



**SUPERRESOLUTION USING SPLINES AND
ANISOTROPIC DIFFUSION**

ELIZABETH BEGIN

Thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Science

**BOSTON
UNIVERSITY**

BOSTON UNIVERSITY
COLLEGE OF ENGINEERING

Thesis

**SUPERRESOLUTION USING SPLINES AND ANISOTROPIC
DIFFUSION**

by

ELIZABETH BEGIN

B.S., Clarkson University, 2004
B.S., University at Albany, 2004

Submitted in partial fulfillment of the
requirements for the degree of
Master of Science

2008

© Copyright by
ELIZABETH BEGIN
2008

Approved by

First Reader

Janusz Konrad, Ph.D.
Professor of Electrical and Computer Engineering

Second Reader

W. Clem Karl, Ph.D.
Professor of Electrical and Computer Engineering

Third Reader

Maja Bystrom, Ph.D.
Associate Professor of Electrical and Computer Engineering

Acknowledgments

First and foremost I would like to thank my advisor Professor Janusz Konrad for his constant support and encouragement throughout this research. Although he is very busy, he has always made time to answer questions and provide guidance.

In addition, a special thanks to Serdar Ince, a previous PhD student, for his collaboration on this work. Serdar provided the optical flow code utilized for motion estimation in this thesis.

I would also like to thank my thesis committee Professor Janusz Konrad, Professor W. Clem Karl and Professor Maja Bystrom for reviewing this thesis and providing feedback. In addition to their comments, I would also like to thank these professors for their excellent courses in signal and image processing.

**SUPERRESOLUTION USING SPLINES AND ANISOTROPIC
DIFFUSION
ELIZABETH BEGIN**

ABSTRACT

High-resolution images are valuable in many applications such as medical imaging, satellite imaging or video surveillance. One way to increase the resolution of an imaging sensor is to decrease the size of pixels while increasing their number. However, decreasing the size of pixels results in increased shot noise. An alternative to modifying the hardware is to improve the resolution of images by means of signal processing. This is often referred to as superresolution image reconstruction.

Superresolution from a single image relies on exploiting specific characteristics of the image, and is very difficult for general imagery. An alternative is superresolution from multiple images that are mutually offset by subpixel shifts ("shaky video"). By combining multiple images together, some of the frequency content lost in the imaging process can be recovered. Superresolution reconstruction exploits this property to obtain a high-resolution image. There are two key steps in superresolution reconstruction, motion estimation and image restoration. In order to combine the low-resolution images, so that additional spectral content can be recovered, the low-resolution images must be aligned, resulting in an irregular set of samples. Once the locations of pixels are known, the high-resolution image can be estimated through image restoration. In this thesis, we focus on the image restoration step.

To date, image restoration for superresolution has been studied in the context of iterative back-projection, projection onto convex sets, maximum *a posteriori* probability estimation, etc. The fundamental difference between different superresolution methods lies in the underlying continuous image model. One of the more recent and effective superresolution methods uses bi-cubic splines as the underlying image model, and 2-D spline fitting under thin-plate regularization. The thin-plate model applies isotropic smoothing

regardless of local image content, thus smoothing across intensity boundaries and limiting the performance.

In this thesis, we extend this method by introducing spatially-adaptive, or anisotropic, regularization. The anisotropic weighting uses local image gradient to adapt the smoothing to intensity edge orientation. Since gradient of the high-resolution image is unknown, we estimate it from the low-resolution images. Although this is only a coarse gradient estimate, it has a beneficial impact on superresolution performance.

For data sets with known motion parameters, our results confirm the edge enhancing effects of the spline-based anisotropic approach. The improvements are verified both quantitatively and visually. However, for data that requires motion estimation this performance gain over isotropic methods is significantly reduced. Therefore, we conclude that without precise motion estimation the benefits of superresolution with anisotropic diffusion are limited.

Contents

1	Introduction	1
1.1	Superresolution	1
1.2	Thesis Overview	4
2	Review of Prior Work	5
2.1	Frequency-Domain Alias Removal	6
2.2	Iterative Back-Projection	8
2.3	Maximum <i>A Posteriori</i> Estimation	10
2.4	Projection onto Convex Sets	12
3	Preliminary Material	14
3.1	Splines	14
3.2	Conjugate Gradient	18
3.3	Optical Flow	20
4	Problem Formulation and Solution	24
4.1	Image Model	24
4.2	Cost Function	26
4.2.1	Error Term	26
4.2.2	Regularization Term	27
4.2.3	Anisotropic Diffusion	32
4.2.4	Gradient Estimation	37
4.2.5	Minimization	40
5	Experimental Results	41
5.1	Synthetic Motion	41

5.1.1	Randomly-Sampled Images	43
5.1.2	Global Translational Motion	44
5.2	Real Motion Data Sets	68
6	Conclusions	72
6.1	Discussion of Results	72
6.2	Suggestions for Future Work	73
	References	75

List of Tables

5.1	Minimum MSE for Each Test Set with Perfect Motion	50
5.2	Minimum MSE for Each Test Set with Estimated Motion	51

List of Figures

1-1	Superresolution Video Restoration	2
1-2	Superresolution Overview	3
3-1	1-D B-Splines	16
3-2	2-D B-Splines	17
3-3	Spline Interpolation Example	18
3-4	Optical Flow Example	23
4-1	Cubic B-spline Derivatives	30
4-2	Discrete and Continuous Comparison	33
4-3	Gradient of <i>Stripes</i> Image	34
4-4	Anisotropic Diffusion Comparison	37
4-5	1-D Isotropic and Anisotropic Example	38
5-1	HR Images Used to Generate Synthetic-Motion Data Sets	43
5-2	Mean Square Error for Randomly-Sampled Data Set	44
5-3	Results of Reconstruction for Randomly-Sampled <i>Squares</i> Image	46
5-4	Pre-Filter and No Pre-Filter Comparison	48
5-5	LR and HR Pixel Locations	49
5-6	MSE Results for <i>Squares</i> Image	54
5-7	MSE Results for <i>Brick</i> Image	56
5-8	MSE Results for <i>Jeep</i> Image	58
5-9	SR Reconstruction Results for <i>Squares</i> Image	61
5-10	SR Reconstruction Results for <i>Brick</i> Image	64
5-11	SR Reconstruction Results for <i>Jeep</i> Image	67

5-12 Optical Flow Vectors	69
5-13 Isotropic and Anisotropic SR Reconstruction Results for a Real Video Sequence	70
5-14 Zoomed-in Region of Isotropic and Anisotropic SR Reconstruction	71

Nomenclature

A	...	Matrix Relating y to x Through Image Formation Model
A^{BP}	...	Back-Projection Operator
α	...	Regularization Weight
$b(x)$...	Continuous 1-D Spline Function
$b(x, y)$...	Continuous 2-D Spline Function
$\beta^n(x)$...	n^{th} Order 1-D B-Spline
C	...	Spline Coefficients
C_k	...	Convex Constraint k
$[\delta_{xn}, \delta_{yn}]$...	Spatial Shift for Global-Translational Motion for Frame n
ϕ'_{xx}	...	Discrete Convolution Kernel for 2^{nd} Spline-Based Derivative in x
ϕ'_{xy}	...	Discrete Convolution Kernel for Spline-Based Derivative in x then y
ϕ'_{yy}	...	Discrete Convolution Kernel for 2^{nd} Spline-Based Derivative in y
D_{xx}	...	Convolution Matrix of ϕ'_{xx}
D_{xy}	...	Convolution Matrix of ϕ'_{xy}
D_{yy}	...	Convolution Matrix of ϕ'_{yy}
\mathcal{E}	...	Optical Flow Cost Function
\mathcal{E}_c	...	Optical Flow Error Term
\mathcal{E}_b	...	Optical Flow Regularization Term
$E_R(C)$...	Regularization/Prior Term
$E_L(C)$...	Log-Likelihood Term
$E(x, y, t)$...	Image Intensity Values

$F(u, v)$...	Continuous Fourier Transform of $f(x, y)$
$f(x, y)$...	Continuous Underlying Image
$g(I) = e^{-\left(\frac{ I }{K}\right)^2}$...	Anisotropic Weighting Function, K is a Constant
$H[\mathbf{n}]$...	Anisotropic Weighting Matrix for Pixel \mathbf{n}
H_{xx}	...	$H[1, 1]$ for All \mathbf{n} Stacked in Lexicographic Order
H_{xy}	...	$H[1, 2]$ for All \mathbf{n} Stacked in Lexicographic Order
H_{yy}	...	$H[2, 2]$ for All \mathbf{n} Stacked in Lexicographic Order
I_x	...	Derivative of Spline Function in x
I_y	...	Derivative of Spline Function in y
$I_{xx} = D_{xx}C$...	2^{nd} Derivative of Spline Function in x
$I_{xy} = D_{xy}C$...	Derivative of Spline Function in x then y
$I_{yy} = D_{yy}C$...	2^{nd} Derivative of Spline Function in y
$J(C) = E_L(C) + \alpha E_R(C)$...	Cost Function
Λ	...	High-Resolution Spline Coefficient Lattice
$[M_1, M_2]$...	Dimensions of High-Resolution Image
$\mathbf{n} = [n_x n_y]^T$...	Discrete Spline Coefficient Indices
N	...	Additive Noise
$[N_1, N_2]$...	Dimensions of Low-Resolution Images
p	...	Continuous Spline Derivative Kernel
P	...	Convolution Matrix of p
P_k	...	Projection Operator k
$\phi^n(x, y) = \beta^n(x)\beta^n(y)$...	n^{th} Order 2-D B-Spline
Φ	...	B-Spline Matrix Related y to C
R	...	Number of Low-Resolution Images
$s[x]$...	Discrete Samples of Continuous Signal
$[T_x, T_y]$...	Sampling Period in x and y

$[v_x[\mathbf{n}], v_y[\mathbf{n}]]$...	Optical Flow Vectors for Pixel \mathbf{n}
$[V_x, V_y]$...	$[v_x, v_y]$ for All Pixels in Lexicographic Order
x	...	Estimate of High-Resolution Image
$[X_{HR}, Y_{HR}]$...	x and y Coordinates of High-Resolution Lattice
$[X_r, Y_r]$...	R Regularly Spaced Low-Resolution Lattices
$[X_{LR}, Y_{LR}] = [X_r + V_x, Y_r + V_y]$...	Low-Resolution x and y Pixel Locations
y	...	LR Images Stacked in Lexicographic Order

Chapter 1

Introduction

The objective of superresolution reconstruction is to recover a high-resolution (HR) frame from multiple low-resolution (LR), degraded images. The basic premise of superresolution is that there exist LR images which are spatially offset by subpixel amounts (Park et al., 2003). Superresolution refers to the ability to recover higher frequency content than what is available in any of the individual LR images. This additional frequency information is provided by the offset of each of the LR images. Superresolution can be used to construct a single HR image or multiple images for a video sequence. For video reconstruction, a moving window of frames is utilized to create each of the HR video frames (see Figure 1.1). Superresolution consists of two steps, both of which are extensive fields of study, image registration and image restoration. In this thesis, we focus on image restoration.

1.1 Superresolution

It is evident that SR is valuable on the basis that a user or system is able to obtain more detailed information from a higher resolution image. It is not always feasible to increase the resolution of a sensor through hardware modifications. For example, although the resolution of an image can be improved by decreasing the size of a pixel and increasing the number of pixels in a sensor, this solution is not always an option due to shot noise (Park et al., 2003). As pixel size decreases, the amount of light captured by each pixel also decreases, resulting in the possibility of shot noise. Superresolution is an effective way to increase resolution through signal processing, an excellent alternative to modifying the dimensions of the pixels.

There exists a variety of applications which can benefit from SR reconstruction. For

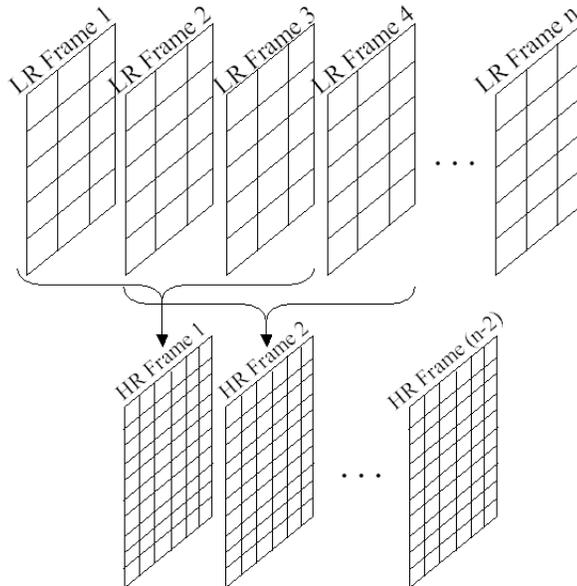


Figure 1.1: Superresolution video restoration uses a moving window of frames to construct the HR images.

example, in medical imaging, such as magnetic resonance imaging (MRI), a higher resolution image helps doctors more accurately diagnose a patient. In radar and satellite applications, SR improves a sensor’s ability to distinguish separate objects and correctly classify them. In digital camera and video applications, a user may want to zoom-in on a particular region of an image. Through the use of SR, artifacts introduced by the zoom can be suppressed. SR can be applied to almost any imaging application in which multiple spatially offset images are being acquired.

The primary requirement for SR reconstruction is the existence of multiple LR images shifted by subpixel amounts. Integer-shifted images contain the same information and therefore do not provide any additional information for the SR reconstruction of the underlying HR image. Combining multiple subpixel-shifted images results in LR samples with higher spatial density than an individual LR image. Ideally, all samples, derived from R frames, are unique in position, meaning there is no spatial redundancy in samples over time. An implicit assumption is that the scene is static, meaning there is no local motion

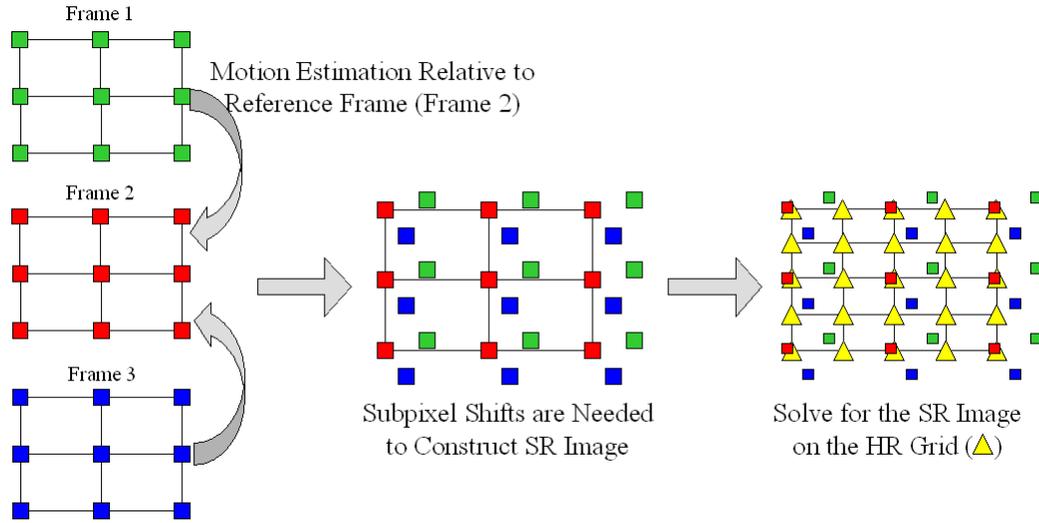


Figure 1.2: Basic steps for superresolution.

between images. Prior to image restoration, the low resolution pixels must be aligned through image registration (see Figure 1.2), often referred to as motion estimation, the term we will use in this thesis. The synthetic-motion data sets presented in this thesis (Chapter 5) consist of globally translated LR images.

