Irani 2002). Furthermore, the dimension-4 constraint also holds for the case of multiple-plane-over-multiple-view (Zelnik-Manor and Irani 1999; Zelnik-Manor and Irani 2002). Exploitation of this low rank constraint is potentially very attractive for subsequent tasks (e.g., the 3-D structure of the scene being imaged, image mosaicing etc.). However, this paper will show that the denoising gains, of exploiting the low rank constraint in this context, are fundamentally more limited than the potential gains of exploitation of the low rank constraint in Structure From Motion.

Structure from Motion (SFM)

In the SFM context, one extracts from the image sequence the coordinates of various tracked points (or other geometric features such as lines). These coordinates may be assembled into a measurement matrix, which is essentially low dimensional despite the matrix (itself) usually being physically huge. For example, under the affine models, the measurements are generally restricted to a dimension-4 subspace (Tomasi and Kanade 1992; Poelman and Kanade 1997; Irani 1999; Kahl and Heyden 1999; Kanatani 2001; Irani 2002). (Although the registered measurement matrix can be of rank 3 (Tomasi and Kanade 1992; Poelman and Kanade 1997; Kahl and Heyden 1999; Kanatani 2001).) In such situations one can use the low-rank constraint, imposed by an SVD based rank projection, to effectively “denoise” the measurement matrix (and to also provide a starting factorization into the product of a matrix related to structure and a matrix related to motion). Standard techniques (Tomasi and Kanade 1992; Poelman and Kanade 1997; Irani 1999; Kahl and Heyden 1999; Kanatani 2001; Irani 2002) can then be used to factor the “denoised” matrix into structure and motion factors. We are not concerned with the factorization stage here.

It is important to realize that in Structure from Motion, the exploitation of the low rank constraint (to reduce noise) is only one part of a complete system. Moreover, there are many practical issues that need to be tackled: first and foremost, the data matrix is invariably going to be incomplete as the tracker loses track or loses sight of features during the tracking stage. In an already published paper (Chen and Suter 2004a), we tackled this particular issue and the solution we provided was informed by the theory/analysis we present here. In that context, the theory enabled us to develop a

heuristic that could be used to decide when the denoising capacity gained by adding extra observations (extra tracked features or extra images containing previously tracked features) was outweighed by the possible adverse affects that extra data may produce in terms of the need to impute the missing values in tracks. The current paper focuses on the more general application of, and theoretical justification of, the theory behind the heuristic in (Chen and Suter 2004a). For our purposes here, and for space limitations, we focus only on the denoising capacity aspects and ignore implementation issues of a “complete” SFM algorithm.

PCA based Face Recognition

Another particularly active area of computer vision research, also employing subspace analysis, is that of PCA-based face recognition* (Turk and Pentland 1991; Hallinan 1994; Eipstein *et al.* 1995). A human face, in typical applications, must be recognized despite illumination changes between the target image (to be recognized) and the database of candidate images. It has been observed that: “the variations between the images of the same face due to illumination and viewing direction are almost larger than image variations due to change in face identity” (Moses *et al.* 1994).

Here, we have to clarify the difference between the common PCA (Turk and Pentland 1991; Hallinan 1994; Eipstein *et al.* 1995) and linear subspace analysis (Belhumeur *et al.* 1997; Basri and Jacobs 1999; Basri and Jacobs 2003). In face recognition and related applications, several terminologies, like PCA (Turk and Pentland 1991), eigenface (Turk and Pentland 1991) and eigenimage (Eipstein *et al.* 1995), have been used for such dimensionality reduction techniques. PCA (Turk and Pentland 1991; Hallinan 1994; Eipstein *et al.* 1995) works on the correlation matrix, where the mean of the images is first subtracted. While, in linear subspace analysis, we work directly on the original data (Belhumeur *et al.* 1997; Basri and Jacobs 1999; Basri and Jacobs 2003), without subtracting their mean. Recently, some theoretical analysis (Basri and Jacobs 1999; Ramamoorthi and Hanrahan 2001; Ramamoorthi 2002; Basri and Jacobs 2003) and experimental result (Belhumeur *et al.* 1997) have proved that better performance can be obtained directly by using the linear subspace analysis, without subtracting the mean. In section 5, we will analyse the performance of the linear subspace analysis, without subtracting the mean, as in (Ramamoorthi and Hanrahan 2001).

In order to tackle this issue, PCA has been utilized to model the lighting variation in images; because it has been proved, experimentally (Hallinan 1994; Murase and Nayar 1994; Nene *et al.* 1994; Eipstein *et al.* 1995; Murase and Nayar 1995; Yuille *et al.* 1999) and theoretically (Basri and Jacobs 1999; Ramamoorthi and Hanrahan 2001; Ramamoorthi 2002; Basri and Jacobs 2003), that the possible images of the same Lambertian object, under different lighting conditions, approximately concentrate in a low-dimensional subspace.

Similar approaches can be used in general object recognition and pose determination systems. A particularly influential example of such was the SLAM system (Murase and Nayar 1994; Nene *et al.* 1994; Murase and Nayar 1995), which captured the variations due to pose and illumination by a 20-dimensional (or less) subspace. Recently, it was proved, by using spherical harmonics, that “all Lambertian reflectance functions obtained with arbitrary distant sources lie in close to a 9D linear subspace”: Basri and Jacobs (Basri and Jacobs 1999; Basri and Jacobs 2003) and Ramamoorthi and Hanrahan (Ramamoorthi and Hanrahan 2001; Ramamoorthi 2002).

Our noise SVD analysis can be applied to PCA based face recognition in both the training/learning phase and in the testing/retrieval phase. Existing publications (Chen and Suter 2004b; Chen and Suter 2004c) have already *used* the theory presented here to make some improvement to Face Recognition approaches (however, those publications concentrate on other aspects of subspace based face recognition, such as shadow and other outlier removal issues, and did not provide the theoretical analysis contained in this paper).

Other Low Rank Constrained Computer Vision Problems

Rank constraints can be applied to many other problems, for example in relation to fundamental matrix and trifocal tensors (Ma *et al.* 2004). We have chosen to concentrate, here, on the above three sets of problems, partly for focus and brevity, but also to avoid detailed consideration of issues such as heteroscedastic noise, which would complicate the analysis.

1.2. Low Rank Constraint and SVD

In the presence of noise, the matrix in question becomes full rank. Thus, the matrix has to be fitted to its closest low-rank approximation. The SVD gives the best solution to this problem (Golub and Loan 1996), measured by the Frobenius norm and 2-norm. The result is guaranteed to be optimal (Press *et al.* 1992) if the noise is i.i.d. Gaussian. Not surprisingly, therefore, the SVD has become a widely used tool. For example, the factorization method (Tomasi and Kanade 1992; Poelman and Kanade 1997) achieves a *Maximum Likelihood* affine reconstruction from multiple (>2) views, as pointed out in (Hartley and Zisserman 2000).

From a related point of view, the low-rank approximation can be regarded as a “denoising” tool, where we refer to the measure of the sum of squared difference (SSD)* between the noise-corrupted matrix (or the “denoised” matrix) and the noise-free matrix. Compared with a noisy matrix that is always of full rank, its low-rank approximation matrix, obtained by SVD, is always closer to the noise-free matrix, i.e. the underlying ground truth. For example, the multiview subspace constraint was utilized to improve the

* In image denoising, we usually use the terminology of mean square error (MSE).

accuracy of estimated homographies, especially for those that have small regions (Zelnik-Manor and Irani 1999; Zelnik-Manor and Irani 2002).

Thus, linear subspace approximation is sometimes a model simplification and sometimes a denoising process (and often both, simultaneously).

Despite the well-deserved popularity of the SVD, little attention has been placed on carefully analyzing the performance of the SVD-based projections. A form of sensitivity analysis can be performed through studying the Jacobian of the SVD (Mathai 1997). Indeed, (Papadopoulos and Lourakis 2000) used such an approach in the estimation of the Epipolar uncertainty and the estimation of the covariance of rigid 3D motion.

However, in this section we pose questions that are best tackled through a statistical perturbation analysis.

Denoising capacity of SVD

It is well known (Golub and Loan 1996) that one can, by SVD, obtain the best solution to the low-rank approximation, measured by 2-norm or Frobenius-norm. However, we do not know its capacity of separating the signal from the noise. Supposing the noise level is small enough, how much signal is retained by keeping the largest r components? Or, how much noise has been reduced by discarding the other components? In this sense, we are *blindly* using a SVD, *without* knowing its *denoising capacity*: How close is the low-rank approximation matrix to the noise-free matrix, or how close is the SVD-based subspace to the ground-truth subspace?

The lack of such performance analysis impedes the careful design of optimal systems. A natural issue arising is to characterize the achieved accuracy with the growth

in data (in the SFM context, this can be either through a growth in the number of frames analyzed, or by a growth in the number of features tracked). In the factorization approach to SFM, it is widely accepted that more frames produce more accurate result than a few frames ("few" typically being little more than 3). It was even claimed that the 3D scene could be reconstructed to arbitrary accuracy given enough frames (Thomas and Oliensis 1999).

However, *what is the gain of adding the data from one extra frame to a very large measurement matrix? What happens as the number of the frames approaches infinity? Can the 3D scene be truly reconstructed with arbitrary accuracy? Can such arbitrary accuracy be achieved only by the increase of the frames (while the number features do not increase)? Is an increase in the number of frames the most efficient way to obtain an increase in accuracy?*

In the example of SFM, as suggested above we can also possibly augment the number of feature points, or we can augment the number frames, or we can do both: i.e., both the row and the column of the matrix can grow towards the infinite in size. However, in a related problem, the induced homography matrix is restricted to a class of $m \times 9$ matrices (Shashua and Avidan 1996; Zelnik-Manor and Irani 1999; Zelnik-Manor and Irani 2002). Such a matrix can only "grow" in one dimension, not both. We introduce some terminology to describe this difference: We call the matrix potentially-double-infinite if it has infinite rows and columns, and potentially-single-infinite for those who has constant rows (columns) and infinite columns (rows). This raises another question: *What is the difference between these two types of matrices in terms of the precision that can be achieved?*

In summary, the first aim of this paper is to analyze the *denoising* capacity of SVD, i.e., to identify the error that still resides in the low-rank approximation matrix and how this error relates to the growth of additional data.

Learning capacity of linear subspace analysis

Questions, different from those posed above, arise from the face recognition and similar applications (including the object recognition and pose determination, and related applications). In the PCA-based face recognition approach, the eigenimage representation relies on a compact approximation of the large image database (or "learning" set), by spanning this set (approximately) with a few orthogonal basis images.

In this paper, we consider the data (images) as being column vectors so that the learning phase involves trying to model the class of faces by a learned column space. In more detail, we adopt the simple description of the LSA-based face recognition algorithm in (Belhumeur *et al.* 1997; Georghiades *et al.* 1998; Georghiades *et al.* 2001). It consists of two steps: the off-line learning stage and the on-line recognition stage. In the learning stage, the image basis is obtained in this way: a set of learning images for one face is arranged as a learning matrix \mathbf{A} so that each image is regarded as one column of the learning matrix \mathbf{A} . Suppose the face image has a dimension of m , and n learning samples are collected. $\mathbf{A} \in R^{m,n}$. The r ($r \ll m$ and $r \leq n$) basis images can be obtained as the first r left singular vectors of \mathbf{A} , which correspond to the r largest singular values. In the on-line recognition stage, a test image is projected on the r basis images and its distance to the image basis is used for recognition. Note, that in the cones-based approaches

(Georghiades *et al.* 1998; Georghiades *et al.* 2001) and in the Fisherfaces approach (Belhumeur *et al.* 1997), the learning set is a set of images of different persons.

Of course, if the space of images of a face was exactly an r -dimensional linear subspace then only r linearly independent images of a person would be required to completely and accurately capture that subspace. However, such a view is overly optimistic. For example, although the "illumination cone" (Belhumeur and Kriegman 1998) (see also (Zhao and Yang 1999)) can be obtained by as little as three images, the result is usually not accurate enough. First, there is inevitably some noise in the images, like quantization error. Second, it is difficult to satisfy the conditions in proposition 3 in that paper (Belhumeur and Kriegman 1998). Even if we can have three distinct light sources that can shed light on all the points of the surface, we can't, in practice, exclude other light sources that cause attached or cast shadows on the subject. These considerations, plus general noise, have generally resulting in researchers trying to "learn" the eigenimages by a data reduction step applied to many "learning samples".

In such a setting, we pose and then attempt to answer, several questions: *What is the relationship between the learning capacity and the size of the learning samples?*

Note that the test image itself contains noise. *Thus the noise in the LSA-based recognition system comes from two sources: one from the basis and the other from the test image. Do these two types of noise interfere with each other?*

In summary, the second aim of this paper is to present some theoretical analysis of the learning capacity of LSA-based recognition systems. Specifically, the error (measured by the sum of squared differences – SSD) of the LSA-based recognition system can be separated into two parts: one from the basis and the other from the test image, and we

obtain some analytical results about their effects on the performance of the recognition system. We show that it is possible, theoretically, to design the optimal recognition system if we know the expectation of the test images.

1.3. Paper Outline

In this paper, we answer the above questions by analyzing the performance of SVD in a noise-corrupted environment using matrix perturbation theory as a major tool. In section 2, we first present our results, concerning the denoising capacity of a large low rank matrix and the learning capacity of the LSA system. (In appendix A.1, some preliminary knowledge concerning the SVD and matrix perturbation theory is summarized. In appendices A.2 and A.3, we deduce our results, by employing matrix perturbation theory as the major tool.) In section 3, some simulation results are presented to testify to the correctness of the denoising capacity. Likewise, in section 4, we present some simulations illustrating our results on the performance of the LSA-based recognition system. In section 5, using the theories in this paper, we re-examine some phenomena, observed by other researchers.

2. Major results

In this section we state our major results. Sections 3 and 4 outline the justification and discuss the importance of these results. The proofs are deferred to the appendices.

Result 1 (*Denoising capacity of SVD*): Suppose a matrix $\mathbf{A} \in R^{m,n}$ lies in a low-dimensional, r , subspace. It is corrupted by i.i.d. Gaussian noise producing another matrix \mathbf{B} , which is directly observed. Then, the error that still resides in the rank- r approximation matrix, \mathbf{B}^r , is

$$E |b_{i,j}^r - a_{i,j}| = \sigma \sqrt{\frac{r(m+n) - r^2}{mn}} \quad (1)$$

if the noise level σ , compared with the signal level, is small enough. Specially, as $m, n \rightarrow \infty$, the rank- r approximation of \mathbf{B} approaches \mathbf{A} , i.e. $\mathbf{B}^r \rightarrow \mathbf{A}$; and if $n \equiv k$ ($k \geq r$) and $m \rightarrow \infty$,

$$E |b_{i,j}^r - a_{i,j}| \rightarrow \sigma \sqrt{\frac{r}{k}} \quad (2)$$

Formulas (1) and (2) characterize the *denoising capacity* of the SVD, in terms of the noise level, the sizes of the measurement matrix and the rank. This result, without proof, was used in (Chen and Suter 2004a) to design a heuristic for SFM.