According to Hadamard's definition of ill-posedness, image restoration for superresolution is considered an ill-posed inverse problem (Borman, 2004; Hadamard, 1923). A problem is considered well-posed by Hadamard if the solution:

1. exists,
2. is unique,
3. and depends continuously on the data.

Given a set of LR frames, the goal of SR is to recover the original underlying scene from the LR observations. Superresolution is considered ill-posed due to the fact that there is no unique solution. Multiple solutions occur when the size of the restored image is larger than the set of observed images. Therefore, there are multiple underlying scenes which

could yield the same set of observed images. Some superresolution problems may also fail to meet the other two conditions, existence of a solution and continuous dependence. However, the superresolution applications in this thesis focus on basic cases which do not have this degree of ill-posedness. We will be examining simple data sets in which the observations are digital images of a well-behaved scene with limited noise.

Superresolution can be broken down into two main steps, frame registration (motion estimation) and image restoration. Motion estimation and image restoration are two broad areas of image processing and are often researched independent of one another. Many authors assume perfect motion estimation and focus solely on image restoration, or vice versa. In a number of cases we will use perfect motion parameters to verify the implementation of our algorithm. The perfect motion cases will allow us to analyze the effectiveness of the SR implementation without biases caused by inaccuracies in the motion estimation. We will then examine the impact of motion estimation on the SR reconstruction. The motion estimation algorithm utilized will be discussed in Chapter 3.

1.2 Thesis Overview

This thesis is organized as follows. First, Chapter 2 contains a review of prior SR work. Next, in Chapter 3 we cover some preliminary topics which are utilized throughout the remainder of the thesis. In Chapter 4 we formulate the problem and propose a solution based on the conjugate gradient approach. In Chapter 5 experimental results are shown. The SR method presented utilizes splines and anisotropic diffusion to obtain a high-quality image. This method will be compared to a simple cubic interpolation scheme as well as the isotropic reconstruction using mean square error (MSE) and visual comparison. The results of the anisotropic superresolution using real video sequences are also presented. Lastly, Chapter 6 includes concluding remarks and suggestions for future work.

Chapter 2

Review of Prior Work

In this section, we review previous superresolution work. The methods can be broken down into two broad categories, frequency-domain methods and spatial-domain methods. Spatial-domain methods are the most flexible when it comes to the inclusion of *a priori* knowledge and are therefore the most extensively studied. The benefits and drawbacks of the methods are compared.

Frequency-domain methods, in general, are fairly simple and efficient. However, they are often very restrictive when it comes to the inclusion of *a priori* information and therefore are frequently outperformed by spatial-domain methods. Since the frequency-domain methods rely on the shifting properties of the Fourier transform to estimate motion, the complexity of the motion model for superresolution is very limited. In terms of frequency-domain restoration methods, the objective of SR reconstruction is to obtain frequency information beyond the Nyquist rate of the low-resolution image. A single LR frame is aliased if the sample rate is below the Nyquist rate of the underlying HR image (twice the bandwidth) (Borman, 2004). However, superresolution utilizes multiple frames which are slightly offset, and therefore obtains additional spatial information which can recover aliased data and form a high-resolution image.

Spatial-domain techniques are a bit more flexible than frequency-domain methods. One of the main benefits of spatial-domain methods is that they allow for the inclusion of *a priori* constraints through regularization. Regularization alleviates the degree of ill-posedness of the superresolution inverse problem by constraining the space of possible solutions based on valuable *a priori* information. Through regularization, we can make use of known characteristics of the desired solution such as smoothness or known edge

locations. Spatial-domain methods also allow for more complex motion models. Since spatial methods do not rely on Fourier shifting properties, the motion of the low resolution frames is not restricted to basic motion models. Although spatial-domain methods are generally superior to frequency-domain techniques in terms of the reconstruction quality, they usually require a bit more computation and are not as straightforward as frequency-domain methods.

2.1 Frequency-Domain Alias Removal

The frequency-domain methods for resolving superresolution images rely on the shifting and sampling properties of the Fourier transform (Borman, 2004; Tsai and Huang, 1984). Although this reliance allows for intuitive and fast results, it significantly restricts the complexity of the motion and observation models. The objective of superresolution in the frequency domain is to obtain an image which contains information with higher frequency than the Nyquist rate of the low resolution frames. Tsai and Huang published the first work on the topic of superresolution in 1984 (Tsai and Huang, 1984). Their work was motivated by the need to construct high-resolution images from a Landsat satellite. The images from the satellite were slightly spatially offset, therefore each of the images contained new information about the scene although they were generally of the same region. Tsai and Huang utilized this information to obtain a HR image.

Tsai and Huang exploited the shifting properties of the Fourier transform to perform image restoration. A shift in position translates to a frequency-domain shift in phase. It is this property which makes the frequency-domain SR reconstruction possible. The authors proposed a method which relates the continuous world, modeled as a continuous frequency-domain function, to the discrete Fourier transform coefficients of the LR images. The continuous Fourier representation is related to the discrete Fourier transform through aliasing. Naturally, a sampled signal is band-limited to the sampling frequency. Therefore, any part of the real world which contains information higher than the sampling frequency will be aliased in the LR images.

Let $f(x, y)$ represent the continuous underlying image. The globally-translated image of frame number n can be represented as,

$$f_n(x, y) = f(x + \delta_{xn}, y + \delta_{yn}). \quad (2.1)$$

where $(\delta_{xn}, \delta_{yn})$ represents the spatial shift of the n^{th} frame in x and y . Therefore, if $F(u, v)$ represents continuous Fourier transform of a $f(x, y)$, then frame n , which is shifted relative to this frame, can be written in frequency domain as:

$$F_n(u, v) = e^{i2\pi(\delta_{xn}u + \delta_{yn}v)} F(u, v) \quad (2.2)$$

Tsai and Huang assume impulse sampling of $f(x, y)$, which results in the following formulation for the observed images,

$$y_n[n_x, n_y] = f(T_x n_x + \delta_{xn}, T_y n_y + \delta_{yn}), \quad (2.3)$$

where $n_x \in (0, 1, \dots, N_x - 1)$ and $n_y \in (0, 1, \dots, N_y - 1)$. N_x and N_y are the dimensions of the low-resolution sampled images, T_x and T_y are the sampling periods in x and y . The relationship between the continuous Fourier transform and the discrete Fourier transform of the shifted and down sampled images can therefore be expressed as (Tsai and Huang, 1984; Smith, 2007):

$$Y_n[u, v] = \frac{1}{T_x T_y} \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} F_n\left(\frac{u}{T_x N_x} + \frac{m}{T_x}, \frac{v}{T_y N_y} + \frac{n}{T_y}\right) \quad (2.4)$$

where $[u, v]$ represent discrete frequency pairs. If the original continuous function $f(x, y)$ is bandlimited, then the infinite sum can be reduced to some finite number of sums. Tsai and Huang assume this to be true and as a result, are able to construct a matrix set of equations combining equations (2.4) and (2.2).

$$Y = \Psi F \quad (2.5)$$

Y is the vector of discrete Fourier transform coefficients for the observed low resolution

image, Ψ is a matrix which relates coefficients of the discrete-space Fourier transform to samples of the unknown continuous Fourier transform coefficients contained in vector F . In order to construct Ψ the global motion parameters, δ_x and δ_y , must be known. Therefore, this SR reconstruction method is broken into two steps (as are most methods), motion estimation and image restoration. In most cases, the accuracy of reconstruction is limited by motion estimation.

There are several drawbacks to the alias removal method proposed by Tsai and Huang. The major hindrance being the requirement of global translational motion. Global translational motion in the spatial domain appears as phase shifts in the frequency domain. However, more often than not, global translational motion models are not accurate enough to represent the motion in image sequences. It is this fact which makes the spatial domain methods more attractive.

Another drawback of Tsai and Huang’s method is that it assumes impulse sampling. Any blur added by the camera is not accounted for in Ψ . Tekalp, Ozkan and Sezan address this issue by including a point-spread function in the observation model (Tekalp et al., 1992). Therefore, the low-resolution sampled images are modeled as impulse sampled points of the blurred continuous world. The blur kernel is convolved with the continuous representation and subsequently down-sampled to yield the low resolution images. There are a variety of other frequency domain approaches including (Kim et al., 1990; Kim and Su, 1993; Prendergast and Nguyen, 2005).

2.2 Iterative Back-Projection

In 1990, Irani and Peleg formulated the iterative back-projection algorithm for superresolution (Irani and Peleg, 1990). Back-projection is commonly associated with tomography and the formulation for superresolution is similar to the tomography implementation. The basic idea is to estimate how the low-resolution images were formed and then construct LR images from this model. The LR observations are simulated for each estimate of the HR image. The residual error between the simulated image and the measured low-resolution

image is computed. Back-projection is utilized to update the estimate of the underlying scene using the residual error. In order to ensure that the high-resolution image is being updated properly, it is imperative that the observation model be accurate.

The basic relationship between the observed LR images y and the underlying HR image x can be written as,

$$y = Ax \tag{2.6}$$

where y contains the low resolution images stacked in lexicographic order and x represents the high resolution underlying image also in lexicographic order. A is a matrix which relates the high-resolution samples to the low-resolution samples through the image formation model. For example, if the image formation were modeled as a Gaussian point spread function (PSF), the entries of the A matrix would contain samples of the continuous Gaussian kernel. Once A and an estimate of x are obtained, the observed LR images are easily simulated through linear equation (2.6).

Let y^n represent the n^{th} estimate of y . For each estimate of y^n , the residual is computed and back projected using the back projection operator A^{BP} . The back projection operator is usually an estimate of the inverse of the forward model A . The estimate of the underlying HR image, x , is updated using the following update equations:

$$x^{n+1} = x^n + A^{BP}(y - y^n) \tag{2.7}$$

$$= x^n + A^{BP}(y - Ax^n) \tag{2.8}$$

One drawback of this approach is the lack of regularization. Regularization limits the range of possible solutions and promotes convergence. Without regularization, the above iterative steps may not converge and could vary between several solutions.

As with most superresolution techniques, the iterative back-projection solution is only as accurate as the model. Inaccuracies in the image model will limit the possibility of fully recovering the HR image. The residual estimated between the simulated LR images y^n and measured LR images y will always have some bias introduced by the inaccuracies

in A . Therefore, the HR reconstruction will never be able to fully recover the underlying scene if the residual errors being back projected are inaccurate. Additional iterative back-projection methods are proposed in (Irani and Peleg, 1993; Mann and Picard, 1994; Dai et al., 2007)

2.3 Maximum *A Posteriori* Estimation

Schultz and Stevenson first introduced the Bayesian approach to superresolution in 1994 (Schultz and Stevenson, 1994). The Maximum *A Posteriori* Probability (MAP) method is a common method for solving stochastic problems. MAP maximizes the *a posteriori* probability density function by utilizing Bayes' rule. The MAP formulation allows for the inclusion of *a priori* information and therefore helps regularize the solution. As before, let y represent the observed low-resolution image stacked in lexicographic order. Through basic stochastic modeling, y can be written as

$$y = Ax + N \quad (2.9)$$

where A is the matrix that relates the high resolution underlying image x to y , and N is additive noise. The authors assume that the motion has been estimated, allowing them to construct A . Using Bayes' rule, which relates conditional probabilities, the MAP expression can be formulated as follows,

$$x_{MAP} = \operatorname{argmax}_x (P(x|y)) \quad (2.10)$$

$$x_{MAP} = \operatorname{argmax}_x \left(\frac{P(y|x)P(x)}{P(y)} \right) \quad (2.11)$$

where x_{MAP} is the MAP estimate of the HR image. The maximum is independent of y , therefore the denominator can be removed from the equation. Since both remaining probabilities are positive, the *log* can be taken resulting in the following expression:

$$x_{MAP} = \operatorname{argmax}_x (\log(P(y|x)) + \log(P(x))) \quad (2.12)$$

The first term is referred to as the log-likelihood function and the second is the prior term. Without the prior, this formulation would be the Maximum Likelihood (ML) estimation. Assuming a Gaussian noise model for the conditional probability yields the following equation:

$$P(y|x) = \frac{1}{(\sqrt{2\pi}\sigma)^{M_1M_2}} e^{-\frac{1}{2\sigma^2}\|y-Ax\|^2} \quad (2.13)$$

where $[M_1, M_2]$ represents the dimensions of the HR image. The \log of equation (2.13) results in the following equation:

$$\log(P(y|x)) = \log\left(\frac{1}{(\sqrt{2\pi}\sigma)^{M_1M_2}}\right) - \frac{1}{2\sigma^2}\|y-Ax\|^2 \quad (2.14)$$

The first term of the log likelihood function is a constant and can be dropped from the overall cost function. Similarly, $\frac{1}{2\sigma^2}$ is also a constant and will only scale the cost function. This term may also be dropped, resulting in the following maximization:

$$x_{MAP} = \underset{x}{\operatorname{argmax}}(-\|y-Ax\|^2 + \log(P(x))) \quad (2.15)$$

The next step is to choose a prior distribution, $P(x)$. A typical choice for a prior distribution is a Markov Random Field (MRF) model which is represented by a Gibbs distribution:

$$P(x) = \left(\frac{1}{Z}\right) e^{-\alpha E_R(x)} \quad (2.16)$$

$$\log(P(x)) = \log\left(\frac{1}{Z}\right) - \alpha E_R(x) \quad (2.17)$$

where $E_R(x)$ is a Gibbs energy function. The energy function, applied to the solution x , is chosen based on the desired solution constraints, some examples include an L_2 norm, the magnitude of the gradient or an L_1 norm. α is the regularization weight which determines the behavior of the solution and the amount each term affects the solution. Once again the constant, $\frac{1}{Z}$, can be dropped from the maximization. Substituting equation (2.17) into

equation (2.15) yields the following MAP maximization:

$$x_{MAP} = \operatorname{argmax}_x(-\|y - Ax\|^2 - \alpha E_R(x)) \quad (2.18)$$

Flipping the signs, equation (2.18) can be rewritten as a minimization:

$$x_{MAP} = \operatorname{argmin}_x(\|y - Ax\|^2 + \alpha E_R(x)) \quad (2.19)$$

Notice that the MAP estimate of x , using an L_2 norm as the prior model, results in a Tikhonov formulation. Therefore, a MAP formulation with a Gaussian conditional probability function, a MRF prior distribution and a L_2 norm regularization function results in a Tikhonov minimization (Karl, 2000).

$$x_{MAP} = \operatorname{argmin}_x(\|y - Ax\|^2 + \alpha \|x\|_2^2) \quad (2.20)$$

To minimize the cost function, the derivative with respect to x is set equal to zero. The solution to the set of normal equations can then be found using an iterative minimization method such as the conjugate gradient algorithm. There are a variety of modifications to this method which have been proposed, using different prior models, observation models and minimization techniques (Cheeseman et al., 1994; Borman and Stevenson, 1999; Segall et al., 2004; Hardie and Droege, 2007; Chantas et al., 2007). One of the benefits of the MAP method is that it allows for the inclusion of *a priori* information, thereby limiting the range of possible solutions to the ill-posed inverse problem.

2.4 Projection onto Convex Sets

The two dimensional implementation of Projection onto Convex Sets (POCS) was first proposed by Stark and Oskoui in 1989 and ever since has become a very popular method in superresolution (Stark and Oskoui, 1989). POCS restricts the space of solutions to the intersection of a set of convex constraints. Therefore, within the set of all possible super-resolution images, the solution must lie within the intersection of the convex constraints.

The set of constraints C_k are determined by modeling how the images are captured. The simplicity of the POCS method allows for easy addition of various regularization terms. Examples of constraints include positivity, minimum difference between the estimated and observed data and smoothness. POCS, like most methods, is solved iteratively to find the solution which satisfies the set of constraints.

$$x^{n+1} = P_m P_{m-1} \dots P_1 x^n \quad (2.21)$$

where P_k is the operator which projects the current estimate of x onto the convex set C_k . Some commonly used constraints are limiting the energy of the solution,

$$C_{energy} = \{x : \|x\|^2 \leq E\} \quad (2.22)$$

where E is a constant. Another example is to limit the solution to lie within some distance to a rough estimate of the solution, such as a cubic interpolation estimate of the HR image, \hat{x} .

$$C_{estimate} = \{x : \|x - \hat{x}\|^2 \leq E\} \quad (2.23)$$

There are a variety of other constraints which may be included in the POCS method (Elad and Feuer, 1997; Patti and Altunbasak, 1998; Ozkan et al., 2001; Hsu et al., 2004; Stasinski and Konrad, 2006).

Another benefit of the POCS method is that nearly any kind of motion model may be used provided that the motion estimation is accurate. As with the previously-discussed methods, it is assumed that the motion is previously estimated.

Chapter 3

Preliminary Material

Prior to revealing the details of the superresolution reconstruction formulation of this thesis, there are several preliminary topics that need to be reviewed. First, splines are examined. They will later be used as the basis for the SR image model. Second, a brief description of the conjugate gradient algorithm is covered. The conjugate gradient method will be used to find a solution which minimizes the cost function developed in Chapter 4. Lastly, a review of the optical flow algorithm developed by Horn and Schunk will follow. Although motion estimation is not the focus of this work, it is a critical step to obtaining a superresolution image and consequently worth some discussion.

3.1 Splines

Splines were first described by Schoenberg in his 1946 paper, just before the introduction of Shannon's sampling theory (Unser, 1999). Schoenberg demonstrated the effectiveness of splines for interpolation of a regularly-spaced sample set. Although splines were introduced prior to the work of Shannon, the band-limited functions introduced by Shannon were far more popular until the 1960's. Splines became very popular in a number of applications, however the area of signal processing did not begin to take advantage of splines until more recently. According to Unser, the interest among the signal processing community began with the increased popularity of wavelet theory. There are several properties which make splines an attractive choice for interpolation and continuous-signal representation.

The first and foremost beneficial property of splines is that they provide a simple way of obtaining a continuous representation of a discrete signal. There are a number of signal processing applications in which it is necessary to estimate or model the continuous signal

from a set of discrete samples. In the case of superresolution, the LR images captured are discrete samples of the continuous scene. Ideally, we would like to recover the exact function which describes the continuous signal, however in most cases this is highly improbable given the complexity of an image. Therefore, we approximate the scene using a continuous spline function. Thanks to the continuous representation, the HR image can be constructed at any resolution by sampling the spline function.

Splines are smoothly-connected, well-behaved, piecewise polynomials. They are designed such that they are $n - 1$ times continuously differentiable, where n represents the degree of the spline. A spline representation is a linear combination of equally-spaced basic splines, also referred to as B-splines. Therefore, although they are continuous functions, they can be solved for discretely.

$$b(x) = \sum_{k \in \mathbb{Z}} C[k] \beta^n(x - k) \quad (3.1)$$

where $b(x)$ represents the continuous spline, C represents the 1-D spline coefficients, β^n represents the n^{th} order B-spline. B-splines are compactly supported. Consequently, there is a limited computational cost of using splines, an important property to consider when modeling a signal. A continuous representation of a signal is easily described as a linear combination of these compactly-supported B-splines, making the implementation of splines for signal modeling straightforward.

B-splines are defined as symmetric functions which result from the $(n + 1)$ -fold convolution of a rectangular pulse β^0 ,

$$\beta^0(x) = \begin{cases} 1, & -\frac{1}{2} < x < \frac{1}{2} \\ \frac{1}{2}, & |x| = \frac{1}{2} \\ 0, & \text{otherwise} \end{cases} \quad (3.2)$$

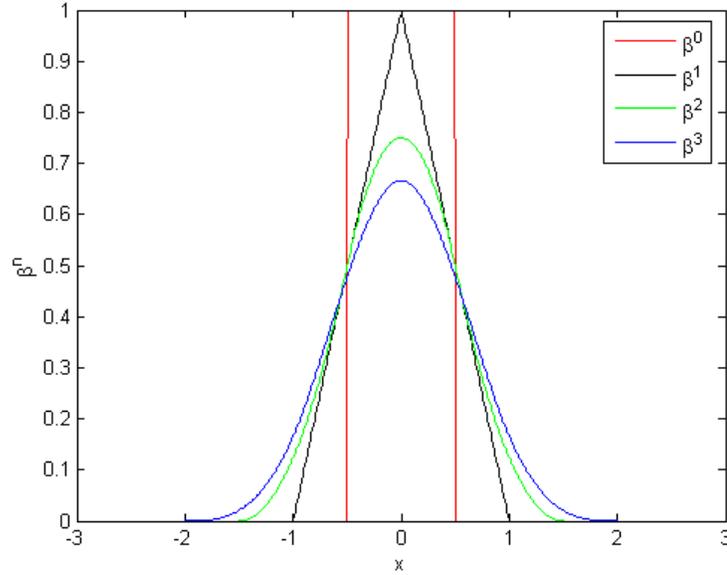


Figure 3-1: 1-D B-splines of order 0 to 3

Therefore, β^0 must be convolved $(n + 1)$ times to obtain a B-spline of degree n ,

$$\beta^n(x) = \underbrace{\beta^0 * \beta^0 * \dots * \beta^0(x)}_{(n+1)} \quad (3.3)$$

Cubic splines are the most popular for interpolation,

$$\beta^3(x) = \begin{cases} \frac{2}{3} - |x|^2 + \frac{|x|^3}{2}, & 0 \leq |x| < 1 \\ \frac{(2-|x|)^3}{6}, & 1 \leq |x| < 2 \\ 0, & 2 \leq |x| \end{cases} \quad (3.4)$$

Figure 3.1 shows a plot of B-splines from order 0 to 3.

In order to model an image, we must use 2-D splines. A 2-D B-spline, ϕ^n , can easily be constructed from the 1-D B-spline functions using a tensor product in x and y .

$$\phi^n(x, y) = \beta^n(x)\beta^n(y) \quad (3.5)$$

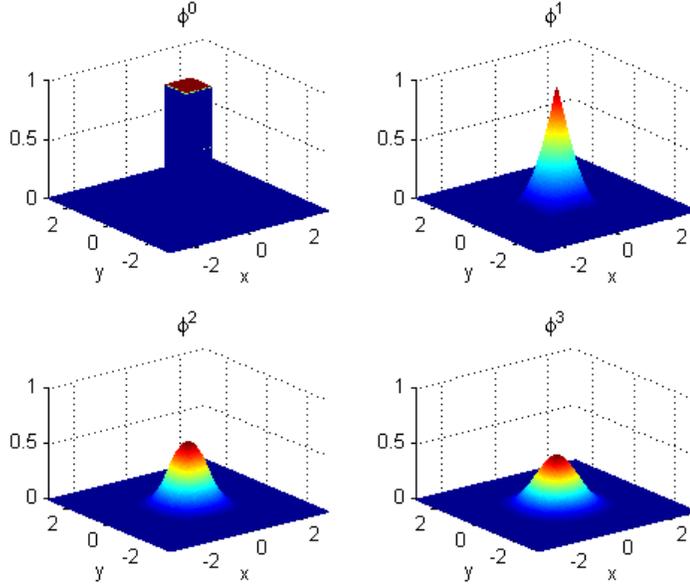


Figure 3-2: 2-D B-splines of order 0 to 3

A 2-D spline representation is the sum of shifted and scaled 2-D B-splines,

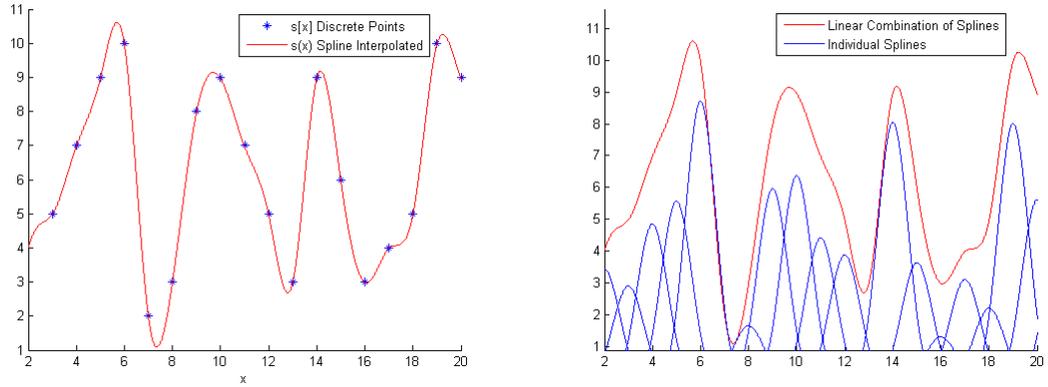
$$b(x, y) = \sum_{\mathbf{n} \in \Lambda} C[n_x, n_y] \phi^n(x - n_x, y - n_y) \quad (3.6)$$

where Λ represents a regularly sampled grid and $\mathbf{n} = [n_x, n_y]^T$.

In the formulation proposed in this thesis, the spline coefficients of the underlying high-resolution image will be solved for to obtain a continuous representation of the scene. Once the coefficients are estimated, they can be convolved with the 2-D B-spline function to acquire the super-resolved image.

Solving for the coefficients of splines of degree 0 or 1 is straightforward since their coefficients are equal to signal sample values. However, solving for a higher order spline is less trivial. To determine the coefficients for higher-order splines, a digital filtering technique was developed (Unser, 1999). Two filters are used to determine the coefficients, c^+ which is causal and the other, c^- , which is anti-causal,

$$c^+[k] = s[k] + z_1 c^+[k - 1], \text{ for } (k = 1, \dots, N - 1) \quad (3.7)$$



(a) Cubic spline interpolation of a discrete signal

(b) Individual b-splines shown in blue. Linear combination of splines shown in red.

Figure 3.3: (a) Result of the cubic spline filtering algorithm; and (b) plot of the contribution of each spline to the linear combination.

$$c^- [k] = z_1(c^- [k + 1] - c^+ [k]), \text{ for } (k = N - 2, \dots, 0) \quad (3.8)$$

This iterative algorithm is initialized with

$$c^+ [0] = \frac{1}{(1 - z_1^2)} \sum_{k=0}^{2N-3} s(k)z_1^k \quad (3.9)$$

$$c^- [N - 1] = \frac{z_1}{(1 - z_1^2)} (c^+ [N - 1] + z_1 c^+ [N - 2]) \quad (3.10)$$

where $z_1 = -2 + \sqrt{3}$, N is the number of samples in the discrete signal s and the final coefficients are obtained from $C = 6c^-$ (Unser, 1999). Figure 3.3a shows a 1D example of cubic spline interpolation using this filtering algorithm. Figure 3.3b demonstrates the contribution of each spline to the overall linear combination.

3.2 Conjugate Gradient

The conjugate gradient algorithm is one of the most common iterative methods for solving systems of linear equations (Shewchuck, 1994). It is effective in solving equations of the form

$$y = Ax \quad (3.11)$$

where x and y are column vectors and A is a matrix. The method of conjugate gradients works best for cases in which A is positive-definite, otherwise the minimization may have difficulty determining the solution. The conjugate gradient algorithm can be used to solve systems where A is not positive-definite. In such cases, there may be no solution to the equation, however, the conjugate gradient method will find a solution that minimizes the error. The method of conjugate gradient is a combination of two other methods, steepest descent and conjugate directions.

The method of steepest descent begins at a point and takes a step in the steepest direction of $q(x)$. In this case, $q(x)$ represents the quadratic form of a vector (Shewchuck, 1994),

$$q(x) = \frac{1}{2}x^T Ax - y^T x + c \quad (3.12)$$

where c is a constant. The point which minimizes $q(x)$ provides the solution to $y = Ax$. Along the line of steepest descent, the point which minimizes $q(x)$ is chosen as the next point. Therefore, each step is taken in the direction of the residual. The residual is defined as the $r = y - Ax$, and indicates how far the current step is from the correct value of y . The error, e , indicates how far the current location is from the solution, and is related to the residual through A , $r = -Ae$. The method of steepest descent focuses on stepping in the direction of the residual, therefore, the search direction is always orthogonal to the previous search direction. The drawback of steepest descent is that it often has to take steps in the same direction as previous steps.

The method of conjugate directions is very similar to the steepest descent algorithm. However, the direction of each step is chosen to be A -orthogonal or conjugate. Two vectors, d_1 and d_2 , are considered A -orthogonal if

$$d_1^T A d_2 = 0 \quad (3.13)$$

By stepping in the conjugate direction, the number of steps required to find the minimum is limited to n steps, where n represents the length of the vector x . With each step

an error term along a particular direction is eliminated. Along each conjugate step, the minimum point of $q(x)$ is chosen. The difficulty with conjugate directions is that all of the previous search directions must be stored in order to create a new search vector which is A -orthogonal to all previous directions. The conjugate gradient method alleviates this problem.

The method of conjugate gradient combines both steepest descent and conjugate directions. Conjugate gradient steps in a direction which is constructed by conjugate of the residual (Shewchuck, 1994). As a result, the residual of each step is orthogonal to the previous search direction and the previous residuals. The conjugate gradient method converges in n steps, the same number of steps as the conjugate directions method. The residual of each step is already A -orthogonal to the previous search directions, therefore there is no need to store the previous search directions. It is this property which makes conjugate gradient one of the more popular minimization algorithms. Since the previous search directions do not need to be stored, the number of computations per iteration is reduced.

In this thesis, the conjugate gradient algorithm will be utilized to minimize the superresolution cost function. To verify the convergence of the algorithm, the normalized residual is compared to a residual threshold. When the residual is below the desired threshold, the minimization is considered completed.

3.3 Optical Flow

One of the most critical steps for superresolution reconstruction is frame registration, also referred to as motion estimation. Inaccuracies in motion estimation will result in artifacts in the HR image. If pixels are not mapped to the correct location, the algorithm will diffuse these inaccuracies to the neighboring pixels of the HR image. To create a HR image, a set of LR frames are chosen from a video sequence. The motion of each frame is estimated relative to some reference frame (typically the center frame). The motion estimation algorithm utilized for this work is the Horn and Schunck optical flow method

(Horn and Schunck, 1981).