The proof of this result can be found in appendix A.2. A sketch of the ideas behind the proof is as follows. Standard perturbation results (see Theorems 1 and 2 in the appendix A.1.3) state that for small noise perturbations, the singular values are themselves perturbed by small noise terms. However, the same theorems state that the singular vectors are perturbed by an amount that depends upon the closeness of the singular values

- potentially a very large perturbation if the singular values are close to each other. However, we show that this result does not destroy the structure of the solution as the large potential perturbations, in the case of nearly equal singular values, merely correspond to possible permutations of the basis vectors of the space spanned by the singular vectors of these singular values. This allows us to derive the following expression for the discrepancy between the rank- r approximation of the noisy data and the clean data (see appendix A.2):

$$E \| \mathbf{B}^r - \mathbf{A} \|_F^2 = E \| \mathbf{Y} \|_F^2$$

$$\text{where } \mathbf{Y} = \begin{bmatrix} c_{1,1} & \cdots & c_{1,r} & c_{1,r+1} & \cdots & c_{1,k} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ c_{r,1} & \cdots & c_{r,r} & c_{r,r+1} & \cdots & c_{r,k} \\ c_{r+1,1} & \cdots & c_{r+1,r} & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ c_{m,1} & \cdots & c_{m,r} & 0 & \cdots & 0 \end{bmatrix}_{m,k} \quad (3)$$

and the non-zero c_{ij} are comparable in size (have the same expectation) as the entries in the original noise perturbation. The formula above holds to first order approximation. This means that the noise in the approximation has reduced degrees of freedom (due to the zero block in \mathbf{Y}) and the noise can be estimated as:

$$E \| \mathbf{B}^r - \mathbf{A} \|_F^2 = E \| \mathbf{Y} \|_F^2 = \sum E y_{i,j}^2 = (rm + rk - r^2) \sigma^2 \quad (4)$$

$$E | b_{i,j}^r - a_{i,j} | = \sigma \sqrt{\frac{rm + rk - r^2}{mk}} \quad (5)$$

Actually, there are second and higher order terms in the zero-block of \mathbf{Y} : they are not exactly zeroes. However, as the noise level is much less than the signal level; these

terms are much smaller compared with the terms denoted by $c_{i,j}$ in the above formula.

We can support to this conclusion by the following simple example in Matlab. Suppose

$\mathbf{A} \in R^{10,10}$ is a signal matrix, with all zero entries except $a_{1,1} = a_{2,2} = a_{3,3} = 100$.

$\mathbf{C} \in R^{10,10}$ is a noise matrix, generated as *randn*(10) in Matlab. The observed signal matrix is $\mathbf{B} = \mathbf{A} + \mathbf{C}$. The matrix in table 1 is a typical example of the approximation error,

$\mathbf{B}^3 - \mathbf{A}$, between \mathbf{A} and the rank 3 approximation of \mathbf{B} , \mathbf{B}^3 . Please note the zero-block in \mathbf{Y} (highlighted) is much smaller than other entries, although they are not exactly zeroes.

Table 1: A typical example of the approximation error, $\mathbf{B}^3 - \mathbf{A}$.

1.188	1.1857	-0.1289	-0.823	0.337	0.1243	-0.4911	-1.0207	0.467	-1.5508
2.2023	1.0554	-0.6569	-0.2284	0.8511	0.0955	0.8644	1.5689	-0.3891	-0.6872
-0.9865	-1.4727	1.1689	1.0365	-0.5205	-0.5185	-0.0025	0.063	-0.714	-1.9757
0.5235	-0.0745	0.4441	-0.0005	0.0012	0.0018	0.0032	0.0062	0.0005	0.0165
-0.3597	1.1914	0.2415	-0.0025	-0.0078	0.0005	-0.012	-0.0224	0.0079	0.0076
-0.2397	0.0212	1.2484	-0.0152	0.0074	0.0069	-0.0012	-0.0034	0.0102	0.021
-0.0112	1.1295	1.3475	-0.0115	-0.0025	0.006	-0.0097	-0.0184	0.014	0.0345
1.0295	1.3365	-0.9429	0.0221	-0.0204	-0.0077	-0.0066	-0.01	-0.0066	0.0072
0.9273	0.255	-0.0308	0.0087	-0.0057	-0.0016	0.0023	0.0054	-0.0036	0.0158
0.3838	-0.9499	0.677	-0.0064	0.0106	0.0041	0.0101	0.0184	-0.0005	0.0126

Further examples, supporting the theory, can be found in the following section 3.

Result 2 (Learning capacity of LSA): For a rank- r LSA-based recognition system, the "error measure" (the SSD) comes from two independent sources: the noise in the basis images and the noise in the test image. That is, because no data is free of measurement errors, the learning phase will not learn the "true" or ideal subspace. Moreover, the test

data will also contain noise – so a given test image will not lie in the learned subspace as both the test image itself, and the learned subspace, are perturbed by noise. Formally, suppose n m -dimensional images are presented as learning samples, and they are arranged as the matrix \mathbf{A} , $\mathbf{A} \in R^{m,n}$. Suppose, further, the noise-free image columns are restricted to a dimension r subspace. According to SVD decomposition, $\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^T$, where $\mathbf{U} \in R^{m,r}$, $\mathbf{V} \in R^{n,r}$, $\mathbf{U}^T\mathbf{U} = \mathbf{I}$, $\mathbf{V}^T\mathbf{V} = \mathbf{I}$ and $\mathbf{S} = \text{diag}\{\lambda_1, \dots, \lambda_r\}$. Because of the noise in the learning samples, a perturbed r -dimension subspace \mathbf{U}'' is obtained, rather than \mathbf{U} . The test image, \mathbf{q} , $\mathbf{q} = \mathbf{U}[f_1, \dots, f_r]$ is also corrupted with noise, observing \mathbf{p} .

Then, the projection error of the noisy test image, \mathbf{p} , upon the perturbed subspace \mathbf{U}'' , is:

$$\left\| \mathbf{p} - \mathbf{U}'' \mathbf{U}''^T \mathbf{p} \right\|_F^2 \cong (m-r)\sigma_t^2 + (m-r)\sigma_l^2 \sum_{i=1}^r \frac{f_i^2}{\lambda_i^2} \quad (6)$$

where σ_t and σ_l are the noise levels of the test image and the learning samples respectively.

Specifically, if \mathbf{q} is taken from the learning samples, i.e., \mathbf{q} is one of the columns of \mathbf{A} , with uniform probability of selecting the column number, the expectation of the SSD in (6), is:

$$(m-r)\sigma_t^2 + (m-r)r\sigma_l^2 / n \quad (7)$$

Conversely, if the expectation of the test image \mathbf{q} is known, i.e., $\{f_i \mid i = 1, \dots, r\}$ is known, the optimal learning set have the following property:

$$\frac{f_i}{\lambda_i^2} \equiv \text{Cons} \quad (8)$$

Moreover, for a *random* test image, (7) is *optimal* among the size- n learning sets (we cannot obtain lower error than this on average); and the size- n learning set is optimal *iff* it has r equal singular values. Here, the *randomness* of a test image means that $f_i = f_j$ because we do not have any knowledge about the test image. Note that the derivation of (6), (7) and (8) is arranged in appendices A.3.2, A.3.3 and A.3.4, respectively.

Formula (7) explicitly computes the reconstruction error when approximating a vector by the linear combination of the basis vectors obtained by SVD. It tells how to effectively control the approximation error. For example, we intuitively know that we can reduce the reconstruction error by reducing the noise levels in the test image and learning samples, however, it is not clear how these two noise sources affect the overall error. According to formula (7), these two sources of noise have different effects on the reconstruction error: the effect of the noise in the learning samples can be controlled by the size of the learning samples, while the noise in the test image works, independent of the learning process. Result (7) was used, without proof, in (Chen and Suter 2004b; Chen and Suter 2004c) to design improvements to PCA based face recognition systems.

3. Denoising capacity of SVD

From result 1, we can now draw a significant conclusion: As $m \rightarrow \infty$, while k is a constant, $E |B_{i,j}^r - A_{i,j}| \rightarrow \sigma \sqrt{\frac{r}{k}}$, a non-zero constant. In contrast, if both $m \rightarrow \infty$ and $n \rightarrow \infty$, the theory shows that it is possible to continue decreasing the noise in the data all the way down to zero.

One of the consequences, in the context of SFM is that it is impossible to reconstruct a 3D scene to arbitrary accuracy by the factorization method (using an affine camera model), by increasing the number of the frames (while keeping the number of the feature points unchanged). Indeed, using our theory we can actually put a lower bound on the noise remaining in such a situation (see the simulations late in this section).

This consequence contrasts with the claim that 3D scene could be reconstructed to arbitrary accuracy given enough frames (Thomas and Oliensis 1999). However, we recognize the need for caution - our setting is not exactly the same as that of (Thomas and Oliensis 1999), where the perspective model was adopted. Perspective factorization approaches are not simple to characterize/analyze and are beyond the scope of this paper.

Another similar consequence follows in formulations of rank constrained homography estimation where one dimension is inherently fixed to be 9. There is a limit to the denoising capacity in such a situation and, again, we can quantify this lower bound.

We note that it is very difficult to have real data with high precision ground truth. Thus, in this section, we present some *simulations* to verify result 1 of section 2.

In this experiment we work with a set of rank-3 matrices. For square matrices, the size of the matrices increases from 3 to 200; while for rectangular matrices, the number of the columns remains unchanged, staying at 40. The noise level is 0.1. Figure 1 shows the simulation results of SVD's denoising performance, compared with the expectation from result 1. It can be easily observed that the expected curve almost coincides with the simulation result. The curves of the simulation result and the expectation from (1) are so close that it is even difficult to separate them in Figure 1. Their closeness to each other confirms the theory about the denoising capacity of the SVD. In contrast with Figure 1 (d-f) (rectangular matrices), the curves in Figure 1 (a-c) (square matrices) can be observed to continue towards zero error, while the error for the rectangular matrices changes little after the number of the rows increases to 20 or 40.

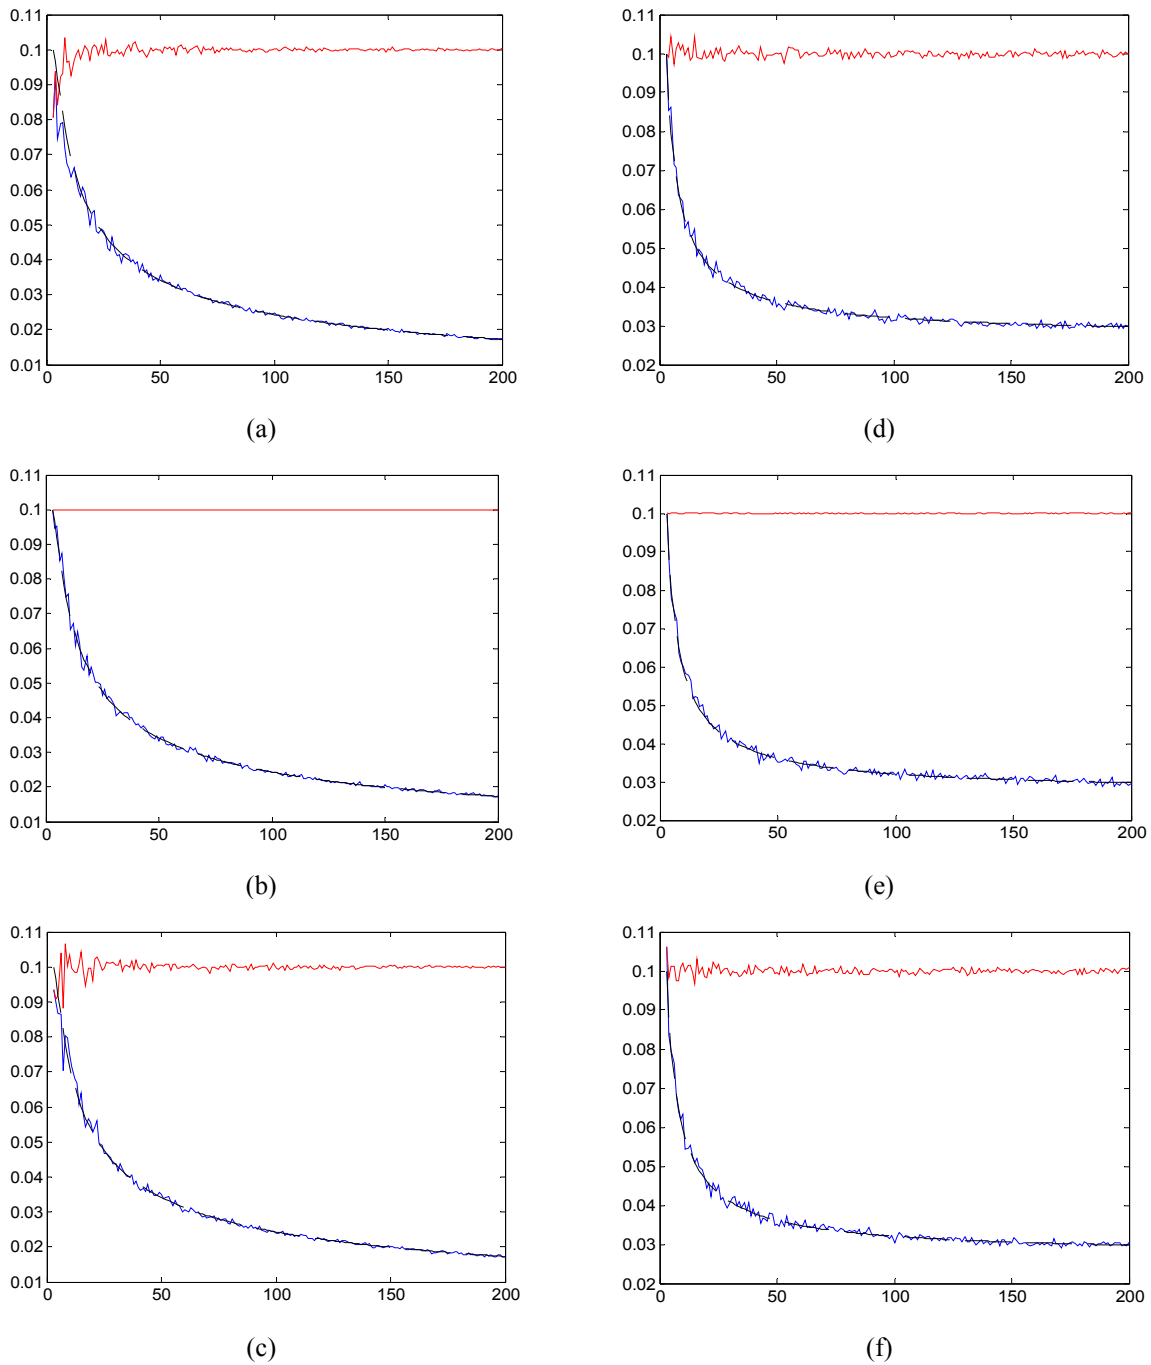


Figure 1: The average error that still resides in the rank-projected matrix. The abscissa denotes the number of the rows of the matrices, and the error is on the ordinate. (a-c) are for the square matrices, and (d-f) are for the rectangular matrices (fixed 40 columns). There are three curves in every sub-figure: the (approximately) straight curve in the upper part denotes the original noise in the noise corrupted matrix, and the smooth/non-smooth curves are the expectation/actual error in the approximation matrix respectively. In (a) and (d), the signal and the noise are randomly generated. In (b) and (e), the noise levels are normalized, so that the average energy in each entry of the matrices is 0.01. In (c) and (f), the signal matrices have 3 equal singular values.