Given two images, optical flow describes the motion of each of the pixels between the two images. The Horn and Schunck method of estimating optical flow applies a smoothness constraint to the solution. This method employs a minimization function to solve for the motion parameters. This function uses image intensity values $E(x, y, t)$ and velocity horizontal and vertical components $u = \frac{dx}{dt}$ and $v = \frac{dy}{dt}$, respectively. The first term of the Horn and Schunck cost function is formulated on the basis that the intensity of a pixel will not change between the images, it will only move. Therefore, the first term seeks to make the change in intensity of paired pixels between the two images

$$\mathcal{E}_b = E_x u + E_y v + E_t \quad (3.14)$$

as close to zero as possible, where E_x , E_y and E_t are the derivatives of image intensity values in x , y and t (time). The second term is designed to apply smoothness to the solution. Pixels of most images do not move independently of one another and therefore their movement between frames should be dependent on the neighboring pixels. Thus, a regularization term is introduced:

$$\mathcal{E}_c^2 = \left(\frac{\partial u}{\partial x}\right)^2 + \left(\frac{\partial u}{\partial y}\right)^2 + \left(\frac{\partial v}{\partial x}\right)^2 + \left(\frac{\partial v}{\partial y}\right)^2 \quad (3.15)$$

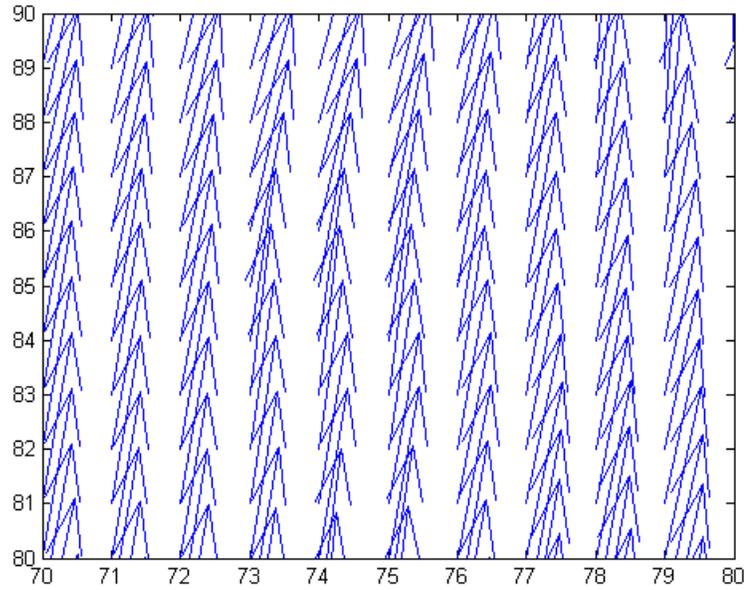
which is small only for locally smooth velocities. Combining the two terms, the total error to be minimized is,

$$\mathcal{E} = \iint (\mathcal{E}_b^2 + \alpha^2 \mathcal{E}_c^2) dx dy \quad (3.16)$$

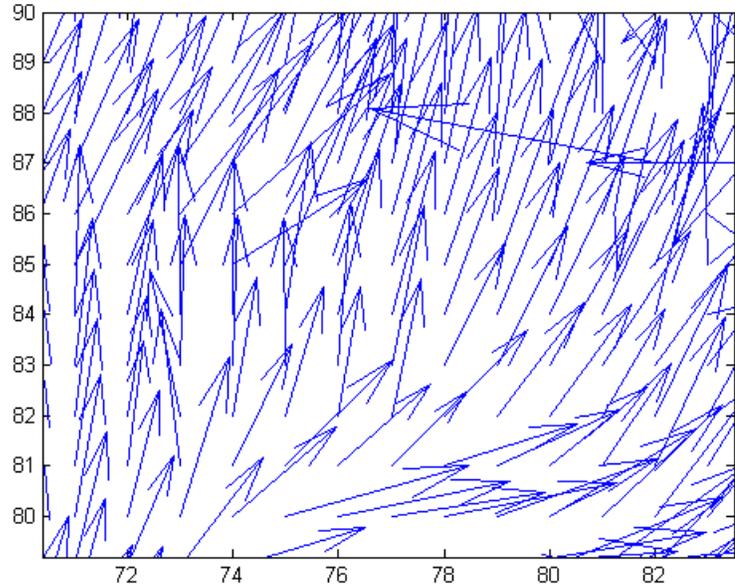
where α denotes the regularization parameter. This controls the amount of smoothing applied to the solution. The Horn-Schunck method is utilized in this thesis to estimate optical flow for synthetic and real video sequences which have global translational motion. In these cases, since the motion is the same across the entire image, a large value for the regularization parameter must be used to ensure a smooth estimate of the motion. Figure 3.4 shows two examples of optical flow estimation between the same two images. Figure

3.4a corresponds to an optical flow estimation with a large regularization parameter and Figure 3.4b utilizes the same images but with a small α . These images are globally offset in x and y . Consequently, the optical flow vector for every pixel should be the same. It is apparent that the small regularization parameter yields significantly more errors than the estimate with a large regularization weight.

When computing optical flow, it is important that the images do not contain artifacts such as aliasing. In these cases, the optical flow algorithm will try to match these artifacts between images and produce undesirable errors in the estimated vector fields.



(a) Optical flow vectors for large regularization parameter.



(b) Optical flow vectors for small regularization parameter.

Figure 3-4: Optical flow vectors computed for the same pair of images, but with varying regularization parameter. The images are globally translated relative to one another. Therefore, the optical flow vectors should be the same at every pixel, as in (a)

Chapter 4

Problem Formulation and Solution

In this chapter we will discuss the focus of this thesis, an improved superresolution method. Ideas from several existing methods have been combined to create an anisotropic, spline-based superresolution image reconstruction.

First, it is important to define key notation. Throughout this derivation C represents spline coefficients of the high resolution image, with $C[\mathbf{n}]$ being a coefficient at $\mathbf{n} = [n_x, n_y]^T \in \Lambda$. The matrices D_{xx} , D_{yy} and D_{xy} are defined as convolution matrices to perform "spline-based derivatives" where xx refers to the second-order derivative in the x direction, yy refers to the second-order derivative in the y direction and xy refers to the derivative in the x and then y direction. I_{xx} , I_{yy} and I_{xy} denote the second-order derivatives of the spline function, in other words $I_{xx} = D_{xx}C$.

4.1 Image Model

The image model used for the superresolution reconstruction is based on the work of Vazquez (Vazquez, 2002; Vazquez et al., 2003). Vazquez utilizes 2-D splines to model the underlying high-resolution scene. One advantage of using splines is that they are a simple way of constructing a continuous function. Given discrete samples of a signal, the coefficients of a spline can easily be solved for by using well-known filtering algorithms (see Section 3.1). Similarly, the continuous spline function can be computed given the spline coefficients. The low-resolution images are modeled as discrete samples of a continuous spline function. Therefore, to obtain a superresolution reconstruction, the coefficients of this underlying function must be computed.

Let y represent the low resolution images stacked in lexicographic order. If Φ relates

the spline coefficients to the low resolution image, and C represents the spline coefficients, then,

$$y = \Phi C \quad (4.1)$$

Matrix Φ is a very large and sparse matrix of size $RN_1N_2 \times M_1M_2$, where R is the number of low resolution frames, $[N_1, N_2]$ are the dimensions of the low resolution images, and $[M_1, M_2]$ are the dimensions of the high resolution spline grid.

For simplification, we assume perfect impulse sampling. Consequently, no point spread function (PSF) is included in the Φ matrix. We chose to exclude the PSF in order to strictly evaluate the effectiveness of the anisotropic spline-based reconstruction. Once the PSF is known, or approximated, it can easily be incorporated into the model by convolving the PSF with the discrete spline model.

To compute the entries of the Φ matrix, the 2-D spline function must be sampled at each of the low-resolution points. This requires the motion of each of the frames to be computed. For each pixel, \mathbf{n} , in the low resolution frame the optical flow vector, $[v_x[\mathbf{n}], v_y[\mathbf{n}]]$, is computed. This vector maps a pixel in the low-resolution image to a location on the high-resolution grid. If V_x and V_y represent matrices of the vector components $v_x[\mathbf{n}]$ and $v_y[\mathbf{n}]$, the location of the low-resolution images relative to the reference frame (see Figure 1.2) can be obtained. Thus, the following relationship is derived,

$$[X_{LR}, Y_{LR}] = [X_r + V_x, Y_r + V_y] \quad (4.2)$$

where X_r and Y_r are the pixel positions of the reference frame repeated R times stacked in lexicographic order. X_{LR} and Y_{LR} are matrices of the LR pixel locations in the x and y direction relative to the regular HR grid. One of the benefits of using splines is that once we obtain a continuous estimate of the underlying HR image, we can sample the spline at any resolution to obtain the HR image. Throughout this work, we will use the same HR sampling lattice as the spline coefficient lattice. Therefore, the location of the spline coefficients is also the location of the HR pixels.

The density of LR samples varies depending on the type of motion and the up-sampling factor of the HR grid. For simple global translational motion, the points will be irregular, however there will be a repeating sample pattern. Using these points, the distance of each irregular sample to the high resolution spline coefficient grid points is computed. This distance determines the contribution of a each 2-D cubic B-spline to each of the irregular points. The 2-D spline is zero outside of the region -2 to +2, therefore, only distances less than 2 will contribute to a particular point. Using the position of the low resolution points relative to the high resolution regular grid the Φ matrix is computed as follows,

$$\Phi[i, j] = \phi(X_{HR}(i) - X(j), Y_{HR}(i) - Y(j)) \text{ for } i = (1, 2, \dots, RN_1N_2), j = (1, 2, \dots, M_1M_2) \quad (4.3)$$

where X_{HR} and Y_{HR} represents the HR lattice, and ϕ represents the 2-D cubic B-spline function (see equation (3.5)). Fortunately, cubic B-splines are compactly supported, making Φ a sparse matrix.

4.2 Cost Function

4.2.1 Error Term

The cost function used for this work is similar to the MAP formulation described in Section 2.3 but with a different prior. The first term, the log-likelihood term, quantifies the deviation of the solution from the sampled data

$$E_L(C) = \|y - \Phi C\|^2 \quad (4.4)$$

where y represents the low resolution frames stacked in lexicographic order, C represents the estimate of the spline coefficients in lexicographic order and Φ is the matrix which relates the spline coefficients to the irregular samples in y . Since there is nothing in E_L to control splines in areas without samples, the minimization would likely lead to intensity values which extend beyond the dynamic range of the image in these regions. As a result, we include a regularization term to control the smoothness of the solution. There are many

possible choices for the regularization term. This thesis focuses on the implementation of a spline thin-plate model, similar to Vazquez's work (Vazquez, 2002; Vazquez et al., 2003).

4.2.2 Regularization Term

The regularization term increases the control over regions with little or no data. The regularization term proposed in this work is similar to the continuous formulation Vazquez proposed (Vazquez, 2002; Vazquez et al., 2003). However, to simplify the inclusion of *a priori* information (gradients), the spline-based regularization term will be discretized. Vazquez utilizes a spline thin plate model to control the curvature of the solution through the second order derivative of the spline. Since this thesis uses a slightly different approach than Vazquez has proposed, it is worthwhile to compare the convolution kernels of the two formulations. This is done to verify that there is no loss in fidelity. First, the comparison will begin with the derivation of the continuous formulation.

Vazquez formulates his regularization term continuously. The continuous thin-plate regularization function, $E_R(C)$, is defined as follows:

$$E_R(C) = \alpha^2 \iint \left(\left(\frac{\partial^2 b}{\partial x^2} \right)^2 + 2 \left(\frac{\partial^2 b}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 b}{\partial y^2} \right)^2 \right) dx dy \quad (4.5)$$

where $b(x, y) = \sum_{\mathbf{n} \in \Lambda} C[\mathbf{n}] \phi^3(x - n_x, y - n_y)$ and Λ represents the regular high-resolution grid. To control the smoothness of the image, the regularization parameter α is utilized. If α is too large, the solution will be over-smoothed. If α is too small, the solution will be under-smoothed.

Plugging in the equation for the 2-D spline, the first term of the integral equation can be rewritten as,

$$\left(\frac{\partial^2 b}{\partial x^2} \right)^2 = \left(\sum_{\mathbf{n} \in \Lambda} C[n_x, n_y] \phi_{xx}^3(x - n_x, y - n_y) \right)^2 \quad (4.6)$$

$$= \sum_{\mathbf{n} \in \Lambda} \sum_{\mathbf{m} \in \Lambda} C[n_x, n_y] C[m_x, m_y] \phi_{xx}^3(x - n_x, y - n_y) \phi_{xx}^3(x - m_x, y - m_y) \quad (4.7)$$

Expanding the regularization function (4.5), the following equation is obtained:

$$\begin{aligned}
E_R(C) &= \alpha^2 \sum_{\mathbf{n} \in \Lambda} C[x - n_x, x - n_y] \sum_{\mathbf{m} \in \Lambda} C[x - m_x, x - m_y] \\
&\quad \iint \phi_{xx}^3(x - n_x, y - n_y) \phi_{xx}^3(x - m_x, y - m_y) dx dy \\
&\quad + 2 \iint \phi_{xy}^3(x - n_x, y - n_y) \phi_{xy}^3(x - m_x, y - m_y) dx dy \\
&\quad + \iint \phi_{yy}^3(x - n_x, y - n_y) \phi_{yy}^3(x - m_x, y - m_y) dx dy \tag{4.8}
\end{aligned}$$

Substituting the 1-D splines results in the following equation,

$$\begin{aligned}
E_R(C) &= \alpha^2 \sum_{\mathbf{n} \in \Lambda} C[x - n_x, x - n_y] \sum_{\mathbf{l} \in \Lambda} C[x - m_x, y - m_y] \\
&\quad \iint \beta_{xx}^3(x - n_x) \beta^3(y - n_y) \beta_{xx}^3(x - m_x) \beta^3(y - m_y) dx dy \\
&\quad + 2 \iint \beta_x^3(x - n_x) \beta_y^3(y - n_y) \beta_x^3(x - m_x) \beta_y^3(y - m_y) dx dy \\
&\quad + \iint \beta^3(x - n_x) \beta_{yy}^3(y - n_y) \beta^3(x - m_x) \beta_{yy}^3(y - m_y) dx dy \tag{4.9}
\end{aligned}$$

Once \mathbf{n} and \mathbf{m} are chosen, the integrals can be evaluated. For our implementation, \mathbf{n} and \mathbf{m} are chosen to be on a regularly spaced grid with integer spacing. All of the parameters within the integrals of equation (4.8) are known and therefore the equation can be solved for analytically.

One of the nice properties of B-splines is that their derivatives are easily computed using the following formula (Unser, 1999):

$$\frac{d\beta^n(x)}{dx} = \beta^{n-1}\left(x + \frac{1}{2}\right) - \beta^{n-1}\left(x - \frac{1}{2}\right) \tag{4.10}$$

The first and second order derivatives of the cubic B-spline, using equation (4.10), are

(Vazquez et al., 2003):

$$\beta_x^3(x) = \begin{cases} 0 & x < -2 \\ \frac{(x+2)^2}{2} & -2 \leq x < -1 \\ -\frac{(3x^2+4x)}{2} & -1 \leq x < 0 \\ \frac{(3x^2-4x)}{2} & 0 \leq x < 1 \\ -\frac{(x-2)^2}{2} & 1 \leq x < 2 \\ 0 & 2 \leq x \end{cases} \quad (4.11)$$

$$\beta_{xx}^3(x) = \begin{cases} 3|x| - 2, & 0 \leq |x| < 1 \\ 2 - |x|, & 1 \leq |x| < 2 \\ 0, & 2 \leq |x| \end{cases} \quad (4.12)$$

Since B-spline functions are compactly supported from -2 to 2, the thin-plate spline integrals only need to be evaluated from -2 and 2. The first-order and second-order derivatives of the cubic B-spline are zero at points greater than or equal to 2. The coefficients of the 2-D splines which are contributing to a particular integral evaluation only extend from -3 to 3, therefore, $[n_x, n_y]$ and $[m_x, m_y]$ are evaluated for integer values between -3 and 3.