3.1. Application in refining homography estimation

In (Chen and Suter 2004a), we extended the denoising capacity of a complete matrix to an incomplete matrix with missing data. By this, we provided a method to recover the most reliable imputation, in terms of deciding when the inclusion of extra rows or columns, containing significant numbers of missing entries, is likely to lead to poor recovery of the missing parts. In this section, we apply the theory of *denoising capacity* of a large low-rank matrix to refine the estimated homographies over two views. For background on projective geometry and homographies we refer the reader to the textbook (Hartley and Zisserman 2000).

Rank-4 constraint of the induced homographies

In (Shashua and Avidan 1996), it has been proved that the induced homography matrix over two views is embedded in a 4-dimensional subspace. Suppose that n ($n \geq 4$) planes are observed over two views. For the i^{th} plane, there exists a homography \mathbf{H}_i :

$$\mathbf{H}_i = \begin{bmatrix} h_{1,i} & h_{2,i} & h_{3,i} \\ h_{4,i} & h_{5,i} & h_{6,i} \\ h_{7,i} & h_{8,i} & h_{9,i} \end{bmatrix}, \text{ which projects the points of the first view upon the second view.}$$

The homography \mathbf{H}_i can be estimated from ≥ 4 points (Hartley and Zisserman 2000).

We arrange the induced vector $\mathbf{h}_i = [h_{1,i}, h_{2,i}, \dots, h_{9,i}]^T$ as the i^{th} column of \mathbf{H} . The fact that $\text{rank}(\mathbf{H}) = 4$, as is one of the major results in (Shashua and Avidan 1996). Figure 2 shows such an example, where the induced noise-free collection of homographies \mathbf{H} exactly has a rank of 4 (shown as the dashed curve). In practical circumstances, the homography will be estimated from noisy points and the collection of such homographies quickly becomes full-rank, as shown by the solid curve.




Figure 2: Rank-4 constraint for the induced matrix of from multiple-plane-over-two-view homographies. The dashed curve denotes for the 9 singular values of a noise-free homography matrix; and the solid for the same set of homographies with added noise.

The matrix \mathbf{H} has a constant width of $\mathfrak{9}$ (Shashua and Avidan 1996). According to the denoising capacity (2), if there are enough planes observed, the estimated homographies can be refined by the rank-4 constraint: In the limiting case, the error in the homographies can be reduced from σ to $\sigma\sqrt{\frac{4}{9}}$, where σ is the error level in estimated homographies.

Note that the error in the estimated homography usually loses the i.i.d. Gaussianity. Worse, the homography estimation suffers from a well-known heteroscedastic noise problem (Leedan and Meer 1999; Leedan and Meer 2000): The heteroscedastic noise arises due to the linearization in the direct linear transformation (DLT) algorithm (Hartley and Zisserman 2000).

As pointed out in appendix A.2, in deducing formula (3), we do not impose the i.i.d. Gaussianity upon the noise in the observed matrix. This means that the zero-block

of $-Y$ in (3) holds for *any* type of noise, if and only if the noise level is small enough compared with the signal level. In the following, we apply the rank-4 constraint to refine the estimated homographies.

Since we need to quantitatively assess our results in this setting, we resort to simulations (rather than using real data), where we know the ground truth for the convenience of performance evaluation.

Experimental settings

The first camera: We assume its center is placed at the coordinate origin and the x and y axes of its image plane coincide with the x and y axes of the coordinate system. The image plane is assumed to be placed between the camera center and the objects. Without loss of generality, its focal length is assumed f , which will be determined in the normalization. Assuming a pinhole model, the projection matrix of the first camera is $\mathbf{P}_1 = [\text{diag}(fI_2, 1), \mathbf{0}]$.

The second camera is randomly placed in a cube of size 2, and it rotates around z axis by a random angle.

The planes are determined in the following way: Each of them passes a point on z axis: $(0, 0, p)$, where p is a random value in $[35, 45]$. The angles between z axis and the planes are 70~90 degrees. The points in the planes are randomly restricted to $[-10:10, -10:10, ?]$, where the question symbol “?” means that the z -coordinate is determined by the plane, and the x and y coordinates.

Consideration of the projection error and normalization: If we only need to testify to the correctness of formula (2), the focal length of f can be arbitrarily set to 1. Here, we want to show that the refined homographies, improved by imposing the rank-4

constraint, have better performance, measured by the projection error between the noise-free feature points in the second view and the corresponding projections from the first view.

Note, the coefficients of the homography \mathbf{H}_i , projecting (x, y) of the first view upon the second view (x', y') involves parameters both in a numerator and in a

denominator: $(x', y') = \left(\frac{h_{1,i}x + h_{2,i}y + h_{3,i}}{h_{7,i}x + h_{8,i}y + h_{9,i}}, \frac{h_{4,i}x + h_{5,i}y + h_{6,i}}{h_{7,i}x + h_{8,i}y + h_{9,i}} \right)$. The worst case noise

amplification is likely to occur when a small error in $\{h_{7,i}, h_{8,i}, h_{9,i}\}$ leads to a comparatively large error in the projection of x' or y' . Thus, we set f (the focal lengths of the cameras) so that the average distance of all the points in all planes is $\sqrt{2}$, which is a normalization process not dissimilar to what is traditionally done in the linear approach to fundamental matrix estimation..

Experimental results

First, we study the denoising capacity of the rank-4 homography matrix, and make comparisons with formula (1). We use 20 points for each plane. By the SVD (Golub and Loan 1996), we calculate the rank-4 approximation of \mathbf{H} , whose i^{th} column, \mathbf{h}_i , is the induced homography vector for the i^{th} plane. Suppose the underlying ground truth of \mathbf{H} is $\tilde{\mathbf{H}}$, which is available in the simulation. What interests us is the ratio of r :

$$r = \frac{\|\tilde{\mathbf{H}} - \mathbf{H}^4\|_F}{\|\tilde{\mathbf{H}} - \mathbf{H}\|_F} \quad (9)$$

which reflects the denoising capacity of a large low-rank matrix. In order to make the comparison meaningful, we normalize each column of \mathbf{H} , i.e., $\|\mathbf{h}_i\|=1$ and make $h_{1,i}$

positive (if $h_{1,i} = 0$, we make another coefficient positive). The same procedure is applied to $\tilde{\mathbf{H}}$ and \mathbf{H}^4 . However, improvement in estimating the *coefficients* of a homography is perhaps not as directly observable as the error in the *projections* defined by the homography. For this reason we also report the average error in the projections of randomly chosen positions in the first image, when projected by the homographies into the second image. More specifically, we compare two errors e and e' : e is the projection error of the projection using the collection of homographies in \mathbf{H} , and e' is produced by the refined homographies in \mathbf{H}^4 . Here, we also use ratio of $\frac{e'}{e}$ to evaluate whether the refined homography improves the projection.

The solid/dotted curves in Figure 3 show the simulation result of r and $\frac{e'}{e}$, respectively. Note, in Figure 3, each experiment is repeated 100 times, then the root of the mean square is computed as the final r or $\frac{e'}{e}$. Figure 4 shows the histogram of r and $\frac{e'}{e}$ for the case with 100 planes (of course, such a case is perhaps unrealistic in practice as there are unlikely to be so many planes of significant size in any two views). The experiment is repeated 1000 times to create the averages. Note, from (1), r should be 0.6831; while the root of the mean square, in Figure 4, is 0.6677, in good agreement with the prediction.

From Figure 3 and Figure 4, the rank-4 constraint can improve the accuracy of the estimated homographies, as measured by the *coefficients* in the homography and this also leads to a reduction in the projection error (although not by the same sort of margin as can be seen from the histogram in Figure 4.)




Figure 3: The simulation of the improvements of the homography coefficients and the projection accuracy, as the number of planes increases.

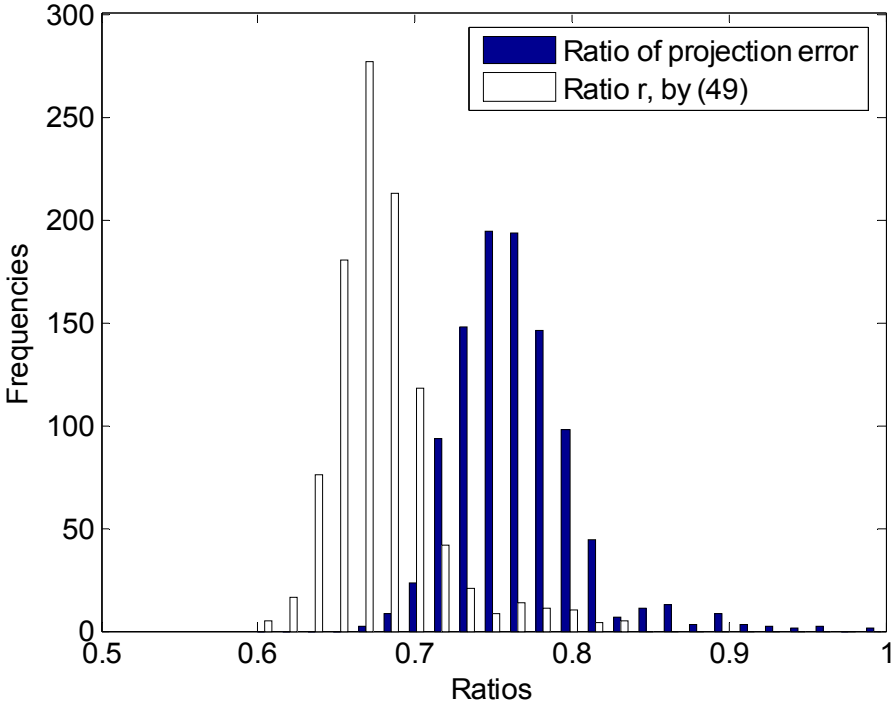


Figure 4: The histograms of r and the improvement of the projection error when there are 100 planes. The experiment was repeated 1000 times to calculate the averages.

4. Simulation of the *learning* capacity for LSA-based recognition

In this section, we present some simulation results illustrating our results on performance of the LSA-based recognition system, as stated in section 2.

In this example, the parameters in (7) are set as follows: $m=100$ (*size of the "images"*), $r=3$ (*dimension of the linear subspace spanning the learning set, before noise was added to the learning set*), $\sigma_s = 100$ (*signal strength – average sum of squares of the components of the "images" – both learning and testing*), and $\sigma_l = \sigma_t = 1$ (*noise variance added to the learning images and the test images, respectively*). In more detail, to generate the learning set we simply generated a number of image vectors (of length 100) that are restricted to a dimension-3 subspace, and scaled so their norm is 100. Noise was then added to these vectors to form the learning vectors.

For testing images/vectors we used two sets. The first set was essentially the original learning set (taking the noise-free learning vectors but adding different noise samples, with the same variance as the learning set, to these vectors). The second set was obtained in a similar fashion (generating random vectors in the same rank-3 subspace) but prior to adding the noise we equalized the 3 singular values of this test set.


As expected by our theory, as the learning sample size approaches infinite, the SSD on these two sets approaches a limiting (non-zero) value, as shown in *Figure 5-a*. On the test set that is just the learning set (with different noise samples), the performance (solid curve) almost coincides with the expectation from the theory (dashed curve). Because the 3 singular values of the second test set have been artificially equalized, the best performance on this set can be obtained only if the learning set has 3 equal singular values, from (8). However, the learning set always has 3 distinct singular values. Thus,

the performance on the random test set is worse than the optimal curve (dashed curve), especially for the small-size learning samples. (In fact, the performance for the recognition system is very bad, at 5,771.6, 788.1 and 588.1 respectively, when the learning sample sizes are only 3, 4 and 5. In order to make the curves clear, these points have been omitted in *Figure 5-a*.)


In the next series of experiments we create a single random test set and investigate the performance of different learning sets (different basis images): an optimal learning set (complying with (8)) and a random learning set (with 3 equal singular values). The performance of the random learning set, denoted by the solid curve, can be expected to coincide with the expectation (7), denoted by the dashed curve, as shown in *Figure 5-b*. From *Figure 5-b*, the optimal learning set has a better performance than the random learning set, especially for small learning sizes. Note, if the learning set is *truly* randomly generated, it probably has a very bad SSD performance, especially for a small-size learning set. For example, the r^{th} basis image is comparatively *dirty*, because the r^{th} singular value of the learning set is comparatively small; while most of the energy of the test image comes from this basis image. In such cases, the error from the basis images, especially from the r^{th} basis image, will dominate the total error, as can be seen from (6).

In Figure 6, we show the effects of the three parameters in (7), the size of the learning samples (n), the noise level in the learning set (σ_l), and the noise level in the test set (σ_t). In this example, the recognition system works on the *learning* samples. Figure 6 (a) shows the performance of SSD when σ_t is 0.5. It can be easily observed: the *square* dependency on σ_l in the small-size learning sets and the decreasing effects of σ_l as the

learning size increases. Figure 6 (b) shows the performance of SSD when σ_l is 0.5. We can observe the *square* dependency on σ_l and also that this effect is almost independent of the learning size. Figure 6 (c) and (d) show the effects of σ_l and σ_t when the learning sizes are 3 and 125 respectively. When the learning size is 3, σ_l has almost a same effect on SSD as σ_t . When the learning size is 125 ($\gg 3$), the noise in learning set can be almost neglected if it is not much stronger than the noise in the test set. This result readily follows from equation (7).

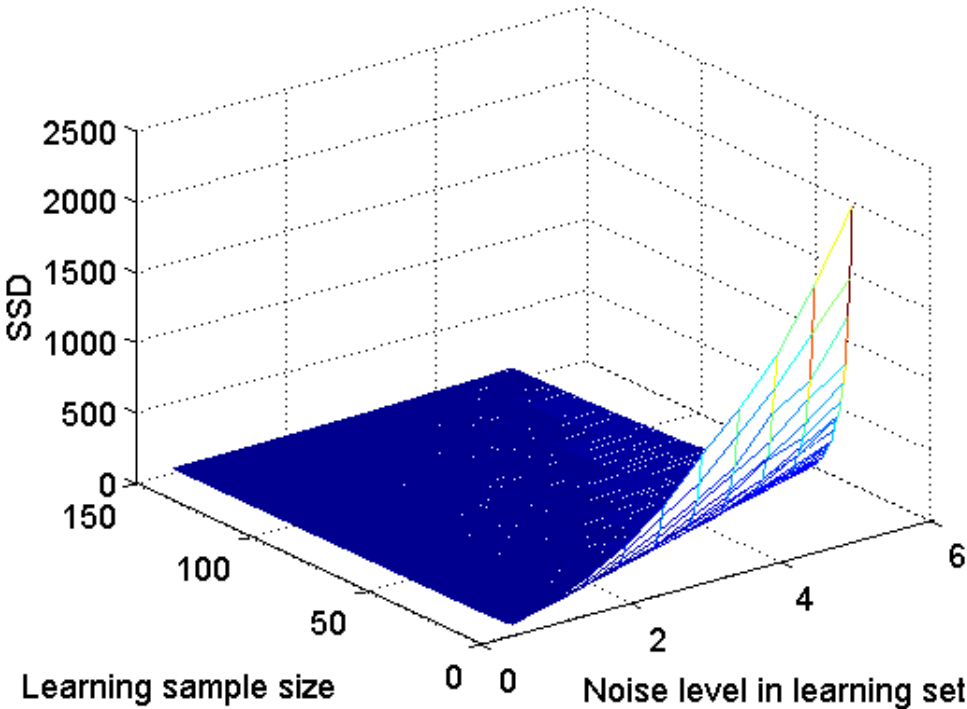


(a)




(b)


Figure 5, The dependency of SSD on the size of learning samples. (a) for a single learning set, tested with two test sets: (solid) the learning set from which the basis images are obtained (both with different samples of noise to the actual learning set), and (dotted) another set from the same subspace as the learning set but one that has 3 equal singular values (and, again, different noise samples); (b) for a single test set, tested by two learning sets: (dotted) the optimal learning set and (solid) another learning set that has 3 equal singular values. In both subfigures, the dashed curves denote the expectation from (7).



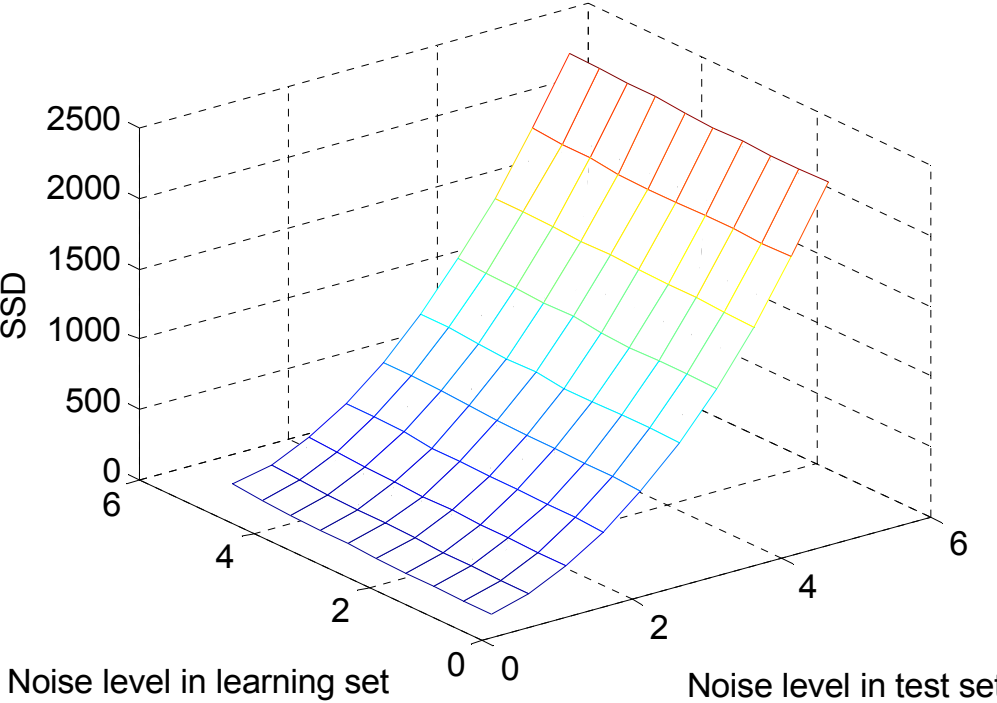
(a)



(b)



(c)



(d)

Figure 6: The effects of the three parameters in (7) on SSD. For details, see the description in the text.

5. A re-examination of some existing results

In this paper, we have presented and proved theorems regarding the noise related behavior of SVD important classes of problems. We have supported the theorems with simulations. However, we can do more than this. We can re-examine some of the existing results, reported in the literature, in the light of the new results. For example, in SFM, the root-mean-square error of the recovered shape with respect to the true shape was reported in (Morita and Kanade 1997). Fig. 6 in that paper shows that the error approaches a constant value after the number of the frames increases to 20 or 40, and the observation is of the same form as the results we reported here.

In terms of LSA, we examine the results in (Basri and Jacobs 2003). There it was reported that no significant deterioration of the performance was found for LSA-based face recognition, if the images were subsampled by 16×16 squares. This subsampling results in a decrease, of the row number of \mathbf{A} (in our terminology), by $1/256$. However, the reduced row number is still very large, about 1000 ($\gg 4$ or 9). From (80): the performance, measured by the angle between the test image and the basis images, is almost independent of m if $m \gg r$, which in essence predicts the observation in (Basri and Jacobs 2003).

The last example (maybe even the most important), we give, relates to the observation that “recognition of an object under a particular lighting and pose can be performed reliably provided the object has been previously seen under similar circumstances” (Georghiades *et al.* 2001). An explanation can be found from (6) and (8). For a test image, if it or similar cases have been observed in the learning samples, its $\{f_i^2\}$ will tend to have a monotonic relationship with $\{\lambda_i^2\}$, i.e., for a larger λ_i^2 , f_i^2 is

also larger, and vice versa. More formally, if (8) holds, the recognition system has the optimal performance. However, a test image may be produced under very different lighting conditions from those in the learning set. In such situations, it is quite possible that much of the energy of the test image concentrates in one or more of the dimensions, associated with small singular values, i.e., the dimensions that are most easily perturbed by noise. From (6), the recognition error would then be very large. This not only explains the drawback of PCA-based face recognition, pointed out in (Georghiades *et al.* 2001), but also gives a possible solution, as suggested by (8). For a *random* test set, the best learning samples should be selected this way: to equalize the first r largest singular values as possible. However, we do not present any specific strategies for this *open*, and perhaps promising, issue.

Finally, we note that (Chen and Suter 2004b; Chen and Suter 2004c) and in (Chen and Suter 2004a) provide further examples of the utility of subspace noise analysis: addressing SFM and Face Recognition respectively. In those papers, considerable space and detail was devoted to particular problems that need addressing (e.g. "missing data" in SFM and outliers in Face Recognition) in a more complete solution to particular computer vision problems. However, the fundamental issue behind the success of those methods was the careful exploitation of our subspace noise analysis.

6. Conclusion

The main contribution of this paper is to present a theoretical analysis of SVD-based low rank projections: the *denoising* capacity of SVD (where we characterized the error that still resides in the *SVD-denoised* matrix) and the *learning* capacity of LSA-based recognition systems (where we showed that the projection error can be decomposed into two independent sources, one from the test image and the other from the basis image). In doing so, a minor contribution is that we fill an apparent gap in the literature: the perturbation theory concerning multiple eigenvalues (singular values).

The theory was supported by a number of experiments presented in this paper, as well as a number of observations that interpreted the results from previous papers in the light of this theory. Though we concentrated here on a few settings: refinement of the estimation of a group of related homographies, SFM, and face recognition, the theory is quite general and could well find more applications in computer vision and pattern recognition tasks.

Appendix A.1: Preliminary knowledge: SVD and perturbation theory

A.1.1 Notation

In the following, a matrix will be denoted by a bold capital letter, like \mathbf{M} , and a bold lowercase letter represents a vector, e.g. \mathbf{x} . The i^{th} column of \mathbf{M} is denoted by \mathbf{m}_i . A scalar entry in a vector or in a matrix will respectively be denoted by, for example, x_1 or $m_{1,2}$. In some cases, we use $[\mathbf{M}]_{m,n}$ to specify the sizes of \mathbf{M} for clarity, i.e., $\mathbf{M} \in R^{m,n}$. $\mathbf{M}_{i:j,k:l}$, a notation from *Matlab*, denotes for the submatrix of \mathbf{M} : the intersection of the i -to- j rows and the k -to- l columns. \mathbf{I}_n denotes the $n \times n$ identity matrix, and $\mathbf{0}_{m,n}$ for a $m \times n$ zero-matrix. \mathbf{e}_i is the i^{th} column of \mathbf{I}_n . A matrix \mathbf{U} , $\mathbf{U} \in R^{m,n}$, is said to have orthonormal columns, *iff* $\mathbf{U}^T \mathbf{U} = \mathbf{I}_n$. The set of such $m \times n$ matrices, which are always represented by \mathbf{U} or \mathbf{V} , is denoted by $O^{m,n}$. A square matrix \mathbf{U} is orthogonal *iff* $\mathbf{U} \in O^{m,m}$. Two matrices, \mathbf{M} and \mathbf{N} , with same sizes, are said to be orthonormal to each other *iff* $\|\mathbf{M}\|_F = 1$, $\|\mathbf{N}\|_F = 1$, and $\sum m_{i,j} n_{i,j} = 0$. The Frobenius norm of a matrix \mathbf{M} (or a vector) will be denoted as $\|\mathbf{M}\|_F$, where $\|\mathbf{M}\|_F = \sqrt{\sum_{i,j} m_{i,j}^2}$. \mathbf{M}^r denotes the closest rank- r approximation of $\mathbf{M} \in R^{m,n}$, where $r \leq \min(m,n)$, as will be explained in appendix A.1.2. Similarly, the entry in \mathbf{M}^r is denoted by $m_{i,j}^r$. The symbol “ \approx ” means the first order perturbation, explained in appendix A.1.3. And, “ \cong ” means the equality, in the sense of statistical expectation.

A.1.2 Singular value decomposition

The principle behind the SVD (Golub and Loan 1996) states that any matrix, $\mathbf{M} \in R^{m,n}$, can be decomposed into

$$\mathbf{M} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T \quad (10)$$

where $\mathbf{U} \in O^{m,m}$, $\mathbf{V} \in O^{n,n}$ and $\mathbf{\Sigma} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_p) \in R^{m,n}$, with $p = \min(m, n)$ and $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p \geq 0$. Without loss of generality, suppose $m \geq n$. $\{\lambda_i^2 \mid i=1,2,\dots,n\}$ are the eigenvalues of $\mathbf{M}^T\mathbf{M}$, or the first n largest eigenvalues of $\mathbf{M}\mathbf{M}^T$. The first n left singular vectors of \mathbf{M} are $\{\mathbf{u}_i \mid i=1,2,\dots,n\}$, where \mathbf{u}_i is the eigenvector, associated with the eigenvalue of λ_i^2 , of $\mathbf{M}\mathbf{M}^T$. Similarly, the right singular vectors of \mathbf{M} are $\{\mathbf{v}_i \mid i=1,2,\dots,n\}$, where \mathbf{v}_i is the eigenvector, associated with the eigenvalue of λ_i^2 , of $\mathbf{M}^T\mathbf{M}$. Another important fact (Golub and Loan 1996) is that one can easily construct \mathbf{M}^r , the closest rank r approximation of \mathbf{M} measured by 2-norm or Frobenius-norm, by:

$$\mathbf{M}^r = \sum_{i=1}^r \lambda_i \mathbf{u}_i \mathbf{v}_i^T \quad (11)$$

Specifically,

$$\|\mathbf{M} - \mathbf{M}^r\|_2 = \lambda_{r+1} \quad (12)$$

$$\|\mathbf{M} - \mathbf{M}^r\|_F = \sqrt{\sum_{j=r+1}^n \lambda_j^2} \quad (13)$$

A.1.3 Perturbation theory

Only the perturbation theory concerning singular values/vectors is needed in this paper. However, we also include the perturbation theory for the eigenvalues/eigenvectors as a useful way to arrive at our results. The proofs of theorems **1** and **3** help one

understand the proofs of theorems 2 and 4. To the best of our knowledge*, the perturbation expansion of the eigenvectors/singular-vectors is available only for those that correspond to a *simple* eigenvalue or singular value (Wilkinson 1965; Stewart and Sun 1990). In this section, we review such theory, and in the next section, we present our new results for those that correspond to a *multiple* eigenvalue or singular value. In order to have a complete description of the perturbation theory, we give detailed proofs, including those available in the textbooks (Wilkinson 1965; Stewart and Sun 1990). Importantly, we provide a characterization of “message” conveyed by these theorems.

Theorem 1 (Wilkinson 1965): *Consider a symmetric matrix, $\mathbf{M} \in R^{m,m}$. Suppose \mathbf{M} has m distinct eigenvalues, $\{\lambda_i | i=1,2,\dots,m\}$ and the associated eigenvectors are $\{\mathbf{u}_i | i=1,2,\dots,m\}$. Matrix \mathbf{M} is perturbed by a matrix $\Delta\mathbf{M}$ so we observe $\mathbf{A} = \mathbf{M} + \Delta\mathbf{M}$, and the eigenvalues and eigenvectors of \mathbf{A} are $\{\lambda'_i | i=1,2,\dots,m\}$ and $\{\mathbf{u}'_i | i=1,2,\dots,m\}$ respectively. Assuming every entry in $\Delta\mathbf{M}$ is small enough, the first-order perturbed** eigenvalues and eigenvectors are:*

$$\lambda'_i = \lambda_i + c_{i,i} \tag{14}$$

$$\mathbf{u}'_i = \mathbf{u}_i + \sum_{j \neq i} \frac{c_{j,i}}{\lambda_i - \lambda_j} \mathbf{u}_j \tag{15}$$

where $\mathbf{C} = \mathbf{U}^T \Delta\mathbf{M} \mathbf{U}$.

Proof:

* Here, we'd like to express our appreciation to Prof. G. W. Stewart [32], who, by private correspondence, pointed this out to us.

** In the following, we omit the “first-order” when we refer to “first-order perturbed” if the omission does not create any ambiguity.

If \mathbf{M} is corrupted with $\Delta\mathbf{M}$ and $\Delta\mathbf{M}$ is small enough, the first-order perturbations of the eigenvalue and the eigenvector, denoted as $\Delta\lambda$ and $\Delta\mathbf{u}$ respectively, will be small enough, from Ostrowski's continuity theorem (Wilkinson 1965). Denote higher-order terms as $\delta\lambda$ and $\delta\mathbf{u}$, respectively. From $\mathbf{A}\mathbf{u}'_i = (\mathbf{M} + \Delta\mathbf{M})(\mathbf{u} + \Delta\mathbf{u} + \delta\mathbf{u}) = (\lambda + \Delta\lambda + \delta\lambda)(\mathbf{u} + \Delta\mathbf{u} + \delta\mathbf{u})$, we have the first-order perturbation, by dropping the higher-order terms:

$$\mathbf{M} \cdot \mathbf{u} + \mathbf{M} \cdot \Delta\mathbf{u} + \Delta\mathbf{M} \cdot \mathbf{u} \approx \lambda \cdot \mathbf{u} + \lambda \cdot \Delta\mathbf{u} + \Delta\lambda \cdot \mathbf{u} \quad (16)$$

We can write $\mathbf{u}'_i = \mathbf{u}_i + \sum_{j \neq i} \Delta f_{j,i} \mathbf{u}_j$ (since these vectors form a basis), and $\lambda'_i = \lambda_i + \Delta\lambda_i$.

Our task is to find an expression for $\Delta f_{j,i}$ and $\Delta\lambda_i$. From the first-order perturbation, we

have $\mathbf{M}\mathbf{u}_i + \mathbf{M} \sum_{j \neq i} \Delta f_{j,i} \mathbf{u}_j + \Delta\mathbf{M}\mathbf{u}_i \approx \lambda_i \mathbf{u}_i + \lambda_i \sum_{j \neq i} \Delta f_{j,i} \mathbf{u}_j + \Delta\lambda_i \mathbf{u}_i$, and

$$\sum_{j \neq i} \Delta f_{j,i} (\lambda_j - \lambda_i) \mathbf{u}_j + \Delta\mathbf{M}\mathbf{u}_i = \Delta\lambda_i \mathbf{u}_i \quad (17)$$

Because \mathbf{M} is symmetric and has m distinct eigenvalues, $\{\mathbf{u}_i\}$ are orthogonal to each other. Pre-multiplying (17) by \mathbf{u}_i^T , we obtain $\Delta\lambda_i = \mathbf{u}_i^T \Delta\mathbf{M}\mathbf{u}_i = c_{i,i}$. Pre-multiplying (17) by

\mathbf{u}_j^T , we have $\Delta f_{j,i} = \frac{c_{j,i}}{\lambda_i - \lambda_j}$. □

Notice that the theorem essentially states that the eigenvalues are, to first order, perturbed by the diagonal elements of \mathbf{C} (which is the “noise” transformed by the eigenvectors of the data). If the “noise” has “no structure”, these disturbances will be small. Likewise, the eigenvectors are perturbed by a “mixing” with the other eigenvectors, with the important

message that the degree of mixing is controlled by the closeness of the respective eigenvalues, in addition to the size and structure of the “noise”.