The energy term (4.9) is a discrete summation of the continuous integrals and can be rewritten in a matrix format. Combining the error term (4.4) and the regularization term (4.9) the following cost function is formulated,

$$J(C) = \|y - \Phi C\|^2 + \alpha^2 C^T P C \quad (4.13)$$

where P represents the convolution matrix of p , the spline thin-plate convolution kernel. The 7x7 matrix p , is computed by evaluating the integral equations at each of the grid of

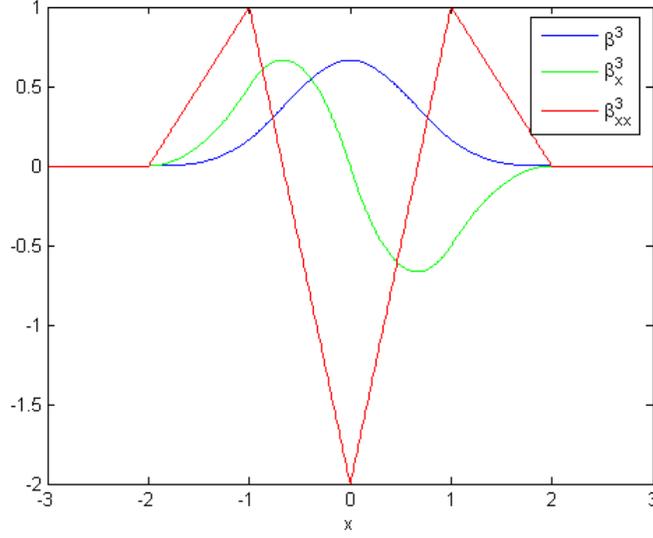


Figure 4-1: Cubic B-spline first-order and second-order derivatives.

points between -3 and 3 in x and y .

$$p = \begin{bmatrix} \frac{19}{92671} & \frac{23}{3150} & \frac{83}{2016} & \frac{131}{1890} & \frac{83}{2016} & \frac{23}{3150} & \frac{19}{92671} \\ \frac{23}{3150} & \frac{2}{25} & \frac{1}{70} & \frac{-64}{315} & \frac{1}{70} & \frac{2}{25} & \frac{23}{3150} \\ \frac{83}{2016} & \frac{1}{70} & \frac{-759}{1120} & \frac{-23}{90} & \frac{-759}{1120} & \frac{1}{70} & \frac{83}{2016} \\ \frac{131}{1890} & \frac{-64}{315} & \frac{-23}{90} & \frac{3256}{945} & \frac{-23}{90} & \frac{-64}{315} & \frac{131}{1890} \\ \frac{83}{2016} & \frac{1}{70} & \frac{-759}{1120} & \frac{-23}{90} & \frac{-759}{1120} & \frac{1}{70} & \frac{83}{2016} \\ \frac{23}{3150} & \frac{2}{25} & \frac{1}{70} & \frac{-64}{315} & \frac{1}{70} & \frac{2}{25} & \frac{23}{3150} \\ \frac{19}{92671} & \frac{23}{3150} & \frac{83}{2016} & \frac{131}{1890} & \frac{83}{2016} & \frac{23}{3150} & \frac{19}{92671} \end{bmatrix} \quad (4.14)$$

To compare Vazquez's continuous formulation to the discrete formulation proposed in this thesis, each term of the integral equation will be examined separately. Simplifying to a discrete formulation allows for straightforward inclusion of anisotropic regularization. First, the integrals are approximated with summations,

$$\iint \beta_{xx}^3(x - n_x) \beta^3(y - n_y) \beta_{xx}^3(x - m_x) \beta^3(y - m_y) dx dy \approx$$

$$\sum_{x=-2}^2 \sum_{y=-2}^2 \beta_{xx}^3[x - n_x] \beta^3[y - n_y] \beta_{xx}^3[x - m_x] \beta^3[y - m_y] \quad (4.15)$$

This can be rewritten using matrices and convolution,

$$= \phi'_{xx} * \phi'_{xx} \quad (4.16)$$

where ϕ'_{xx} denotes ϕ_{xx} , the second-order derivative of the 2-D B-spline, evaluated at discrete locations. The 2-D B-spline is a tensor product of the 1-D B-splines, therefore, ϕ'_{xx} can be written as the tensor product of β' and β'_{xx} , the vectors of the B-spline function and the second-order derivative of the B-spline function evaluated at integer points.

$$\phi'_{xx} = \beta'_{xx} \beta'^T \quad (4.17)$$

Each integral of the regularization term can be approximated in this manner,

$$\iint \beta_{xx}^3(x - n_x) \beta^3(y - n_y) \beta_{xx}^3(x - m_x) \beta^3(y - m_y) dx dy \approx \phi'_{xx} * \phi'_{xx} \quad (4.18)$$

$$2 \iint \beta_x^3(x - n_x) \beta_y^3(y - n_y) \beta_x^3(x - m_x) \beta_y^3(y - m_y) dx dy \approx 2 \phi'_{xy} * \phi'_{xy} \quad (4.19)$$

$$\iint \beta^3(x - n_x) \beta_{yy}^3(y - n_y) \beta^3(x - m_x) \beta_{yy}^3(y - m_y) dx dy \approx \phi'_{yy} * \phi'_{yy} \quad (4.20)$$

where $\phi'_{xy} = \beta'_x \beta'^T_y$. Figure 4.2 shows the spline derivative kernels for the two formulations. The first column of plots represents the discrete spline derivative kernel, the second corresponds to samples of the continuous spline derivative kernel, and the third displays the difference between the two. These plots demonstrate only slight differences between the continuous and discrete formulations. The reconstruction of the continuous and discrete convolution kernels was compared for several data sets. The results confirm that there is no loss in fidelity with the discrete approximation. There was no visible difference between the reconstructions for the two implementations and only slight differences in MSE. As a result, the approximation of the spline thin plate model closely replicates the continuous formulation proposed by Vazquez.

Using the above definitions, the regularization term can be rewritten as,

$$E_R(C) = \alpha^2(\|D_{xx}C\|^2 + 2\|D_{xy}C\|^2 + \|D_{yy}C\|^2) \quad (4.21)$$

where D_{xx} , D_{xy} and D_{yy} represent the convolution matrices of the spline derivative kernels ϕ'_{xx} , ϕ'_{xy} and ϕ'_{yy} , respectively.

4.2.3 Anisotropic Diffusion

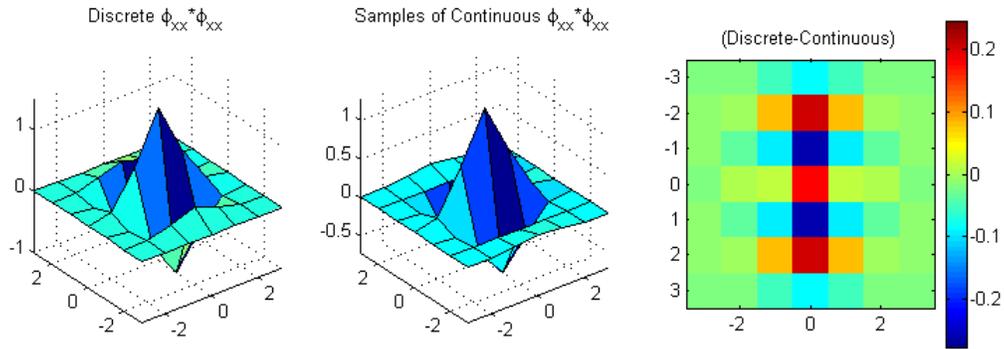
One of the drawbacks of including a regularization term is the possibility of over-smoothing the solution. This risk can be mitigated by adjusting the regularization parameter. Nonetheless, there are likely to be regions which are over-smoothed and other regions which are under-smoothed. Ideally, the image recovered would be smoothed appropriate to the content of the underlying scene. The formulation, up to this point, applies isotropic smoothing (i.e., is independent of scene content). To improve the impact of the regularization term, anisotropic diffusion will be applied.

The objective of anisotropic diffusion is to preserve image intensity boundaries. In other words, the smoothing should be performed along edges and in regions with constant intensity, but not across intensity boundaries. By incorporating anisotropic diffusion in our regularization term, the amount of smoothing can be dynamically controlled.

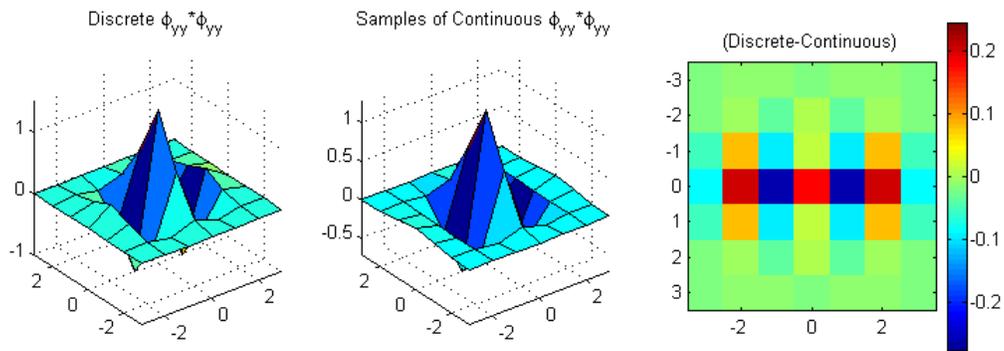
Anisotropic diffusion has been utilized for a number of image processing applications. Perona and Malik apply anisotropic diffusion for image scaling and creating coarser images (Perona and Malik, 1990). Their approach is to apply a separate weighting of the regularization term in the x and y directions. They apply this weighting through monotonically decreasing functions $g(|I_x|)$ and $g(|I_y|)$, where I_x and I_y represent the gradient in the x and y directions. The prior term of the cost function is computed as follows:

$$E_R(C) = \sum_{\mathbf{n} \in \Lambda} \left(\begin{array}{cc} [I_{xx}[\mathbf{n}] & I_{yy}[\mathbf{n}]] \end{array} \begin{bmatrix} g(|I_x[\mathbf{n}]|) & 0 \\ 0 & g(|I_y[\mathbf{n}]|) \end{bmatrix} \begin{bmatrix} I_{xx}[\mathbf{n}] \\ I_{yy}[\mathbf{n}] \end{bmatrix} \right) \quad (4.22)$$

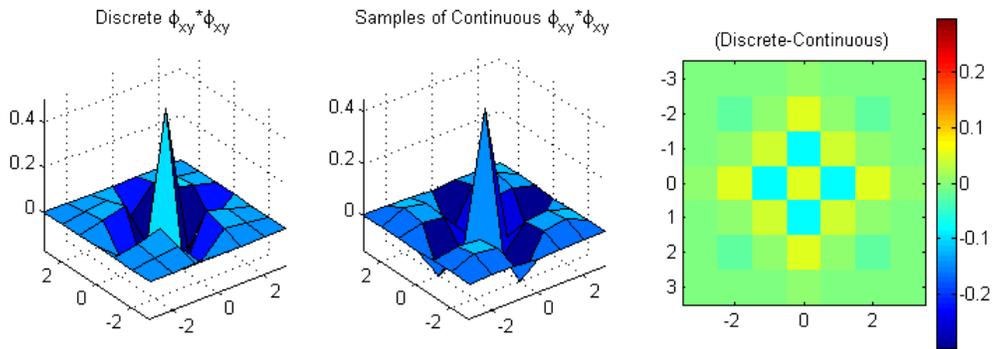
where, $g(|I|) = e^{-\left(\frac{|I|}{K}\right)^2}$, K is a constant, and I_{xx} and I_{yy} represent the second order



(a) Discrete and Continuous spline derivative kernel comparison for the second order derivative in the x direction.



(b) Discrete and Continuous spline derivative kernel comparison for the second order derivative in the y direction.



(c) Discrete and Continuous spline derivative kernel comparison for the second order derivative in the x then y direction.

Figure 4-2: Discrete and continuous spline derivative convolution kernel comparison.

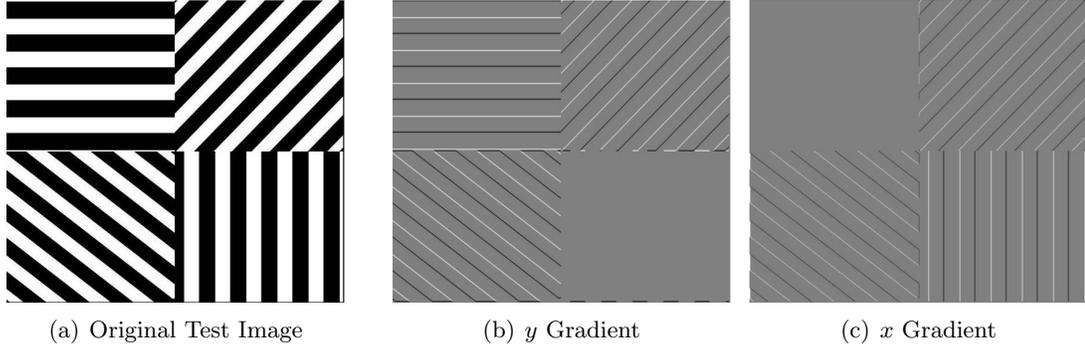


Figure 4-3: Binary *Stripes* image utilized for comparison of anisotropic diffusion methods.

derivatives in the x and y directions. As $K \rightarrow \infty$, the matrix becomes the identity matrix and therefore applies isotropic weighting. For a very large horizontal gradient and small vertical gradient, $g(|I_x|)$ is small and $g(|I_y|)$ is large. Therefore, smoothing is disabled in the horizontal direction and smoothing is performed in the vertical direction. Similarly, strong vertical gradients and small horizontal gradients will result in smoothing only along the horizontal intensity boundaries. Although this method performs well for horizontal and vertical gradients, off-axis edges are not smoothed as well.

Consider the binary image shown in Figure 4.3a, with horizontal, vertical and diagonal stripes with an intensity value of 1 and a background values of 0. The gradient for the horizontal and vertical stripes will result in anisotropic weighting such that smoothing will be performed along the edges and not across. For example, let $K = 0.1$ and let H represent the anisotropic weighting matrix.

$$H[\mathbf{n}] = \begin{bmatrix} g(|I_x[\mathbf{n}]|) & 0 \\ 0 & g(|I_y[\mathbf{n}]|) \end{bmatrix} \quad (4.23)$$

The values of H for horizontal, vertical and diagonal edges are as follows:

$$H_{Horizontal} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad H_{Vertical} = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \quad H_{Diagonal} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \quad (4.24)$$

Therefore, for the horizontal and vertical edges, smoothing will be performed in the appropriate directions. However, the diagonal edges will not be smoothed in either direction. The gradient of the diagonal stripes has an absolute value of 1 (gradient of a binary image) in both the x and y directions. Therefore, $g(|I_x|)$ and $g(|I_y|)$ is small in both directions, resulting in no smoothing. This will result in rough edges along the diagonals (see Figure 4.4a).

Another approach to anisotropic diffusion is presented by Mansouri *et al* (Mansouri et al., 1998). The motivation of Mansouri's work is to reconstruct image disparities through image diffusion. Although the problem is a bit different from the superresolution reconstruction problem, the anisotropic diffusion equation which Mansouri *et al* propose is appropriate for our problem. The anisotropic weighted regularization term they propose is:

$$E_R(C) = \sum_{\mathbf{n} \in \Lambda} \left(\begin{matrix} I_{xx}[\mathbf{n}] & I_{yy}[\mathbf{n}] \end{matrix} \begin{bmatrix} \frac{\mu^2 + I_y[\mathbf{n}]^2}{\mu^2 + \|\vec{\nabla} I[\mathbf{n}]\|^2} & \frac{-I_x[\mathbf{n}]I_y[\mathbf{n}]}{\mu^2 + \|\vec{\nabla} I[\mathbf{n}]\|^2} \\ \frac{-I_x[\mathbf{n}]I_y[\mathbf{n}]}{\mu^2 + \|\vec{\nabla} I[\mathbf{n}]\|^2} & \frac{\mu^2 + I_x[\mathbf{n}]^2}{\mu^2 + \|\vec{\nabla} I[\mathbf{n}]\|^2} \end{bmatrix} \begin{bmatrix} I_{xx}[\mathbf{n}] \\ I_{yy}[\mathbf{n}] \end{bmatrix} \right) \quad (4.25)$$

where μ is a constant and $\vec{\nabla} I[\mathbf{n}] = [I_x[\mathbf{n}], I_y[\mathbf{n}]]^T$. For simplification, let $H[\mathbf{n}]$ represent the anisotropic matrix for pixel \mathbf{n} , and H_{xx} , H_{xy} and H_{yy} represent the entries of the matrix evaluated at each pixel location in lexicographic order.

$$H[\mathbf{n}] = \begin{bmatrix} \frac{\mu^2 + I_y[\mathbf{n}]^2}{\mu^2 + \|\vec{\nabla} I[\mathbf{n}]\|^2} & \frac{-I_x[\mathbf{n}]I_y[\mathbf{n}]}{\mu^2 + \|\vec{\nabla} I[\mathbf{n}]\|^2} \\ \frac{-I_x[\mathbf{n}]I_y[\mathbf{n}]}{\mu^2 + \|\vec{\nabla} I[\mathbf{n}]\|^2} & \frac{\mu^2 + I_x[\mathbf{n}]^2}{\mu^2 + \|\vec{\nabla} I[\mathbf{n}]\|^2} \end{bmatrix} \quad (4.26)$$

$$H_{xx} = \frac{\mu^2 + I_y^2}{\mu^2 + \|\vec{\nabla} I\|^2} \quad (4.27)$$

$$H_{xy} = \frac{-I_x I_y}{\mu^2 + \|\vec{\nabla} I\|^2} \quad (4.28)$$

$$H_{yy} = \frac{\mu^2 + I_x^2}{\mu^2 + \|\vec{\nabla} I\|^2} \quad (4.29)$$

As $\mu \rightarrow \infty$, this matrix becomes an identity matrix yielding an isotropic equation. Similarly, if the gradient is small relative to μ , isotropic diffusion is performed. Once

again, this matrix is evaluated against the simple binary *Stripes* image shown in Figure 4.3. For example let $\mu = 0$, the smoothing parameters for horizontal, vertical and diagonal edges are weighted as follows,

$$H_{Horizontal} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad H_{Vertical} = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \quad H_{Diagonal} = \begin{bmatrix} \frac{1}{2} & \pm\frac{1}{2} \\ \pm\frac{1}{2} & \frac{1}{2} \end{bmatrix} \quad (4.30)$$

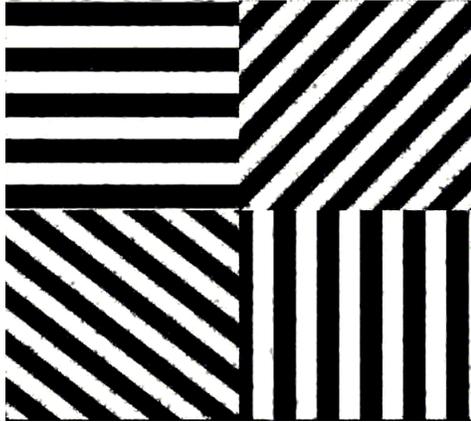
Therefore, in the horizontal case smoothing will only be performed in the horizontal direction and in the vertical direction for the vertical stripes. The diagonal case will have some smoothing in both directions. Figure 4.4b demonstrates the effectiveness of this anisotropic matrix for the simple striped case. It is evident that the diagonal edges in the Perona and Malik formulation are slightly rougher than the edges in the Mansouri *et al* implementation. This equation is reformatted to the matrix format used in the prior term, resulting in the following equations,

$$E_R(C) = \sum_{\mathbf{n} \in \Lambda} (H_{xx} I_{xx}^2[\mathbf{n}] + 2H_{xy} I_{xx}[\mathbf{n}] I_{yy}[\mathbf{n}] + H_{yy} I_{yy}^2[\mathbf{n}]) \quad (4.31)$$

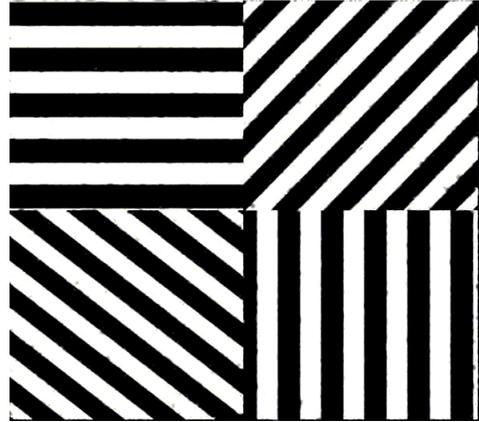
$$E_R(C) = \|D_{xx}C\|_{H_{xx}}^2 + 2\|D_{xy}C\|_{H_{xy}}^2 + \|D_{yy}C\|_{H_{yy}}^2 \quad (4.32)$$

Note that $D_{xx}D_{yy} = D_{xy}D_{xy}$. This results in the same spline thin-plate prior term as previously proposed, however, it now has anisotropic weighting.

Figure 4.5 compares anisotropic and isotropic diffusion for a simple 1-D case. The original signal, containing two square waves, was randomly sampled and reconstructed using the formulation proposed. Figure 4.5a shows the isotropic result and Figure 4.5b displays the anisotropic result. In this case, the gradient is constructed from the original signal, therefore the gradient is exact. The same regularization parameter is used for the two reconstructions. It is apparent that the anisotropic method significantly outperforms the isotropic method in this example. The isotropic reconstruction is overly smoothed and bridges the gap between the two square waves. On the other hand, the anisotropic method creates sharp edges with only small errors where there are few samples.



(a) Perona and Malik method for anisotropic weighting using monotonically decreasing function.



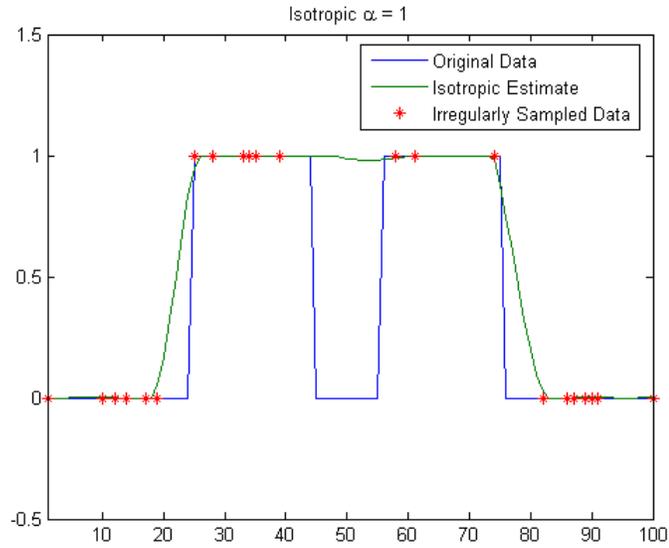
(b) Mansouri *et al* anisotropic weighting using matrix.

Figure 4-4: Comparison of anisotropic diffusion methods. Mansouri *et al* implementation (Mansouri et al., 1998) provides improved smoothing of diagonal edges.

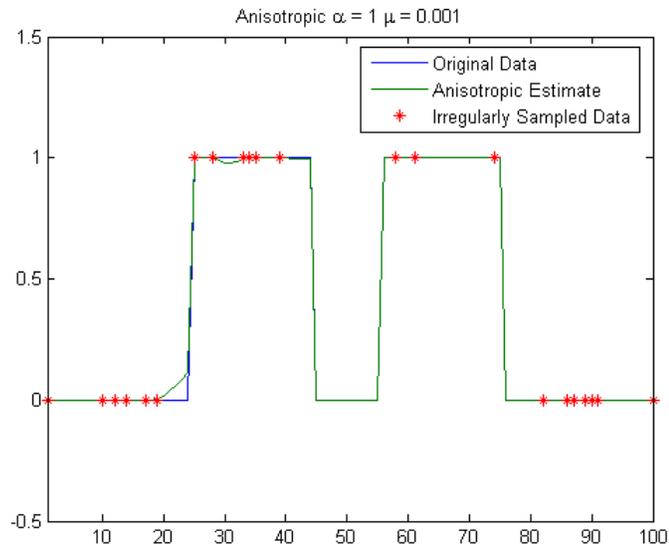
4.2.4 Gradient Estimation

In order to compute the anisotropic weights, we need to estimate the underlying HR image gradient. Any inaccuracies in the gradient estimation will result in undesirable anisotropic weighting and artifacts in the HR image. If it were possible to obtain perfect gradient information from the LR frames, then there would be no need for superresolution since all of the frequency content of the underlying scene would be present.

The anisotropic diffusion methods presented by Perona and Malik, and Mansouri *et al* estimate the edge information for every iteration of the minimization. This works well given their applications. Perona and Malik implement anisotropic diffusion to create high quality coarse images. Given a HR image, their objective, is to create a coarse image with distinct intensity boundaries. Therefore, they begin with 'perfect' gradient information of the high resolution image and with each iteration obtain new edge information based on the anisotropically diffused image. Mansouri *et al* apply anisotropic diffusion to disparity estimation. The basic assumption is that image intensity boundaries correspond to disparity boundaries and, therefore, smoothing should be disabled along these boundaries. Once



(a) 1-D example of isotropic diffusion for randomly sampled data.



(b) 1-D example of anisotropic diffusion for randomly sampled data.

Figure 4.5: 1-D example of isotropic and anisotropic diffusion demonstrates the edge-enhancing effects of the anisotropic implementation.

again, this application has the advantage of knowing the gradient information of each of the frames and therefore will be properly updated with each iteration. The application of anisotropic diffusion to superresolution, however, is a bit different.

Superresolution images are constructed from a set of LR images. Therefore, the gradient information of the HR image is unknown at the onset of the minimization. The only gradient information available is the gradient of the LR, degraded images. The HR image may be estimated using a basic interpolation scheme and then this information used to initialize the minimization. However, any inaccuracies in the interpolation scheme will result in artifacts in the reconstruction. Another problem with this implementation is that there will be holes in the data where there are few irregular samples. As a result, the gradient estimate for each iteration will have artifacts caused by these holes. Based on these realizations, the solution chosen is to estimate the gradient of the high resolution image prior to minimization. Every iteration of the minimization utilizes the same gradient information thus necessitating an accurate superresolution gradient estimation. Several approaches for the estimation of the gradient were experimented with. These methods include estimating the gradient from the estimate of a HR image using:

1. Nearest-Neighbor Interpolation
2. Linear Interpolation
3. Cubic Interpolation
4. Isotropic Diffusion

For these approaches the HR image is estimated, a low pass filter is applied to eliminate any high-frequency errors and then the gradient is computed. The gradient is approximated using finite-differences in the x and y direction.

Three different data sets were utilized for the gradient comparison. The construction of these data sets is detailed in Chapter 5. For each of the data sets, the gradient was approximated with the methods detailed above and then the mean square error (MSE) was

computed. Based on the experimental results, it was found that isotropic reconstruction, followed by low pass filtering and gradient estimation provided the best overall gradient estimation.

4.2.5 Minimization

Combining equations (4.4) and (4.32), we construct the following cost function:

$$J(C) = (\|y - \Phi C\|^2) + \alpha^2 \left[\|D_{xx}C\|_{H_{xx}}^2 + 2\|D_{xy}C\|_{H_{xy}}^2 + \|D_{yy}C\|_{H_{yy}}^2 \right] \quad (4.33)$$

The function is minimized by computing the derivative and setting equal to zero:

$$J(C) = (y - \Phi C)^T (y - \Phi C) + \alpha^2 \left[(D_{xx}C)^T H_{xx} (D_{xx}C) + 2(D_{xy}C)^T H_{xy} (D_{xy}C) + (D_{yy}C)^T H_{yy} (D_{yy}C) \right] \quad (4.34)$$

$$\frac{dJ}{dC} = 2(\Phi^T \Phi C) - 2(\Phi^T y) + 2\alpha^2 \left[D_{xx}^T H_{xx} D_{xx} C + 2D_{xy}^T H_{xy} D_{xy} C + D_{yy}^T H_{yy} D_{yy} C \right] = 0 \quad (4.35)$$

Rearranging the equation we obtain,

$$\Phi^T y = \Phi^T \Phi C + \alpha^2 \left[D_{xx}^T H_{xx} D_{xx} C + 2D_{xy}^T H_{xy} D_{xy} C + D_{yy}^T H_{yy} D_{yy} C \right] \quad (4.36)$$

To determine the solution to equation (4.36), for which $J(C)$ is minimized, the conjugate gradient algorithm is utilized.

Chapter 5

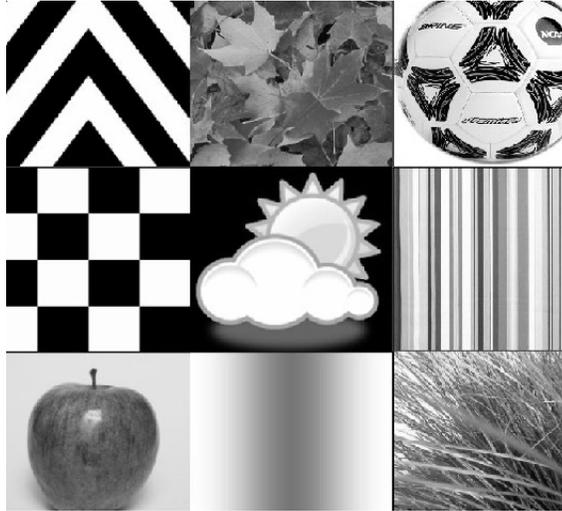
Experimental Results

In this chapter, we evaluate the effectiveness of the proposed spline-based, anisotropic superresolution algorithm on several data sets. First, we present results for ground-truth data that allow a numeric comparison in terms of reconstruction error. While one data set includes random sample positions, the other set includes three LR images undergoing translational motion. Then, we present results for real video sequences.

5.1 Synthetic Motion

We tested our anisotropic spline-based algorithm on several synthetically-generated data sets. These data sets allow us to control the presence of blur (PSF of the capture system) as well as motion parameters used to generate the LR images. Therefore, the SR algorithm could be tested under different conditions such as with perfect motion and with estimated motion, with and without PSF, with perfect gradient and with estimated gradient. In each case, the improvements were quantified using the mean squared error (MSE). Three HR images were used to construct three corresponding LR data sets (see Figure 5.1). The reconstruction for each of these three images and each type of low-resolution modeling was evaluated. To obtain samples for reconstruction, two different methods were used. The first method consists of randomly sampling the HR image and estimating the original image from these irregular samples. In the second method, a more realistic approach, globally-translated LR images were created from the HR image. Each of the modeling techniques and the corresponding results are detailed in Sections 5.1.1 and 5.1.2.

For each synthetic data set the MSE was computed for a variety of regularization parameters. Unfortunately, MSE is not always an accurate measure of image quality.



(a) High-resolution *Squares* image (480x528).



(b) High-resolution *Brick* image (672x896).



(c) High-resolution *Jeep* image (448x832).

Figure 5-1: HR images used to generate synthetic-motion data sets.

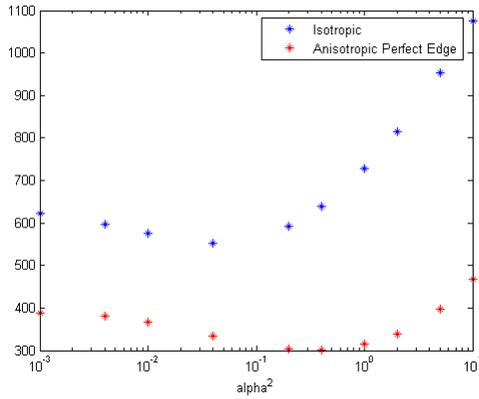
When comparing two images, one image may have a lower MSE than the other image, however the image with a higher MSE could be more visually pleasing. This is likely to occur when there are overshoots at a few pixels only. The cases in this thesis for which this holds true will be identified and the corresponding images will be provided. Given that there were many images generated to evaluate the effectiveness of the SR algorithm, it is impractical to include all of the images in this thesis. Therefore, through MSE and several representative examples the effectiveness of the SR algorithm will be demonstrated.

5.1.1 Randomly-Sampled Images

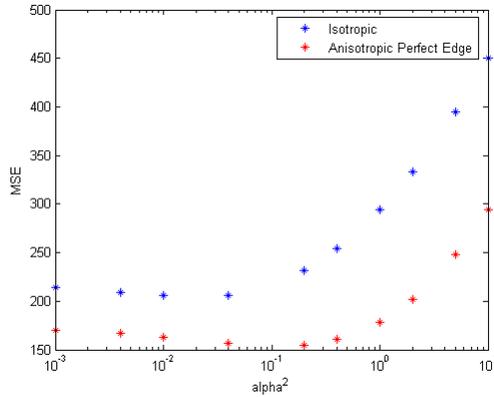
The first and simplest method for testing the SR implementation involved randomly sampling a high resolution-image. The original image is then reconstructed from the irregularly-sampled data. There is no motion estimation needed since only one set of sample points is used (in a sense, one image). The effectiveness of the anisotropic reconstruction is compared to isotropic reconstruction. Anisotropic diffusion was performed using gradient information of the original HR image. This is a very simple test case, however it allows for examination of the impact of anisotropic diffusion under ideal conditions (gradient of the original image). For each of the images tested, 30% of the pixels from original HR image

were used.