Theorem 2 (Stewart and Sun 1990): *Suppose \mathbf{A} (not necessarily symmetric) is corrupted with \mathbf{N} and we observe \mathbf{B} : $\mathbf{B} = \mathbf{A} + \mathbf{N}$. According to SVD, we have $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$, where $\mathbf{U} \in O^{m \times m}$, $\mathbf{\Sigma} = \text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_m\}$, $\mathbf{V} \in O^{m \times m}$. Define $\mathbf{C} = \mathbf{U}^T \mathbf{N} \mathbf{V}$. Suppose λ_i is a simple non-zero singular value of \mathbf{A} . Then, the perturbed singular value λ'_i , left singular vector \mathbf{u}'_i , and right singular vector \mathbf{v}'_i , of \mathbf{B} are respectively*

$$\lambda'_i = \lambda_i + c_{i,i} \quad (18)$$

$$\mathbf{u}'_i = \mathbf{u}_i + \sum_{j \neq i} \frac{\lambda_i c_{j,i} + \lambda_j c_{i,j}}{\lambda_i^2 - \lambda_j^2} \mathbf{u}_j \quad (19)$$

$$\mathbf{v}'_i = \mathbf{v}_i + \sum_{j \neq i} \frac{\lambda_j c_{j,i} + \lambda_i c_{i,j}}{\lambda_i^2 - \lambda_j^2} \mathbf{v}_j \quad (20)$$

The following Lemma will not only provide most of the proof but it will also be independent use:

Lemma 1 Define $\mathbf{D} = \mathbf{\Sigma} + \mathbf{C}$. Because $\mathbf{\Sigma}$ is a diagonal matrix, $\{\lambda_j\}$ and $\{\mathbf{e}_j\}$ are respectively the singular values and the right/left singular vectors of $\mathbf{\Sigma}$. First, focusing on \mathbf{D} , we prove its singular values, κ_i , right singular vectors \mathbf{x}_i , and left singular vectors \mathbf{y}_i are respectively

$$\kappa_i = \lambda_i + c_{i,i} \quad (21)$$

$$\mathbf{x}_i = \mathbf{e}_i + \sum_{j \neq i} \frac{\lambda_j c_{j,i} + \lambda_i c_{i,j}}{\lambda_i^2 - \lambda_j^2} \mathbf{e}_j \quad (22)$$

$$\mathbf{y}_i = \mathbf{e}_i + \sum_{j \neq i} \frac{\lambda_i c_{j,i} + \lambda_j c_{i,j}}{\lambda_i^2 - \lambda_j^2} \mathbf{e}_j \quad (23)$$

Proof of Lemma 1: Writing $\kappa_i = \lambda_i + \Delta\lambda_i$, $\mathbf{x}_i = \mathbf{e}_i + \sum_{j \neq i} \Delta f_{j,i} \mathbf{e}_j$, and

$\mathbf{y}_i = \mathbf{e}_i + \sum_{j \neq i} \Delta g_{j,i} \mathbf{e}_j$ the task is to find expressions for $\Delta f_{j,i}$, $\Delta g_{j,i}$, and $\Delta\lambda_i$

From our definition of \mathbf{D} , κ_i , \mathbf{x}_i and \mathbf{y}_i we have $\mathbf{D}\mathbf{x}_i = \kappa_i \mathbf{y}_i$ and $\mathbf{D}^T \mathbf{y}_i = \kappa_i \mathbf{x}_i$. Equating their first-order terms, we have:

$$\Sigma \mathbf{e}_i + \mathbf{C} \mathbf{e}_i + \Sigma \sum_{j \neq i} \Delta f_{j,i} \mathbf{e}_j \approx \lambda_i \mathbf{e}_i + \Delta\lambda_i \mathbf{e}_i + \lambda_i \sum_{j \neq i} \Delta g_{j,i} \mathbf{e}_j \quad (24)$$

$$\Sigma^T \mathbf{e}_i + \mathbf{C}^T \mathbf{e}_i + \Sigma^T \sum_{j \neq i} \Delta g_{j,i} \mathbf{e}_j \approx \lambda_i \mathbf{e}_i + \Delta\lambda_i \mathbf{e}_i + \lambda_i \sum_{j \neq i} \Delta f_{j,i} \mathbf{e}_j \quad (25)$$

Then

$$\mathbf{C} \mathbf{e}_i + \sum_{j \neq i} \lambda_j \Delta f_{j,i} \mathbf{e}_j = \Delta\lambda_i \mathbf{e}_i + \lambda_i \sum_{j \neq i} \Delta g_{j,i} \mathbf{e}_j \quad (26)$$

$$\mathbf{C}^T \mathbf{e}_i + \sum_{j \neq i} \lambda_j \Delta g_{j,i} \mathbf{e}_j = \Delta\lambda_i \mathbf{e}_i + \lambda_i \sum_{j \neq i} \Delta f_{j,i} \mathbf{e}_j \quad (27)$$

First, by equating the coefficients of \mathbf{e}_i in (26)-(27), we have $\Delta\lambda_i = c_{i,i}$. Likewise, from the coefficients of \mathbf{e}_j ($j \neq i$),

$$\begin{cases} \lambda_i \Delta g_{j,i} - \lambda_j \Delta f_{j,i} = c_{j,i} \\ -\lambda_j \Delta g_{j,i} + \lambda_i \Delta f_{j,i} = c_{i,j} \end{cases} \quad (28)$$

The solution of the two-by-two system of equations in (28) is:

$$\begin{cases} \Delta g_{j,i} = (\lambda_i c_{j,i} + \lambda_j c_{i,j}) / (\lambda_i^2 - \lambda_j^2) \\ \Delta f_{j,i} = (\lambda_j c_{j,i} + \lambda_i c_{i,j}) / (\lambda_i^2 - \lambda_j^2) \end{cases} \quad (29)$$

So far, the Lemma has been proved.

Proof of Theorem 2: Noting that $\mathbf{B} = \mathbf{A} + \mathbf{N} = \mathbf{U}\Sigma\mathbf{V}^T + \mathbf{UCV}^T = \mathbf{UDV}^T$ we have:

$$\mathbf{B} = \mathbf{UDV}^T \approx \mathbf{U}[\mathbf{y}_1, \dots, \mathbf{y}_m] \text{diag}\{\kappa_1, \dots, \kappa_m\} [\mathbf{x}_1, \dots, \mathbf{x}_m]^T \mathbf{V}^T \quad (30)$$

so that \mathbf{B} has κ_i , \mathbf{Vx}_i , and \mathbf{Uy}_i respectively as its singular values, right and left singular vectors. Thus, combining (30) with (21-23), we have proved the theorem. \square

The interpretation of this theorem is similar to that of theorem 1. The singular values are perturbed by an amount that will be small if the noise is small and “has no structure”. The perturbation of the singular vectors involves some form of mixing that will be small if the noise is small (and unstructured) and if the singular values are not very close to each other.

The above perturbation theorem concerning the singular values/vectors, holds only for positive (and significantly large) singular values (Stewart and Sun 1990) (Note: singular values have to be non-negative.) In this paper, we are concerned with perturbations of subspaces associated with the first r largest singular values, where $r \ll m$. Thus, we do not have to consider the behavior of the perturbation for the zero (or near zero) singular values.

Note that it is assumed that the noise level is much smaller than the signal level in this paper. Under this assumption, it is reasonable to assume that the perturbed singular values, $\{\lambda'_i | i = 1, \dots, r\}$ are still the r largest singular values of \mathbf{B} . However, as the noise level increases, this assumption can be violated. In these cases, the perturbed zero singular value can take place of the perturbed smallest nonzero singular value. Such an example can be found in (Faugeras and Luong 2001).

A.1.4 New perturbation theory, associated with a multiple eigenvalue/singular value

In this section, we present our perturbation expansions, corresponding to the case where the matrix has at least one multiple eigenvalue/singular value.

First, we want to shed some light on the perturbation expansions concerning singular vectors that correspond to a multiple singular value. We do this by considering the perturbation expansions of the eigenvectors of a symmetric square matrix:

Theorem 3: *Suppose $\mathbf{A} \in R^{m,m}$, $\mathbf{A} = \mathbf{A}^T$, and it has m eigenvalues $\{\lambda_i\}$ and m eigenvalues $\{\mathbf{u}_i\}$, which are orthonormal to each other*. Without loss of generality, suppose the first r eigenvalues of \mathbf{A} are same, $\lambda_i = \lambda$ for $i = 1, 2, \dots, r$. \mathbf{A} is corrupted with \mathbf{N} , which, compared with \mathbf{A} , is small enough. $\mathbf{B} = \mathbf{A} + \mathbf{N}$ is observed. For the reason of simplicity, assume \mathbf{N} is symmetric. Define $\mathbf{C} = \mathbf{U}^T \mathbf{N} \mathbf{U}$. Then, the first r perturbed eigenvalues and eigenvectors of \mathbf{B} are:*

$$\lambda'_i = \lambda + \delta_i \quad (31)$$

$$\mathbf{u}'_i = \sum_{j=1}^r \hat{u}_{j,i} \mathbf{u}_j + \sum_{j=r+1}^m \frac{c'_{j,i}}{\lambda - \lambda_j} \mathbf{u}_j \quad (32)$$

where $\mathbf{C}_{1r,1r} = \hat{\mathbf{U}} \text{diag}\{\delta_1, \dots, \delta_r\} \hat{\mathbf{U}}^T$ (supposing $\delta_i \neq \delta_j$ if $i \neq j$). Define

$\mathbf{C}' = \begin{bmatrix} \hat{\mathbf{U}}^T & \\ & \mathbf{I}_{m-r} \end{bmatrix} \mathbf{C} \begin{bmatrix} \hat{\mathbf{U}} \\ & \mathbf{I}_{m-r} \end{bmatrix}$. The other $m-r$ eigenvalues/eigenvectors can be obtained

as in theorem 1.

* For an r -ple multiple eigenvalue, we, first, have its r eigenvectors, $\{\mathbf{u}_i \mid i = 1, \dots, r\}$, which may not be orthogonal. Then, the r orthogonal eigenvectors can be obtained by applying Schmidt orthogonalization on $\{\mathbf{u}_i \mid i = 1, \dots, r\}$.

To understand the significance of this theorem, it is useful to consider what happens as the eigenvalue/eigenvector structure approaches the situation where we have multiple equal eigenvalues and a subspace of eigenvectors, of dimension >1 for that eigenvector. Theorem *I* becomes not useful in such a situation since the eigenvector perturbation clearly becomes divergent. However, the present theorem shows that the subspace structure is still boundedly perturbed. (In essence the mixing is between eigenvectors now spanning the subspace associated with the multiple eigenvalue - not outside of it. Put another way, theorem *I* predicts very large perturbations of the eigenvectors of very close eigenvalues but this large perturbation is essentially due to the alternate bases for the same subspace they span.)

Proof: From the perturbation theory about the eigenvectors associated with a multiple

eigenvalue (Wilkinson 1965), we can suppose $\mathbf{u}'_i = \sum_{j=1}^r p_{j,i} \mathbf{u}_j + \sum_{j=r+1}^m \Delta f_{j,i} \mathbf{u}_j$ and

$\lambda'_i = \lambda + \Delta \lambda_i$. Note: $p_{j,i}$ are different from $\Delta f_{j,i}$. $p_{j,i}$ can possibly take any value within $[-1,1]$, while $\Delta f_{j,i}$ approach zeroes if N is small enough.

$$(\mathbf{A} + \mathbf{N})\mathbf{u}'_i = \lambda'_i \mathbf{u}'_i \quad (33)$$

Equating the first-order terms of (33) produces:

$$\mathbf{A} \sum_{j=1}^r p_{j,i} \mathbf{u}_j + \mathbf{A} \sum_{j=r+1}^m \Delta f_{j,i} \mathbf{u}_j + \mathbf{N} \sum_{j=1}^r p_{j,i} \mathbf{u}_j \approx \lambda \sum_{j=1}^r p_{j,i} \mathbf{u}_j + \lambda \sum_{j=r+1}^m \Delta f_{j,i} \mathbf{u}_j + \Delta \lambda_i \sum_{j=1}^r p_{j,i} \mathbf{u}_j \quad (34)$$

Then

$$\sum_{j=r+1}^m \lambda_j \Delta f_{j,i} \mathbf{u}_j + [\mathbf{u}_1, \dots, \mathbf{u}_m] \mathbf{C}_{1:m,1:r} \mathbf{p}_i \approx \lambda \sum_{j=r+1}^m \Delta f_{j,i} \mathbf{u}_j + \Delta \lambda_i \sum_{j=1}^r p_{j,i} \mathbf{u}_j \quad (35)$$

where $\mathbf{p}_i = [p_{1,i}, p_{2,i}, \dots, p_{r,i}]^T$. Equating the coefficients of \mathbf{u}_j for $(j=1, \dots, r)$, we have

$$\mathbf{C}_{1r,1r}\mathbf{p}_i = \Delta\lambda_i\mathbf{p}_i \quad (36)$$

where $\mathbf{C}_{1r,1r}$ is the left-up $r \times r$ submatrix of \mathbf{C} . If $\mathbf{C}_{1r,1r}$ has r distinct eigenvalues, the solution of $\Delta\lambda_i$ and \mathbf{p}_i is unique, as (36). From (36), \mathbf{P} is same as $\hat{\mathbf{U}}$, as defined in the theorem. After substituting $\Delta\lambda_i$ and \mathbf{p}_i in (34), the equality of \mathbf{u}_j for $(j = r+1, \dots, m)$ produces the first order perturbations of $\Delta f_{j,i}$ as in the theorem. \square

Following the same notation as used in theorem 2, we now consider the perturbation expansion, where the matrix has at least one multiple *singular* value.

Theorem 4: \mathbf{A} , \mathbf{B} , \mathbf{C} and Σ are defined as those in theorem 2. Define $\mathbf{D} = \mathbf{C} + \Sigma$. Without loss of generality, suppose the first r singular values of \mathbf{A} are same: $\{\lambda_i = \lambda \mid i = 1, \dots, r\}$.

By SVD, $\mathbf{D}_{1r,1r} = \hat{\mathbf{U}}\text{diag}\{t_1, \dots, t_r\}\hat{\mathbf{V}}^T$. Let $\mathbf{U}'' = \begin{bmatrix} \hat{\mathbf{U}} & \\ & \mathbf{I}_{m-k} \end{bmatrix}$, $\mathbf{V}'' = \begin{bmatrix} \hat{\mathbf{V}} & \\ & \mathbf{I}_{m-k} \end{bmatrix}$, and

$$\mathbf{D}' = \mathbf{U}''^T \mathbf{D} \mathbf{V}''.$$

$$\mathbf{B} = (\mathbf{U}\mathbf{U}'')\mathbf{D}'(\mathbf{V}\mathbf{V}'')^T \quad (37)$$

The perturbed singular values, $\{\lambda'_i\}$, right singular vectors $\{\mathbf{v}'_i\}$, and left singular vectors $\{\mathbf{u}'_i\}$ for $1 \leq i \leq r$, of \mathbf{B} are respectively

$$\lambda'_i = d'_{i,i} = t_i \quad (38)$$

$$\mathbf{u}'_i = \sum_{j=1}^r \hat{u}_{j,i} \mathbf{u}_j + \sum_{j=r+1}^m \frac{\lambda d'_{j,i} + \lambda_j d'_{i,j}}{\lambda^2 - \lambda_j^2} \mathbf{u}_j \quad (39)$$

$$\mathbf{v}'_i = \sum_{j=1}^r \hat{v}_{j,i} \mathbf{v}_j + \sum_{j=r+1}^m \frac{\lambda_j d'_{j,i} + \lambda d'_{i,j}}{\lambda^2 - \lambda_j^2} \mathbf{v}_j \quad (40)$$

The perturbations, associated with other non-zero simple singular values, can be obtained as in theorem 2.