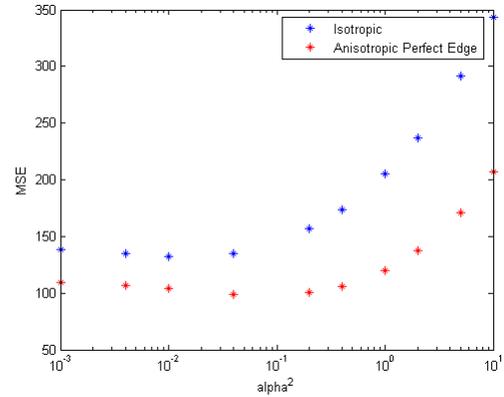
As the MSE plots in Figure 5.2 demonstrate, the proposed anisotropic algorithm based on perfectly-known edges (gradient) significantly outperforms the isotropic reconstruction. This is also clearly visible in reconstructed images (Figure 5.3), especially on the checker-board pattern.



(a) *Squares*



(b) *Brick*

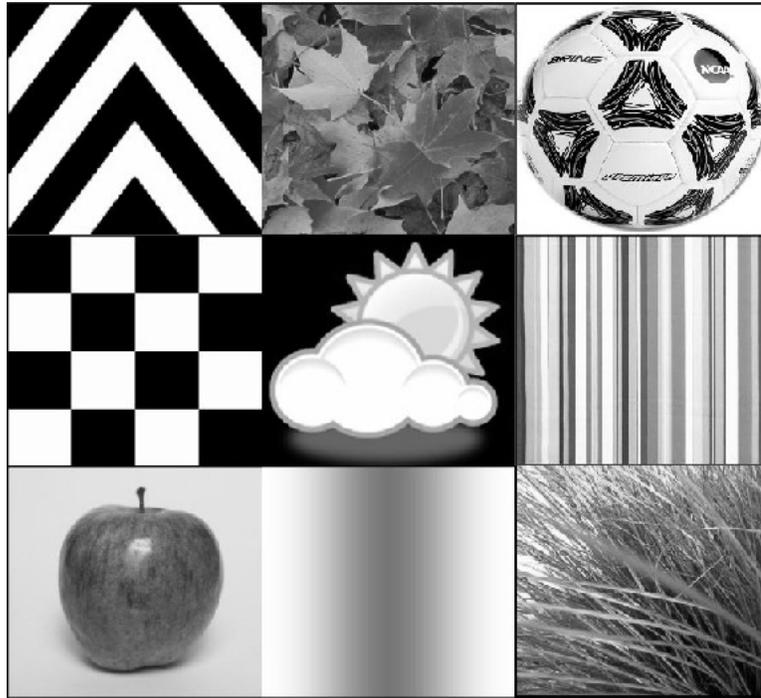


(c) *Jeep*

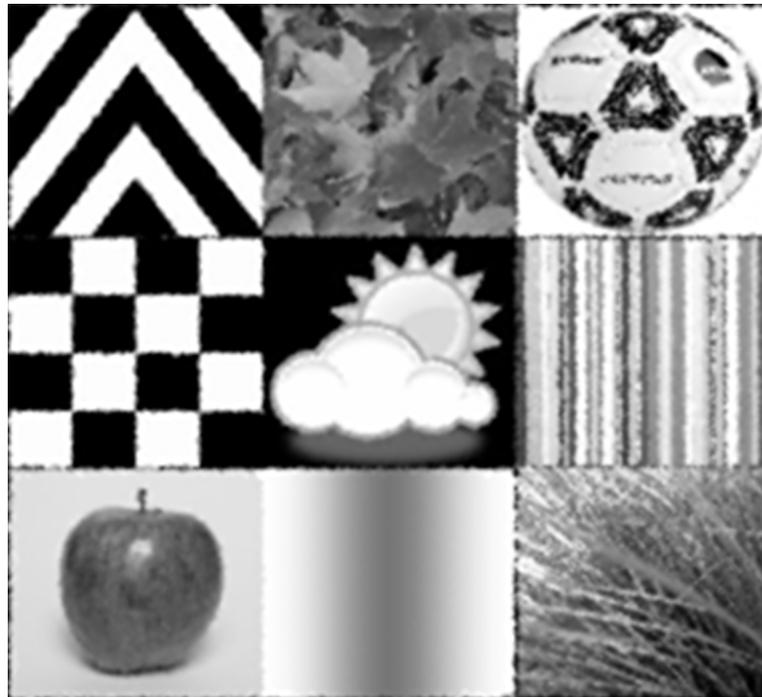
Figure 5.2: Mean square reconstruction error using randomly-sampled data set for both isotropic as well as anisotropic algorithm.

5.1.2 Global Translational Motion

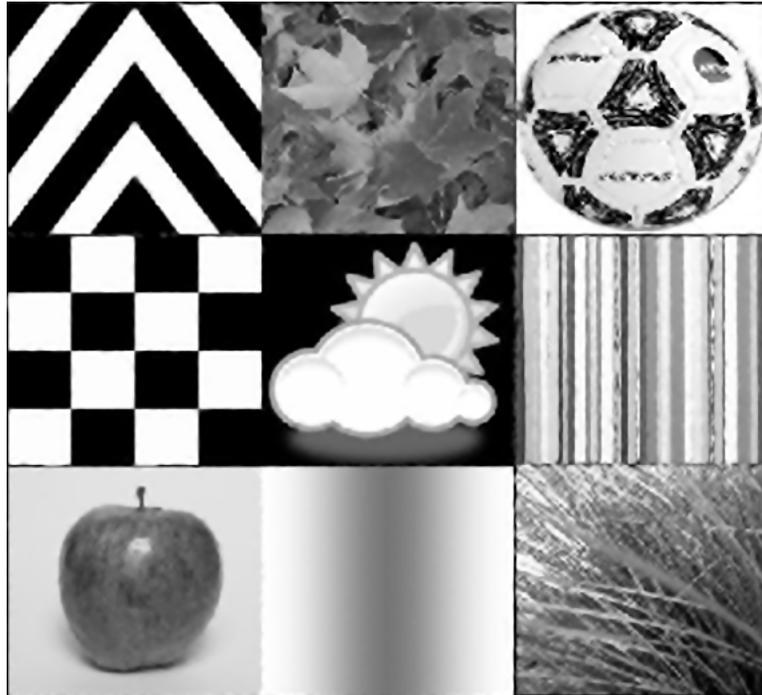
The second set of test images were created using global translational motion. To generate this test set two different approaches were used, one with pre-filtering and the other without



(a) *Squares* original image.



(b) *Squares* isotropic reconstruction with $\alpha = .4$ (MSE = 639).



(c) *Squares* anisotropic reconstruction using perfect gradient information with $\alpha = .4$, $\mu = 5$ (MSE = 301).

Figure 5-3: Results of reconstruction for randomly-sampled *Squares* image: (a) original image; (b) isotropic reconstruction; and (c) anisotropic reconstruction using perfect gradient information. Each image is presented with α resulting in lowest MSE. The edge enhancing effects of the anisotropic diffusion are clearly visible in (c) on the checkerboard pattern.

pre-filtering. Pre-filtering is typically used prior to downsampling to eliminate aliasing. The drawback to using a pre-filter to create the LR frames is that the blur introduced will appear in the HR reconstruction. To demonstrate the impact of blurring, each data set will be tested with and without pre-filtering.

In the case without pre-filtering, the LR images are aliased, however the reconstruction is not limited by the blur introduced by a PSF. In the second test, a Gaussian pre-filter is used to eliminate aliasing (see Figure 5.4 for the impact of pre-filtering). The variance of the pre-filter was chosen by visual inspection. Since this pre-filter (PSF) is not accounted for in the Φ matrix, the HR images constructed from the pre-filtered LR images are expected to be blurred. In both data sets, the HR images were shifted using global translational motion and then downsampled. In order to obtain LR images that are shifted by subpixel amounts, at least one shift, either x or y , cannot be a multiple of the sample period. The images were cropped after down-sampling to eliminate regions without data (caused by the shifting in x and y directions). The SR up-sampling factor for each image is the same as the down-sampling factor used in preparing each data set, thus allowing for computation of the MSE between the SR image and the original HR image.

The number of frames necessary for superresolution reconstruction can vary significantly depending on the up-sampling factor and the displacement between frames. For each test set, the number of LR frames for SR reconstruction is limited to three. The motivation for using only three frames for reconstruction is to show a slightly stressing case (using more frames naturally leads to improved reconstruction results) and highlight the effectiveness of the anisotropic superresolution algorithm.

For all 3 test images, *Squares*, *Brick* and *Jeep*, the HR image was shifted by $[2, 13]$, $[0, 0]$ and $[5, 11]$. In the *Squares* case, the HR image was downsampled by 3 in x and y , resulting in 3 LR images with subpixel shifts of $[0.67, 4.33]$, $[0, 0]$ and $[1.67, 3.67]$. Figure 5.5 shows the locations of the LR pixels relative to the HR grid. The original HR *Squares* image is 480×528 , therefore each LR image is 160×176 . Since 3 frames are utilized for reconstruction, the percentage of HR pixels used for reconstruction is 33%. The *Brick*

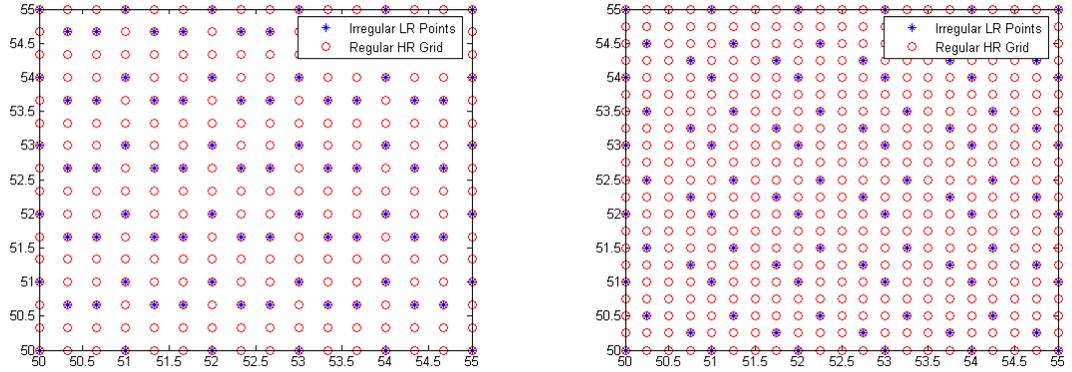


(a)



(b)

Figure 5.4: Impact of pre-filter on down-sampled image quality: LR image obtained (a) with pre-filtering; and (b) without pre-filtering. Gaussian pre-filter minimizes aliasing in the LR frames, but also limits the improvements of the SR algorithm.



(a) Location of LR pixels on HR grid for *Squares* image. For this case, 3 frames were used for the superresolution reconstruction and the images were up-sampled by 3 in both dimensions.

(b) Location of LR pixels on HR grid for *Brick* and *Jeep* images. For these images, 3 frames were used for the superresolution reconstruction and the images were up-sampled by 4 in both dimensions.

Figure 5-5: Location of LR pixels on HR grid.

image was downsampled by 4, resulting in 3 LR images with subpixel shifts of $[0.5, 3.25]$, $[0, 0]$, and $[1.25, 2.75]$. The original HR image is 672×896 and the LR images are 168×224 . In this case, the percentage of HR pixels used for reconstruction is 19%. Similarly, the 448×832 *Jeep* image was downsampled by 4, resulting in 112×208 LR images.

In order to evaluate the effectiveness of the SR reconstruction, the SR image was computed using either perfect or estimated motion parameters. Also, the pre-filter was either used or not. This resulted in the following four test conditions:

1. no pre-filter, perfect motion,
2. no pre-filter, estimated motion,
3. Gaussian pre-filter, perfect motion,
4. Gaussian pre-filter, estimated motion.

Motion estimation was performed in each case using the Horn and Schunck optical flow algorithm described in Section 3.3, implemented over a hierarchy of resolutions (multi-resolution implementation). For each of the above conditions, we evaluated isotropic diffusion, anisotropic diffusion using the gradient of the HR image and anisotropic diffusion

using the gradient estimated from LR images (Section 4.2.4). This comparison permits verification of the improvement of anisotropic diffusion over isotropic diffusion in the ideal case (perfect gradient) and for the proposed method (estimated gradient).

Tables 5.1 and 5.2 show the minimum MSE obtained using various reconstruction algorithms applied to *Squares*, *Brick* and *Jeep* images. Note that we also included results for a simple 2-D cubic interpolation of one LR image to gauge the improvement afforded by SR reconstruction.

Clearly, all SR reconstruction algorithms compared in Tables 5.1 and 5.2 outperform the simple bi-cubic interpolation, which was to be expected. It can be also concluded that anisotropic diffusion outperforms the isotropic diffusion when the motion is precisely known. As expected, the anisotropic diffusion with perfect gradient has the lowest MSE and is significantly lower than that of the isotropic reconstruction. The estimated gradient case does not perform as well as the perfect gradient case, however it still outperforms the isotropic reconstruction. The inaccuracies in the gradient estimation limit the edge enhancement potential of the anisotropic method.

Table 5.1: Minimum MSE for Each Test Set with Perfect Motion

Filtering	Image	Cubic Interpolation	Isotropic	Anisotropic Perfect Gradient	Anisotropic Estimated Gradient
Pre-Filtered	<i>Squares</i>	280.7	245.6	204.4	226.9
	<i>Brick</i>	183.5	172.9	148.6	160.4
	<i>Jeep</i>	107.0	97.4	77.0	85.8
No Pre-Filter	<i>Squares</i>	284.4	237.2	174.4	213.1
	<i>Brick</i>	176.0	175.7	136.2	159.5
	<i>Jeep</i>	116.1	111.8	80.6	97.3

Figures 5.6, 5.7 and 5.8 show complete MSE plots as a function of the regularization parameter α^2 . Clearly, a very small α^2 results in an under-regularized image and therefore a