This theorem completes the picture - now for equal singular values: showing that even in the case of coincident singular values, the perturbations are likely to be “well behaved”.

Proof: Let \mathbf{D} have $\{\kappa_i\}$, $\{\mathbf{x}_i\}$ and $\{\mathbf{y}_i\}$ as its first r singular values, right singular vectors and left singular vectors respectively. For $i > r$, $\{\kappa_i\}$, $\{\mathbf{x}_i\}$ and $\{\mathbf{y}_i\}$ can be obtained as in theorem 2. Thus, we concentrate on the perturbed $\{\kappa_i\}$, $\{\mathbf{x}_i\}$ and $\{\mathbf{y}_i\}$, for $i \leq r$.

Combining the techniques in the proof of theorem 2 and theorem 3, we assume that the perturbed right/left singular vectors, \mathbf{x}_i and \mathbf{y}_i ($1 \leq i \leq r$) respectively, have the following forms:

$$\mathbf{x}_i = \sum_{j=1}^r p_{j,i} \mathbf{e}_j + \sum_{j=r+1}^m \Delta f_{j,i} \mathbf{e}_j \quad (41)$$

$$\mathbf{y}_i = \sum_{j=1}^r q_{j,i} \mathbf{e}_j + \sum_{j=r+1}^m \Delta g_{j,i} \mathbf{e}_j \quad (42)$$

Note: $p_{j,i}$ and $q_{j,i}$ can possibly take any values within $[-1,1]$, while $\Delta f_{j,i}$ and $\Delta g_{j,i}$ approach zeroes if \mathbf{N} is small enough. Because the singular values of the matrix, \mathbf{A} , are the square roots of the eigenvalues of $\mathbf{A}\mathbf{A}^T$, and due to the continuity of the eigenvalues of $\mathbf{A}\mathbf{A}^T$, the singular values of \mathbf{A} also obey Ostrowski’s continuity rule (Wilkinson 1965). Supposing the associated singular value is $\kappa_i = \lambda + \Delta\lambda_i$, the equality of the first-order terms of $\mathbf{D}\mathbf{x}_i = \kappa_i\mathbf{y}_i$ and $\mathbf{D}^T\mathbf{y}_i = \kappa_i\mathbf{x}_i$ produces:

$$\lambda \sum_{j=1}^r p_{j,i} \mathbf{e}_j + \sum_{j=r+1}^m \Delta f_{j,i} \lambda_j \mathbf{e}_j + \sum_{j=1}^r p_{j,i} \mathbf{C} \mathbf{e}_j = (\lambda + \Delta\lambda_i) \sum_{j=1}^r q_{j,i} \mathbf{e}_j + \lambda \sum_{j=r+1}^m \Delta g_{j,i} \mathbf{e}_j \quad (43)$$

$$\lambda \sum_{j=1}^r q_{j,i} \mathbf{e}_j + \sum_{j=r+1}^m \Delta g_{j,i} \lambda_j \mathbf{e}_j + \sum_{j=1}^r q_{j,i} \mathbf{C}^T \mathbf{e}_j = (\lambda + \Delta \lambda_i) \sum_{j=1}^r p_{j,i} \mathbf{e}_j + \lambda \sum_{j=r+1}^m \Delta f_{j,i} \mathbf{e}_j \quad (44)$$

From (43) and (44), we have, by equating \mathbf{e}_s (for $s = 1, \dots, r$):

$$\lambda p_{s,i} + \sum_{j=1}^r p_{j,i} c_{s,j} = (\lambda + \Delta \lambda_i) q_{s,i} \quad (45)$$

$$\lambda q_{s,i} + \sum_{i=1}^r q_{j,i} c_{j,s} = (\lambda + \Delta \lambda_i) p_{s,i} \quad (46)$$

In matrix form, (45) and (46) are:

$$(\mathbf{C}_{1:r,1:r} + \lambda \mathbf{I}) \mathbf{p}_i = (\lambda + \Delta \lambda_i) \mathbf{q}_i \quad (47)$$

$$(\mathbf{C}_{1:r,1:r}^T + \lambda \mathbf{I}) \mathbf{q}_i = (\lambda + \Delta \lambda_i) \mathbf{p}_i \quad (48)$$

where $\mathbf{C}_{1:r,1:r}$ is the left-up r -by- r submatrix of \mathbf{C} , $\mathbf{p}_i = [p_{1,i}, p_{2,i}, \dots, p_{r,i}]^T$ and $\mathbf{q}_i = [q_{1,i}, q_{2,i}, \dots, q_{r,i}]^T$. From (47) and (48), $\lambda + \Delta \lambda_i$, \mathbf{p}_i and \mathbf{q}_i are respectively the singular value, the right and the left singular vectors of $\mathbf{C}_{1:r,1:r} + \lambda \mathbf{I}$. Consequently, \mathbf{Q} and \mathbf{P} are $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ in the theorem, respectively. $\mathbf{C}_{1:r,1:r} + \lambda \mathbf{I}$ just has r singular values, right and left singular vectors, which correspond to $\{\lambda_i\}$, $\{\mathbf{x}_i\}$ and $\{\mathbf{y}_i\}$, for $i \leq r$, of \mathbf{D} .

Equating the \mathbf{e}_t in (43) and (44), for $t > r$, we have

$$\sum_{j=1}^r p_{j,i} c_{t,j} = \lambda \Delta g_{t,i} - \lambda_t \Delta f_{t,i} \quad (49)$$

$$\sum_{j=1}^r q_{j,i} c_{j,t} = \lambda \Delta f_{t,i} - \lambda_t \Delta g_{t,i} \quad (50)$$

Note for $t > r$ and $j \leq r$, $c_{t,j} = d'_{t,j}$ and $c_{j,t} = d_{j,t}$. And, from the fact $\mathbf{D}' = \mathbf{U}''^T \mathbf{D} \mathbf{V}''$ (in the definition of the theorem) and the fact \mathbf{Q} and \mathbf{P} are $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ respectively (as proved above):

$$\lambda \Delta g_{t,i} - \lambda_t \Delta f_{t,i} = \sum_{j=1}^r p_{j,i} c_{t,j} = \sum_{j=1}^r p_{j,i} d'_{t,j} = d'_{t,i} \quad (51)$$

$$\lambda \Delta f_{t,i} - \lambda_t \Delta g_{t,i} = \sum_{j=1}^r q_{j,i} c_{j,t} = \sum_{j=1}^r q_{j,i} d_{j,t} = d'_{i,t} \quad (52)$$

Combining (47-48) and (51-52),

$$\mathbf{x}_i = \sum_{j=1}^r p_{j,i} \mathbf{e}_j + \sum_{j=r+1}^m \frac{\lambda_j d'_{j,i} + \lambda d'_{i,j}}{\lambda^2 - \lambda_j^2} \mathbf{e}_j \quad (53)$$

$$\mathbf{y}_i = \sum_{j=1}^r q_{j,i} \mathbf{e}_j + \sum_{j=r+1}^m \frac{\lambda d'_{j,i} + \lambda_j d'_{i,j}}{\lambda^2 - \lambda_j^2} \mathbf{e}_j \quad (54)$$

are respectively the right and left singular vectors of \mathbf{D} . According to the coordinate transformation of $\mathbf{B} = \mathbf{U} \mathbf{D} \mathbf{V}^T$ and the fact that \mathbf{Q} and \mathbf{P} are $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ in the theorem, the first r left/right singular vectors of \mathbf{B} are as defined in (39,40). \square

A.2 Denoising capacity of SVD

A.2.1 Case of distinct singular values

First, we consider the simplest case: a square matrix with a few distinct non-zero singular values. Recall the definitions in theorem 2 and lemma 1: \mathbf{A} is the signal matrix, \mathbf{N} is the i.i.d. Gaussian noise matrix (with zero mean and σ^2 variance), $\mathbf{B}=\mathbf{A}+\mathbf{N}$, $\mathbf{A}=\mathbf{U}\Sigma\mathbf{V}^T$ and $\mathbf{C}=\mathbf{U}^T\mathbf{N}\mathbf{V}$. Note \mathbf{C} is still an i.i.d. Gaussian noise matrix (with zero mean and σ^2 variance). Further, define $\mathbf{D}=\mathbf{C}+\Sigma$. Then,

$$\mathbf{B} = \mathbf{U}\mathbf{D}\mathbf{V}^T \quad (55)$$

$\{\mathbf{u}'_i\}$ and $\{\mathbf{v}'_i\}$, are respectively the left/right singular vectors of \mathbf{B} . $\{\mathbf{x}_i\}$ and $\{\mathbf{y}_i\}$, are respectively the right/left singular vectors of \mathbf{D} . Because \mathbf{U} and \mathbf{V} in (55) are orthogonal matrices,

$$\mathbf{u}'_i = \mathbf{U}\mathbf{y}_i \text{ and } \mathbf{v}'_i = \mathbf{V}\mathbf{x}_i \quad (56)$$

And, also from (55), the singular values of \mathbf{B} , $\{\lambda'_i\}$, are same as the corresponding singular values of \mathbf{D} .

Suppose that the noise-free matrix \mathbf{A} should have a rank of r , i.e. $\mathbf{A} = \sum_{i=1}^r \lambda_i \mathbf{u}_i \mathbf{v}_i^T$.

Combining (11), (55) and (56), the closest rank- r approximation of \mathbf{B} is

$$\mathbf{B}^r = \sum_{i=1}^r \lambda'_i \mathbf{u}'_i \mathbf{v}'_i{}^T = \mathbf{U} \left(\sum_{i=1}^r \lambda'_i \mathbf{y}_i \mathbf{x}_i^T \right) \mathbf{V}^T = \mathbf{U}\mathbf{D}^r\mathbf{V}^T \quad (57)$$

where $\mathbf{D}^r = \sum_{i=1}^r \lambda'_i \mathbf{y}_i \mathbf{x}_i^T$.

Then

$$\|\mathbf{B}^r - \mathbf{A}\|_F^2 = \left\| \sum_{i,j} (d_{i,j}^r - \Sigma_{i,j}) \mathbf{u}_i \mathbf{v}_j^T \right\|_F^2 \quad (58)$$

Due to the mutual orthonormality among any $\mathbf{u}_i \mathbf{v}_j^T$, we have the following formula:

$$\|\mathbf{B}^r - \mathbf{A}\|_F^2 = \|\mathbf{D}^r - \Sigma\|_F^2 \quad (59)$$

From lemma *I* we can write an expression for the terms $\lambda'_i \mathbf{y}_i \mathbf{x}_i^T$, for the special case where *A* is rank-*r* and thus the λ_i are zero for $i > r$. For example $\lambda'_1 \mathbf{y}_1 \mathbf{x}_1^T$, is:

$$\lambda'_1 \mathbf{y}_1 \mathbf{x}_1^T \approx \begin{bmatrix} \lambda'_1 & \lambda'_1 \boldsymbol{\zeta}^T \\ \lambda'_1 \boldsymbol{\xi} & \mathbf{0} \end{bmatrix} \approx \begin{bmatrix} \lambda_1 + c_{1,1} & \lambda_1 \boldsymbol{\zeta}^T \\ \lambda_1 \boldsymbol{\xi} & \mathbf{0} \end{bmatrix} \quad (60)$$

where $[\boldsymbol{\zeta}]_{m-1,1} = \left[\frac{\lambda_2 c_{2,1} + \lambda_1 c_{1,2}}{\lambda_1^2 - \lambda_2^2} \quad \dots \quad \frac{\lambda_r c_{r,1} + \lambda_1 c_{1,r}}{\lambda_1^2 - \lambda_r^2} \quad \frac{c_{1,r+1}}{\lambda_1} \quad \dots \quad \frac{c_{1,m}}{\lambda_1} \right]^T$, and

$[\boldsymbol{\xi}]_{m-1,1} = \left[\frac{\lambda_1 c_{2,1} + \lambda_2 c_{1,2}}{\lambda_1^2 - \lambda_2^2} \quad \dots \quad \frac{\lambda_1 c_{r,1} + \lambda_r c_{1,r}}{\lambda_1^2 - \lambda_r^2} \quad \frac{c_{r+1,1}}{\lambda_1} \quad \dots \quad \frac{c_{m,1}}{\lambda_1} \right]^T$. Note, 2nd-order and

higher-order terms in (60) have been dropped, and that the above matrix has zeros except for entries in column 1 and row 1. The reader can verify for themselves that, on adding the contribution for $\lambda'_2 \mathbf{y}_2 \mathbf{x}_2$ fills in the second column and second row, with entries similar to the first row but also that where these overlap with column 1 and row 1, the contribution of $\lambda'_2 \mathbf{y}_2 \mathbf{x}_2$ leads to a simplification of the expression in those two entries.

By combining all such terms, filling in rows and columns, it is straightforward to obtain

$$[\mathbf{D}^r]_{m,m} - [\Sigma]_{m,m} \approx [\mathbf{Y}]_{m,m} \quad (61)$$

$$\text{where } \mathbf{Y} = \begin{bmatrix} \begin{bmatrix} c_{1,1} & \cdots & c_{1,r} \\ \vdots & \ddots & \vdots \\ c_{r,1} & \cdots & c_{r,r} \end{bmatrix}_{r,r} & \begin{bmatrix} c_{1,r+1} & \cdots & c_{1,m} \\ \vdots & \ddots & \vdots \\ c_{r,r+1} & \cdots & c_{r,m} \end{bmatrix}_{r,m-r} \\ \begin{bmatrix} c_{r+1,1} & \cdots & c_{r+1,r} \\ \vdots & \ddots & \vdots \\ c_{m,1} & \cdots & c_{m,r} \end{bmatrix}_{m-r,r} & \begin{bmatrix} 0 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 0 \end{bmatrix}_{m-r,m-r} \end{bmatrix}_{m,m}$$

There are two notable aspects of what has just been derived. Firstly, there is a large zero block which reduces the degrees of freedom that can “hold” significant noise. Second, the non-zero part of the matrix is, to first order, composed of elements of the same size as the original noise elements. Thus the rank projected matrix has been “de-noised” by the amount given below:

$$E \|\mathbf{B}^r - \mathbf{A}\|_F^2 = E \|\mathbf{Y}\|_F^2 = \sum E y_{i,j}^2 = (2rm - r^2) \sigma^2 \quad (62)$$

$$E |b_{i,j}^r - a_{i,j}| = \sigma \frac{\sqrt{2rm - r^2}}{m} \quad (63)$$

Obviously, (62) is a special case of (1) for square matrices, where $n=m$.

We also note that, in deducing the formula (61), we do not impose the i.i.d. Gaussianity upon the noise in the observed matrix. This means that the zero-block of \mathbf{Y} in (61) holds for *any* type of noise, if and only if the noise level is small enough compared with the signal level. In fact, the i.i.d. Gaussianity of the noise matrix is required only at the last equality of (62). From another point of view, the noise matrix can be decomposed into m^2 components: $\{\mathbf{u}_i \mathbf{v}_j^T \mid 1 \leq i \leq m, 1 \leq j \leq m\}$. The components of the noise matrix, for $\{(i, j) \mid r < i \leq m, r < j \leq m\}$, “live” outside of the subspace we project onto and therefore can be reduced or *denoised* in the rank- r approximation; while the other

components of the noise, for $\{(i, j) \mid i \leq r \text{ or } j \leq r\}$, are kept. The denoising capacity for non i.i.d. Gaussian noise is dependent on the ratio of the number these two types of components.

A.2.2 Case of multiple singular value

As in the theorem 4, suppose \mathbf{A} has r singular values the same as each other.

Following the notation in theorem 4, we similarly have:

$$\mathbf{B}^r = \sum_{i=1}^r \lambda'_i \mathbf{u}'_i \mathbf{v}'_i{}^T = \mathbf{U} \left(\sum_{i=1}^r \lambda'_i \mathbf{y}_i \mathbf{x}_i{}^T \right) \mathbf{V}^T = \mathbf{U} \mathbf{D}^r \mathbf{V}^T \quad (64)$$

By the same techniques as in the previous section, the first-order perturbation of \mathbf{D}^r has the following form:

$$\mathbf{D}^r \approx \begin{bmatrix} d'_{1,1} & \cdots & d'_{1,r} & d'_{1,r+1} & \cdots & d'_{1,m} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ d'_{r,1} & \cdots & d'_{r,r} & d'_{r,r+1} & \cdots & d'_{r,m} \\ d'_{r+1,1} & \cdots & d'_{r+1,r} & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ d'_{m,1} & \cdots & d'_{m,r} & 0 & \cdots & 0 \end{bmatrix} \quad (65)$$

Then,

$$\begin{aligned}
\mathbf{D}^r &= \mathbf{U}^n \mathbf{D}'^r \mathbf{V}^{nT} \approx \begin{bmatrix} \hat{\mathbf{U}} & \\ & \mathbf{I}_{m-k} \end{bmatrix} \begin{bmatrix} d'_{1,1} & \cdots & d'_{1,r} & d'_{1,r+1} & \cdots & d'_{1,m} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ d'_{r,1} & \cdots & d'_{r,r} & d'_{r,r+1} & \cdots & d'_{r,m} \\ d'_{r+1,1} & \cdots & d'_{r+1,r} & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ d'_{m,1} & \cdots & d'_{m,r} & 0 & \cdots & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{V}}^T \\ & \mathbf{I}_{m-k} \end{bmatrix} \\
&= \begin{bmatrix} d_{1,1} & \cdots & d_{1,r} & d_{1,r+1} & \cdots & d_{1,m} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ d_{r,1} & \cdots & d_{r,r} & d_{r,r+1} & \cdots & d_{r,m} \\ d_{r+1,1} & \cdots & d_{r+1,r} & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ d_{m,1} & \cdots & d_{m,r} & 0 & \cdots & 0 \end{bmatrix}
\end{aligned}$$

Since $\mathbf{D} = \mathbf{\Sigma} + \mathbf{C}$: $d_{i,i} = c_{i,i} + \lambda_i$ for $1 \leq i \leq r$ and $d_{i,j} = c_{i,j}$ if $(i, j) \notin \{(1,1), (2,2), \dots, (r,r)\}$.

Thus, we also have $\mathbf{D}^r - \mathbf{\Sigma} = \mathbf{Y}$, which is the same result as in the case with distinct singular values.

A.2.3 Extension to the case of a rectangular matrix

The analysis of the case where the matrix is rectangular follows the same pattern as described and is omitted for brevity.

A.3 Learning capacity of LSA-based recognition system

A.3.1 Perturbation of the basis images

First, we analyze the learning stage, by using the matrix perturbation theory in appendices A.1.3 and A.1.4. By SVD, the low-dimension subspaces, $\mathbf{U}'' = [\mathbf{u}'_1, \mathbf{u}'_2, \dots, \mathbf{u}'_r]$ and $\mathbf{V}'' = [\mathbf{v}'_1, \mathbf{v}'_2, \dots, \mathbf{v}'_r]$, as defined in theorem 2, are obtained. In some cases, such as in face recognition, the consequent step is contingent on an accurate basis. Here, we only consider the subspace $\mathbf{U}'' : \mathbf{U}'' = \mathbf{UH}$, where

$$\mathbf{H} = \begin{bmatrix} 1 & \frac{\lambda_2 c_{1,2} + \lambda_1 c_{2,1}}{\lambda_2^2 - \lambda_1^2} & \dots & \frac{\lambda_r c_{1,r} + \lambda_1 c_{r,1}}{\lambda_r^2 - \lambda_1^2} \\ \frac{\lambda_1 c_{2,1} + \lambda_2 c_{1,2}}{\lambda_1^2 - \lambda_2^2} & 1 & \vdots & \vdots \\ \vdots & \dots & \ddots & \frac{\lambda_r c_{r-1,r} + \lambda_{r-1} c_{r,r-1}}{\lambda_r^2 - \lambda_{r-1}^2} \\ \frac{\lambda_1 c_{r,1} + \lambda_r c_{1,r}}{\lambda_1^2 - \lambda_r^2} & \dots & \frac{\lambda_{r-1} c_{r,r-1} + \lambda_r c_{r-1,r}}{\lambda_{r-1}^2 - \lambda_r^2} & 1 \\ \frac{c_{r+1,1}}{\lambda_1} & \frac{c_{r+1,2}}{\lambda_2} & \dots & \frac{c_{r+1,r}}{\lambda_r} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{c_{m,1}}{\lambda_1} & \frac{c_{m,2}}{\lambda_2} & \dots & \frac{c_{m,r}}{\lambda_r} \end{bmatrix}_{m,r} \quad (66)$$

Note: From (66), we can see that (roughly) the perturbed $\{\mathbf{u}'_i\}$ have been corrupted to different extent, which depends on their associated singular values. For example, $\mathbf{u}'_1 \approx \mathbf{u}_1 + \sum_{j=2}^r \frac{\lambda_1 c_{j,1} + \lambda_j c_{1,j}}{\lambda_1^2 - \lambda_j^2} \mathbf{u}_j + \sum_{j=r+1}^m \frac{c_{j,1}}{\lambda_1} \mathbf{u}_j$, where the third term $\sum_{j=r+1}^m \frac{c_{j,1}}{\lambda_1} \mathbf{u}_j$ is inversely proportional to λ_1 . If we only consider the perturbations, coming from $\{\mathbf{u}_i \mid (i > r)\}$, \mathbf{u}'_1 can be considered *cleanest*, while \mathbf{u}'_r the *dirtiest*. In the following section A.3.2, we will return to this point when the projection error is analyzed.

Furthermore, to decompose \mathbf{H} into: $\mathbf{H} = \mathbf{E} + \Delta\mathbf{F} + \Delta\mathbf{G}$, where $\mathbf{E} = \begin{bmatrix} \mathbf{I}_r \\ \mathbf{0}_{(m-r) \times r} \end{bmatrix}_{m,r}$

$$\Delta\mathbf{F} = \begin{bmatrix} 0 & \frac{\lambda_2 c_{1,2} + \lambda_1 c_{2,1}}{\lambda_2^2 - \lambda_1^2} & \dots & \frac{\lambda_r c_{1,r} + \lambda_1 c_{r,1}}{\lambda_r^2 - \lambda_1^2} \\ \frac{\lambda_1 c_{2,1} + \lambda_2 c_{1,2}}{\lambda_1^2 - \lambda_2^2} & 0 & \vdots & \vdots \\ \vdots & \dots & \ddots & \frac{\lambda_r c_{r-1,r} + \lambda_{r-1} c_{r,r-1}}{\lambda_r^2 - \lambda_{r-1}^2} \\ \frac{\lambda_1 c_{r,1} + \lambda_r c_{1,r}}{\lambda_1^2 - \lambda_r^2} & \dots & \frac{\lambda_{r-1} c_{r,r-1} + \lambda_r c_{r-1,r}}{\lambda_{r-1}^2 - \lambda_r^2} & 0 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix}_{m,r} \quad (67)$$

$$\Delta\mathbf{G} = \begin{bmatrix} 0 & 0 & \dots & 0 \\ 0 & 0 & \vdots & \vdots \\ \vdots & \dots & \ddots & 0 \\ 0 & 0 & \dots & 0 \\ \frac{c_{r+1,1}}{\lambda_1} & \frac{c_{r+1,2}}{\lambda_2} & \dots & \frac{c_{r+1,r}}{\lambda_r} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{c_{m,1}}{\lambda_1} & \frac{c_{m,2}}{\lambda_2} & \dots & \frac{c_{m,r}}{\lambda_r} \end{bmatrix}_{m,r} \quad (68)$$

A.3.2 Projection of a new test image on the basis images

The underlying noise-free subspace $\mathbf{U}^r = \mathbf{U} \begin{bmatrix} \mathbf{I}_r \\ \mathbf{0}_{(m-r) \times r} \end{bmatrix} = \mathbf{U}\mathbf{E}$. Suppose a noise corrupted test image \mathbf{p} , to be identified, is observed. Its underlying truth is \mathbf{q} and the noise in \mathbf{p} is \mathbf{n} . Vectors \mathbf{q} and \mathbf{n} can be represented in \mathbf{U} as $\mathbf{q} = \mathbf{U}\mathbf{f}$ and $\mathbf{n} = \mathbf{U}\Delta\mathbf{g}$, and consequently $\mathbf{p} = \mathbf{U}(\mathbf{f} + \Delta\mathbf{g})$. Because $\mathbf{q} \in \mathbf{U}^r$, only the first r components of \mathbf{f} are possibly non-zeroes, i.e. $\mathbf{f} = [f_1, f_2, \dots, f_r, 0, \dots, 0]^T$. In practice, the noise-corrupted test image has to be projected on the noise-corrupted basis in the recognition system because

the noise free basis is always unknown. More formally, substitute

$\mathbf{U}'' = \mathbf{U}\mathbf{H} = \mathbf{U}(\mathbf{E} + \Delta\mathbf{F} + \Delta\mathbf{G})$ and $\mathbf{p} = \mathbf{U}(\mathbf{f} + \Delta\mathbf{g})$ into the projection error of \mathbf{p} on \mathbf{U}'' ,

$\mathbf{p} - \mathbf{U}''\mathbf{U}''^T\mathbf{p}$:

$$\begin{aligned} \mathbf{p} - \mathbf{U}''\mathbf{U}''^T\mathbf{p} &= \mathbf{U}(\mathbf{f} + \Delta\mathbf{g}) - \mathbf{U}(\mathbf{E} + \Delta\mathbf{F} + \Delta\mathbf{G})(\mathbf{E} + \Delta\mathbf{F} + \Delta\mathbf{G})^T\mathbf{U}^T\mathbf{U}(\mathbf{f} + \Delta\mathbf{g}) \\ &\approx \mathbf{U}[\Delta\mathbf{g} - \mathbf{E}\mathbf{E}^T\Delta\mathbf{g} - (\mathbf{E}\Delta\mathbf{F}^T + \mathbf{E}\Delta\mathbf{G}^T + \Delta\mathbf{F}\mathbf{E}^T + \Delta\mathbf{G}\mathbf{E}^T)\mathbf{f}] \\ &= \mathbf{U}[\Delta\mathbf{g} - \mathbf{E}\mathbf{E}^T\Delta\mathbf{g} - (\mathbf{E}\Delta\mathbf{G}^T + \Delta\mathbf{G}\mathbf{E}^T)\mathbf{f}] \\ &= \mathbf{U}[\Delta\mathbf{g}' - \Delta\mathbf{G}\mathbf{f}'] \end{aligned} \quad (69)$$

where $\mathbf{f}' = [f_1, f_2, \dots, f_r]^T$; and $\Delta\mathbf{g}'$ has same components as $\Delta\mathbf{g}$, except its first r

zeroes, i.e. $\Delta\mathbf{g}' = [0, 0, \dots, 0, \Delta g_{r+1}, \dots, \Delta g_m]^T$. Note, $\mathbf{E}\Delta\mathbf{F}^T$ is antisymmetric, i.e.

$\mathbf{E}\Delta\mathbf{F}^T + \Delta\mathbf{F}\mathbf{E}^T = \mathbf{0}$; and, the 2-order and higher-order terms in (69) have been dropped:

$\Delta\mathbf{F}$ and $\Delta\mathbf{G}$ can possibly approach $\mathbf{0}$. From (69),

$$\mathbf{p} - \mathbf{U}''\mathbf{U}''^T\mathbf{p} = \mathbf{U}[\Delta\mathbf{g}', \mathbf{c}'_1, \mathbf{c}'_2, \dots, \mathbf{c}'_r][1, -\mathbf{h}^T]^T \quad (70)$$

$$\left\| \mathbf{p} - \mathbf{U}''\mathbf{U}''^T\mathbf{p} \right\|_F = \left\| [\Delta\mathbf{g}', \mathbf{c}'_1, \mathbf{c}'_2, \dots, \mathbf{c}'_r][1, -\mathbf{h}^T]^T \right\|_F \quad (71)$$

where $\mathbf{c}'_i = [0, \dots, 0, c_{r+1,i}, c_{r+2,i}, \dots, c_{m,i}]^T$ and $\mathbf{h} = [f_1/\lambda_1, \dots, f_r/\lambda_r]^T$.

We suppose the basis is obtained from n learning samples, i.e., the learning matrix is $\mathbf{A} \in R^{m,n}$, and each entry of \mathbf{A} has energy of σ_s^2 and is corrupted with i.i.d. Gaussian

noise with energy of σ_l^2 . It is also assumed that the test image has energy of σ_s^2 and is

corrupted with noise of σ_l^2 . $\|\mathbf{A}\|_F^2/mn = \sum_{i=1}^r \lambda_i^2/mn \cong \sigma_s^2$, $\sum_{i=1}^m \sum_{j=1}^n c_{i,j}^2/mn \cong \sigma_l^2$,

$\|\mathbf{q}\|_F^2/m = \|\mathbf{f}\|_F^2/m = \sum_{i=1}^r f_i^2/m \cong \sigma_s^2$, and $\|\mathbf{n}\|_F^2/m = \|\Delta\mathbf{g}\|_F^2/m = \sum_{i=1}^m \Delta g_i^2/m \cong \sigma_l^2$.

$\|\Delta \mathbf{g}'\|_F \cong \sqrt{m-r}\sigma_t$ and $\|\mathbf{c}'_i\|_F \cong \sqrt{m-r}\sigma_i$. Due to the independence among $\{\Delta \mathbf{g}', \{\mathbf{c}'_i \mid i = 1, \dots, r\}\}$, (71) becomes

$$\left\| \mathbf{p} - \mathbf{U}^r \mathbf{U}^{rT} \mathbf{p} \right\|_F^2 \cong (m-r)\sigma_t^2 + (m-r)\sigma_i^2 \sum_{i=1}^r \frac{f_i^2}{\lambda_i^2} \quad (72)$$

which is (6) in result 2 of section 2. And,

$$\|\mathbf{p}\|_F^2 \cong m(\sigma_s^2 + \sigma_t^2) \quad (73)$$

The second term in (72) is caused by the noise in the learning samples because it is independent of the noise level in the test image. Furthermore, this part is contingent on the ratios between $\{f_i\}$ and $\{\lambda_i\}$. Note, however, that this last term is *not independent of the chosen test image because $\{f_i\}$ is the representation of the underlying noise free test image in the true (noise free) subspace that we tried to learn*. More specifically, the error of $(m-r)\sigma_i^2 \frac{f_i^2}{\lambda_i^2}$ in the second term of (72) is *caused* by the perturbation in \mathbf{u}'_i and that subspace perturbation is amplified by the ratios between the true image components and the singular values of the true matrix SVD. Thus, it can be concluded that the basis \mathbf{u}'_1 that corresponds to the largest singular value leads to the least noise amplification and is in that sense the “*cleanest*”, and that the basis \mathbf{u}'_r that corresponds to the least singular value is the “*dirtiest*”.

For a random test image, the best and worst performance is:

$$(m-r)\sigma_t^2 + (m-r)\sigma_i^2 \frac{\sum_{i=1}^r f_i^2}{\lambda_1^2} \leq \left\| \mathbf{p} - \mathbf{U}^r \mathbf{U}^{rT} \mathbf{p} \right\|_F^2 \leq (m-r)\sigma_t^2 + (m-r)\sigma_i^2 \frac{\sum_{i=1}^r f_i^2}{\lambda_r^2} \quad (74)$$

$$(m-r)\sigma_i^2 + (m-r)\sigma_i^2 \frac{m\sigma_s^2}{\lambda_1^2} \leq \left\| \mathbf{p} - \mathbf{U}'^r \mathbf{U}'^{rT} \mathbf{p} \right\|_F^2 \leq (m-r)\sigma_i^2 + (m-r)\sigma_i^2 \frac{m\sigma_s^2}{\lambda_r^2} \quad (75)$$

where $\lambda_r^2 \leq \frac{mn\sigma_s^2}{r} \leq \lambda_1^2$. Define, furthermore, $\lambda_i^2 = k_i mn\sigma_s^2$:

$$(m-r)\sigma_i^2 + \frac{m-r}{nk_1} \sigma_i^2 \leq \left\| \mathbf{p} - \mathbf{U}'^r \mathbf{U}'^{rT} \mathbf{p} \right\|_F^2 \leq (m-r)\sigma_i^2 + \frac{m-r}{nk_r} \sigma_i^2 \quad (76)$$

It can be easily proved that $(m-r)\sigma_i^2 + \frac{(m-r)r\sigma_i^2}{n}$ is the expectation for any

test sets when all of the r (non-zero) singular values of the learning matrix \mathbf{A} are equal.

A.3.3 Performance analysis over the learning samples

Here, we analyze the average performance of the system when we test the basis on the whole learning examples, i.e. all the images that are used to obtain the basis images. Using the notation in section A.3.2: \mathbf{q} and \mathbf{p} are the noise-free and observed test image, and $\mathbf{q} = \mathbf{U}\mathbf{f}$ (\mathbf{f} is a vector of coefficients of the uncontaminated image expressed in the uncontaminated subspace). In fact, \mathbf{q} is one of the columns of \mathbf{A} , $\{\mathbf{a}_i \mid i = 1, \dots, n\}$. Note: \mathbf{p} is an observation of \mathbf{q} that contains noise *and that we allow this noise to have a different scale σ_i to the noise in the observations σ_i used to learn the subspace*. We sample noise corrupted versions of the columns of \mathbf{A} with uniform probability of the column number and each sample has i.i.d. noise σ_i . We then seek the expectation of the error in the projection onto the learned subspace:

$$E_{\mathbf{q} \in \{\mathbf{a}_i \mid i=1, \dots, n\}} \left\| \mathbf{p} - \mathbf{U}'^r \mathbf{U}'^{rT} \mathbf{p} \right\|_F^2 = (m-r)\sigma_i^2 + (m-r)\sigma_i^2 \sum_{i=1}^r \frac{E_{\mathbf{q} \in \{\mathbf{a}_i \mid i=1, \dots, n\}} f_i^2}{\lambda_i^2} \quad (77)$$

From the SVD of \mathbf{A} ,

$$\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_n] = [\mathbf{u}_1, \dots, \mathbf{u}_r] \text{diag}(\lambda_1, \dots, \lambda_r) [\mathbf{v}_1, \dots, \mathbf{v}_r]^T = [\mathbf{u}_1, \dots, \mathbf{u}_r] [\lambda_1 \mathbf{v}_1, \dots, \lambda_r \mathbf{v}_r]^T$$

Perhaps surprisingly,

$$\sum_{\mathbf{q} \in \{\mathbf{a}_i | i=1, \dots, n\}} f_i^2 = \|\lambda_i \mathbf{v}_i\|_F^2 = \lambda_i^2 \quad \text{and} \quad E_{\mathbf{q} \in \{\mathbf{a}_i | i=1, \dots, n\}} f_i^2 = \frac{\lambda_i^2}{n} \quad (78)$$

Substituting (78) into (77), obtain

$$E_{\mathbf{q} \in \{\mathbf{a}_i | i=1, \dots, n\}} \left\| \mathbf{p} - \mathbf{U}^r \mathbf{U}^{rT} \mathbf{p} \right\|_F^2 \cong (m-r)\sigma_i^2 + (m-r)r\sigma_i^2/n \quad (79)$$

which is (7) in result 2 of section 2.

From this formula, (79), we can see clearly the effects of all the parameters in the recognition system. Given that the noise in the learning samples and in the test image, compared with the signal, is small, the performance can be regarded to be independent of the signal level. As m , compared with r , approaches a very large number, the SSD is almost linearly dependent on m . As the number of the learning samples, n , increases, the recognition system improves: the error from the basis images decreases, and as n approaches infinite, the error from the basis images approaches zero. However, the error from the test image can't be reduced except by having a cleaner image.

Another measure, used in the recognition system, is the angle between the test image and the basis images:

$$\frac{\left\| \mathbf{p} - \mathbf{U}^r \mathbf{U}^{rT} \mathbf{p} \right\|_F^2}{\left\| \mathbf{p} \right\|_F^2} = \frac{(m-r)\sigma_i^2 + (m-r)r\sigma_i^2/n}{m(\sigma_s^2 + \sigma_i^2)} \xrightarrow{m \rightarrow \infty} \frac{\sigma_i^2 + r\sigma_i^2/n}{\sigma_s^2 + \sigma_i^2} = \frac{\sigma_i^2}{\sigma_s^2 + \sigma_i^2} + \frac{r\sigma_i^2}{n(\sigma_s^2 + \sigma_i^2)} \quad (80)$$

Supposing $m \gg r$, the angle is independent of the size of the object, and depends on the energy levels of the signal and the noises (in the learning samples and in the test image). As the size of the learning samples, n , increases, the system improves: the error

from the basis images approaches zero and the error from the test image gradually dominates in the total error.

A.3.4 The optimal learning set

Suppose that the expectation of the test images, i.e. $\{f_i^2\}$, in (72), is known. How should we design the recognition system? Specifically, how to select the learning samples, so that the system, concerning the expectation, has the best performance? Obviously, only the second term in (72) is dependent on the learning samples. The problem is:

$$\min \sum \frac{f_i^2}{\lambda_i^2}, \text{ subject to } \sum \lambda_i^2 = C \quad (81)$$

$\sum \lambda_i^2 = C$ means that, when the dimension, m , and the size, n , of the learning samples is large enough, the signal energy, $\sum \lambda_i^2$, should be approximately $mn\sigma_s^2$. By using a Lagrange multiplier, the minimum can be obtained *iff*

$$\frac{f_i}{\lambda_i^2} \equiv \text{Cons} \quad (82)$$

which is (8) in result 2 of section 2.

From (82), we can draw the conclusion that the basis images, obtained from the n samples of \mathbf{A} (i.e., $\{\mathbf{a}_i | 1 \leq i \leq n\}$), are not optimal when the test image set is also $\{\mathbf{a}_i | 1 \leq i \leq n\}$. (Though, this conclusion is somewhat surprising.) The reason is that, the basis, associated with the largest singular value, is *overlearned* in the *learning* process: From (82), the optimal *learning ability*, λ_i^2 , should be proportional to f_i , while λ_i^2 is actually proportional to f_i^2 , as in (78).

However, the expectation of the test images $\{f_i^2\}$ is usually not known beforehand. In such cases without prior knowledge, we can reasonably suppose that $E(f_i^2)$ in (72) are statistically equivalent. From (82), the optimal learning set for such a *random* test image should have r equal singular values, and, for a *random* test image, (79) is the optimal reconstruction error.

References

- Basri, R. and D. W. Jacobs (1999). Lambertian Reflectance and Linear Subspaces. Proc. Int'l Conf. Computer Vision, 383-390.
- Basri, R. and D. W. Jacobs (2003). "Lambertian Reflectance and Linear Subspaces." IEEE Trans Pattern Analysis and Machine Intelligence **25**(2): 218-233.
- Belhumeur, P. N., J. Hespanha, et al. (1997). "Eigenfaces vs. Fisherfaces: Recognition Using Class Specific Linear Projection." IEEE Trans Pattern Analysis and Machine Intelligence **19**(7): 711-720.
- Belhumeur, P. N. and D. Kriegman (1998). "What is the Set of Images of an Object under all Possible Illumination Conditions?" Int'l J. Computer vision **28**(3): 245-260.
- Chen, P. and D. Suter (2004a). "Recovering the missing components in a large noisy low-rank matrix: Application to SFM." IEEE Trans Pattern Analysis and Machine Intelligence **26**(8): 1051-1063.
- Chen, P. and D. Suter (2004b). "Subspace-based face recognition: outlier detection and a new distance criterion." Accepted by International Journal of Pattern Recognition and Artificial Intelligence (Special issue of ACCV 2004).
- Chen, P. and D. Suter (2004c). Subspace-based face recognition: outlier detection and a new distance criterion. Asian Conf. Computer Vision 2004, 830-835.
- Eipstein, R., P. Hallinan, et al. (1995). 5±2 eigenimages suffices: An empirical investigation of low-dimensional lighting models. Proc. IEEE Workshop Physics-Based vision, 108-116.
- Faugeras, O. and Q. T. Luong (2001). The geometry of multiple images : the laws that govern the formation of multiple images of a scene and some of their applications. Cambridge, Mass., MIT Press.
- Georghiades, A., P. N. Belhumeur, et al. (2001). "From Few to Many: Generative Models of Object Recognition." IEEE Trans Pattern Analysis and Machine Intelligence **23**(6): 643-660.
- Georghiades, A., D. Kriegman, et al. (1998). Illumination cones for recognition under variable lighting: faces. Proc. Conf. Computer Vision and Pattern Recognition, 52-58.
- Golub, G. H. and C. F. V. Loan (1996). Matrix Computations. Baltimore, Johns Hopkins University Press.
- Hallinan, P. (1994). A low-dimensional representation of human faces for arbitrary lighting conditions. Proc. Conf. Computer Vision and Pattern Recognition, 995-999.
- Hartley, R. and A. Zisserman (2000). Multiple view geometry in computer vision, Cambridge University Press.
- Irani, M. (1999). Multi-frame optical flow estimation using subspace constraints. Proc. Int'l Conf. Computer Vision, 626-633.
- Irani, M. (2002). "Multi-frame correspondence estimation using subspace constraints." Int'l J. Computer vision **48**(3): 173-194.
- Kahl, F. and A. Heyden (1999). "Affine structure and motion from points, lines and conics." Int'l J. Computer vision **33**(3): 163-180.
- Kanatani, K. (2001). Motion segmentation by subspace separation and model selection. Proc. Int'l Conf. Computer Vision, 301-306.
- Leedan, Y. and P. Meer (1999). Estimation with bilinear constraints in computer vision. Proc. Int'l Conf. Computer Vision, 733-738.
- Leedan, Y. and P. Meer (2000). "Heteroscedastic Regression in Computer Vision: Problems with Bilinear Constraint." Int'l J. Computer vision **37**(2): 127-150.

- Ma, Y., K. Huang, et al. (2004). "Rank conditions on the multiple-view matrix." INT'L J. COMPUTER VISION **59**(2): 115-137.
- Mathai, A. M. (1997). Jacobians of Matrix Transformations and Functions of Matrix Argument, World Scientific Publishers.
- Morita, T. and T. Kanade (1997). "A sequential factorization method for recovering shape and motion from image streams." IEEE Trans Pattern Analysis and Machine Intelligence **19**(8): 858-867.
- Moses, Y., Y. Adini, et al. (1994). Face recognition: The problem of compensating for changes in illumination direction. Proc. European Conf. Computer Vision, 286-296.
- Murase, H. and S. K. Nayar (1994). "Illumination planning for object recognition using parametric eigenspaces." IEEE Trans Pattern Analysis and Machine Intelligence **16**(12): 1219-1227.
- Murase, H. and S. K. Nayar (1995). "Visual learning and recognition of 3-D objects from appearance." Int'l J. Computer vision **14**: 5-24.
- Nene, S. A., S. K. Nayar, et al. (1994). Software library for appearance matching (SLAM). ARPA Image understanding workshop.
- Papadopoulos, T. and M. I. A. Lourakis (2000). Estimating the Jacobian of the Singular Value Decomposition: theory and applications. Proc. European Conf. Computer Vision, 554-570.
- Poelman, C. and T. Kanade (1997). "A paraperspective factorization method for shape and motion recovery." IEEE Trans Pattern Analysis and Machine Intelligence **19**(3): 206-219.
- Press, W. H., S. A. Teukolsky, et al. (1992). Numerical recipes in C, Cambridge University Press.
- Ramamoorthi, R. (2002). "Analytic PCA construction for theoretical analysis of lighting variability in images of a Lambertian Object." IEEE Trans Pattern Analysis and Machine Intelligence **24**(10): 1322-1333.
- Ramamoorthi, R. and P. Hanrahan (2001). "On the relationship between radiance and irradiance: Determining the illumination from images of a convex Lambertian object." Journal of the Optical Society of America (JOSA A) **18**(10): 2448-2459.
- Shashua, A. and S. Avidan (1996). The Rank 4 Constraint in Multiple (≥ 2) View Geometry. Proc. European Conf. Computer Vision, 196-206.
- Stewart, G. W. and J. G. Sun (1990). Matrix perturbation theory, Academic press.
- Thomas, J. I. and J. Oliensis (1999). "Dealing with noise in multiframe structure from motion." Computer Vision and Image Understanding **76**(2): 109-124.
- Tomasi, C. and T. Kanade (1992). "Shape and motion from image streams under orthography: A factorization method." Int'l J. Computer vision **9**(2): 137-154.
- Turk, M. and A. Pentland (1991). "Eigenfaces for recognition." J. Cognitive Neuroscience **3**(1): 71-96.
- Wilkinson, J. H. (1965). The algebraic eigenvalue problem. Oxford, Clarendon Press.
- Yuille, A. L., D. Snow, et al. (1999). "Determining generative models of objects under varying illumination: shape and albedo from multiple images using SVD and integrability." Int'l J. Computer vision **35**(3): 203-222.
- Zelnik-Manor, L. and M. Irani (1999). Multi-View Subspace Constraints on Homographies. Proc. Int'l Conf. Computer Vision, 710-715.
- Zelnik-Manor, L. and M. Irani (2002). "Multi-View Subspace Constraints on Homographies." IEEE Trans Pattern Analysis and Machine Intelligence **24**(2): 214-223.
- Zhao, L. and Y. H. Yang (1999). "Theoretical analysis of illumination in PCA-based vision systems." Pattern recognition **32**: 547-564.