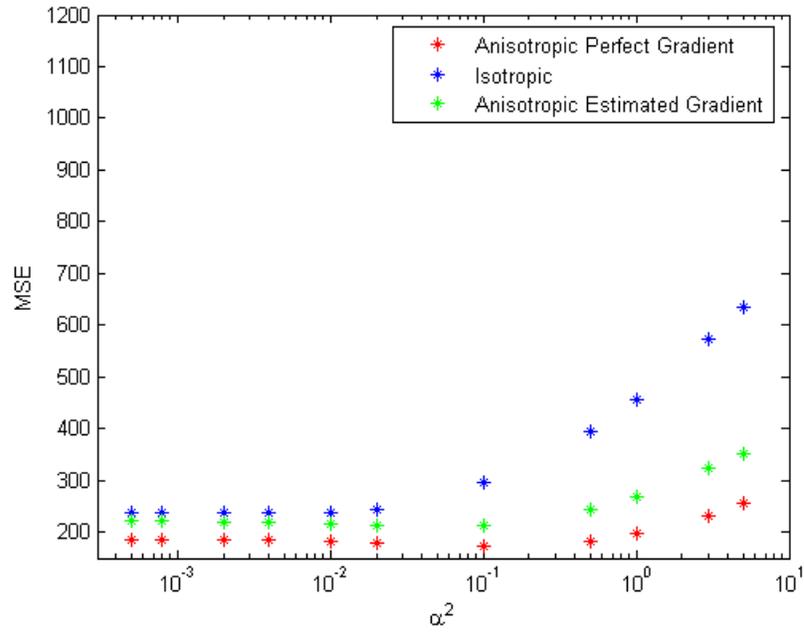
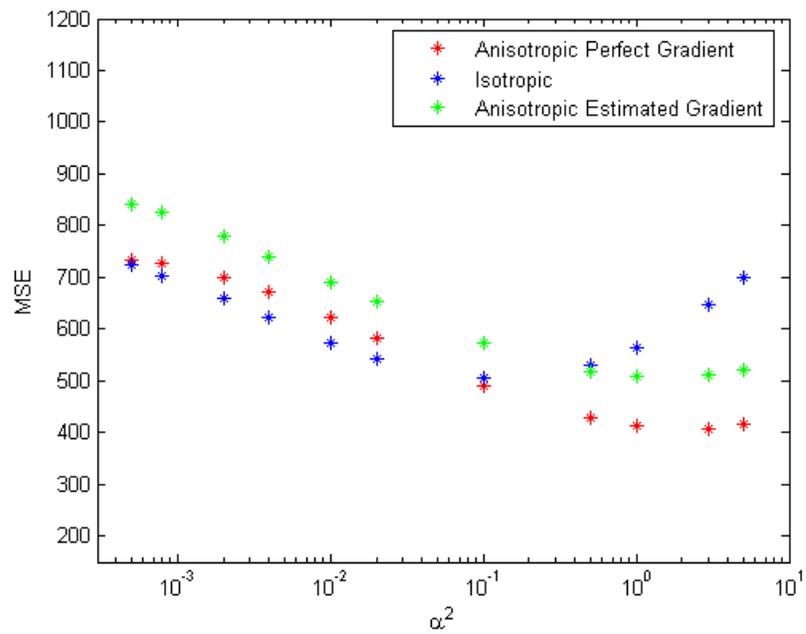
Table 5.2: Minimum MSE for Each Test Set with Estimated Motion

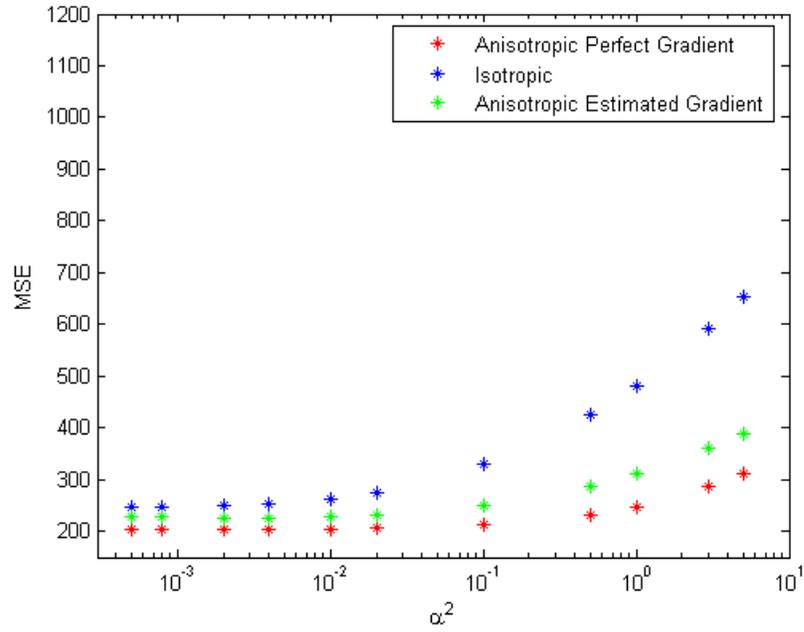
Filtering	Image	Cubic Interpolation	Isotropic	Anisotropic Perfect Gradient	Anisotropic Estimated Gradient
Pre-Filtered	<i>Squares</i>	323.1	307.4	261	304.7
	<i>Brick</i>	217.7	207.2	186.0	198.8
	<i>Jeep</i>	133.1	124.4	107.4	117.6
No Pre-Filter	<i>Squares</i>	565	505.2	407.1	508.8
	<i>Brick</i>	388.7	343.4	318	340.2
	<i>Jeep</i>	210.4	190.4	169.4	187.2

very high MSE. Similarly, a large α^2 results in an over-regularized reconstruction and a high MSE. Each of the MSE plots is zoomed-in to show the range of regularization parameters of interest. Again, note the degrading impact of imprecise motion (inaccurate estimation) on the performance of SR algorithms, especially in the case without pre-filtering.

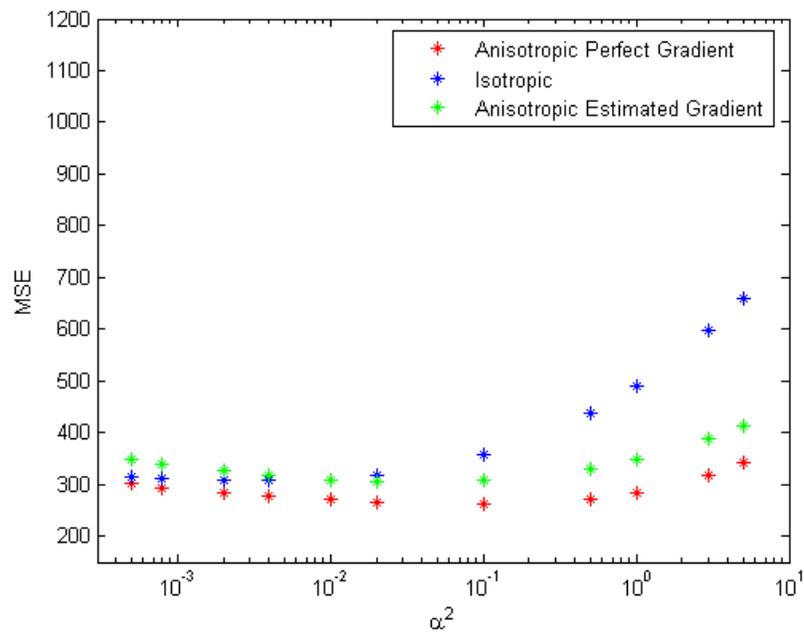
Finally, reconstructed images with the lowest MSE (from Tables 5.1 and 5.2) for all three images are shown in Figures 5.9–5.11. The improvements are subtle but clear, especially in the zoomed-in images (Figures 5.9f, 5.10f, and 5.11f). Based on the subjective comparison of reconstructed images, it can be also concluded that anisotropic diffusion outperforms the isotropic diffusion. The isotropic diffusion image, such as one shown in Figure 5.10b, has highly visible artifacts along the edges. The spacing of these artifacts corresponds to the spacing of the LR samples shown in Figure 5.5. The regularization parameter is too low in these cases to smooth out the regions void of LR samples. As the regularization parameter is increased, these effects are reduced, however the image then becomes severely blurred. Anisotropic diffusion allows the regularization parameter to increase to the point where the sampling artifacts are eliminated without causing an overly blurred image. The Gaussian pre-filtered images lead to similar results, however, the effects are not as prominent due to the blurring introduced by the pre-filter (see Figures 5.9f, 5.10f, and 5.11f).

The comparison of the MSE plots with optical flow estimation reveals the impact of aliasing in the LR images. In the case of no pre-filter, the MSE value increases significantly. Aliasing in the LR images is interpreted as motion by the optical flow algorithm. Therefore, pixels are mapped to incorrect locations causing severe artifacts in the SR images. On the other hand, the optical flow test for the pre-filtered images only slightly increases the MSE values.

(a) Results for *Squares* image: no pre-filter and perfect motion.(b) Results for *Squares* image: no pre-filter and optical flow motion estimation.

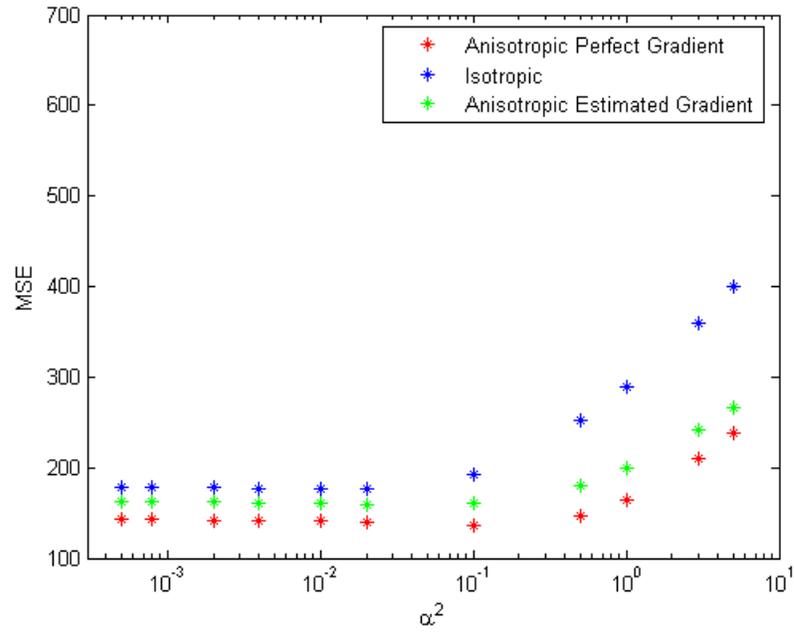
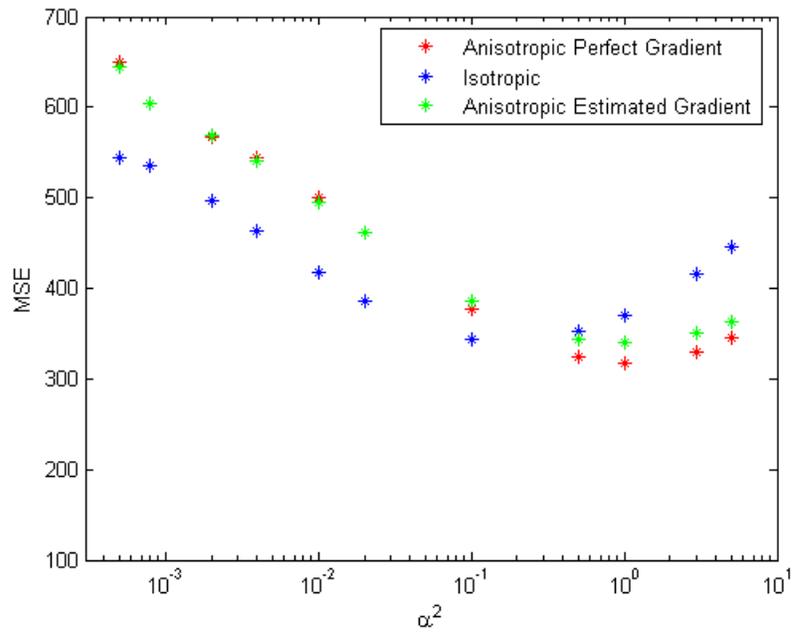


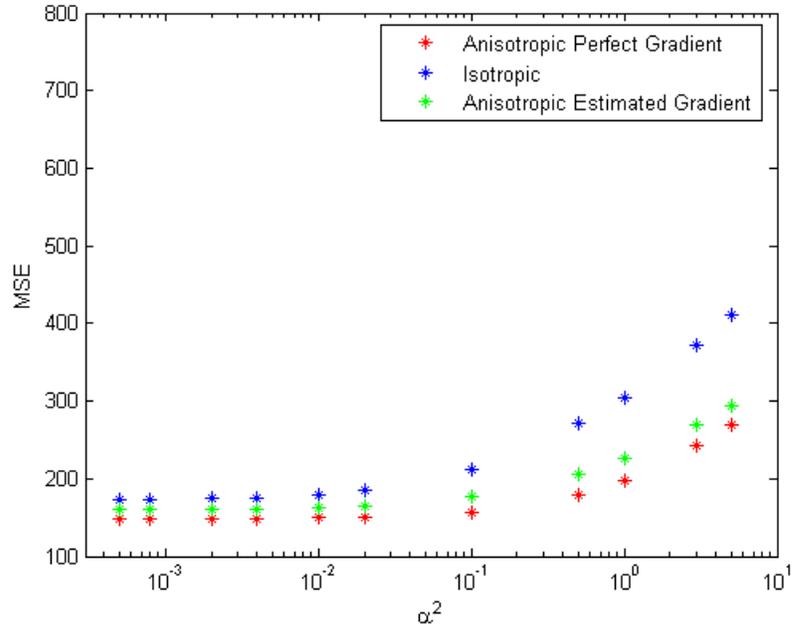
(c) Results for *Squares* image: Gaussian pre-filter and perfect motion.



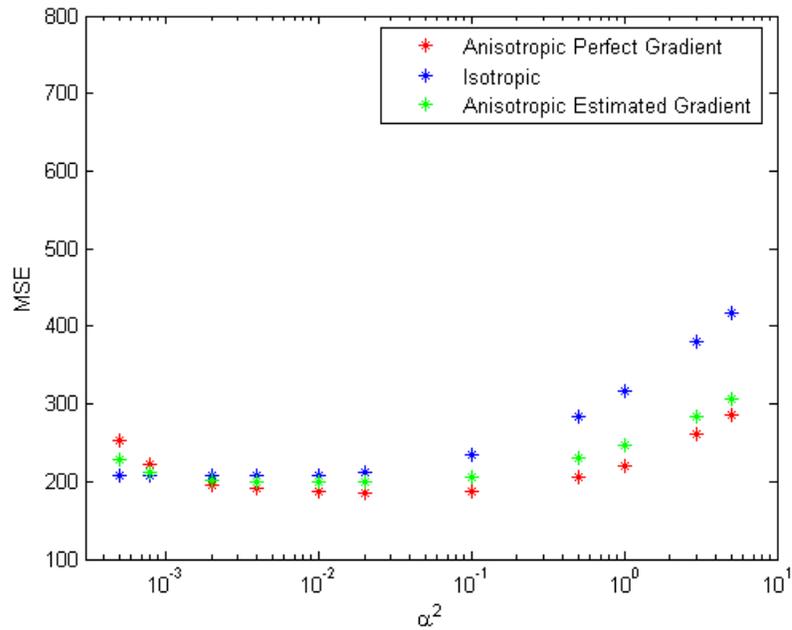
(d) Results for *Squares* image: Gaussian pre-filter and optical flow motion estimation.

Figure 5-6: MSE results for *Squares* image.

(a) Results for *Brick* image: no pre-filter and perfect motion.(b) Results for *Brick* image: no pre-filter and optical flow motion estimation.

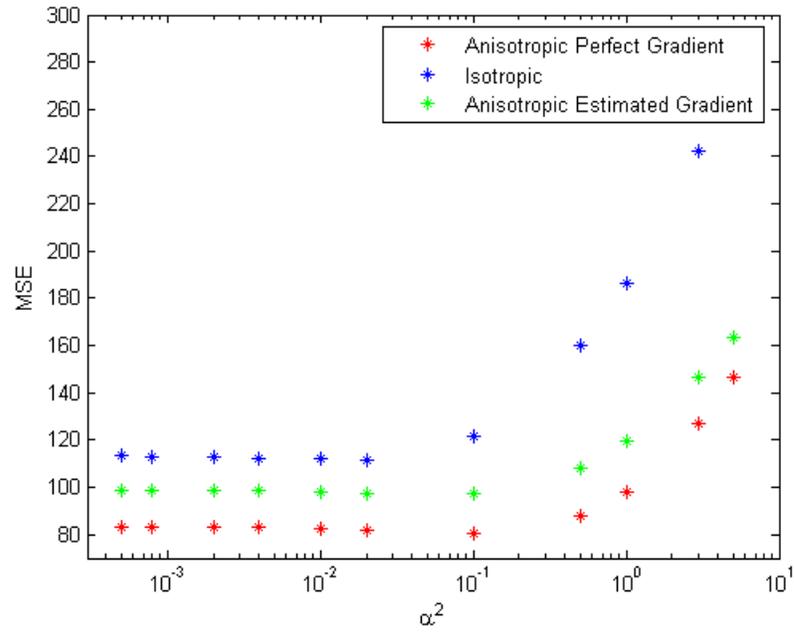
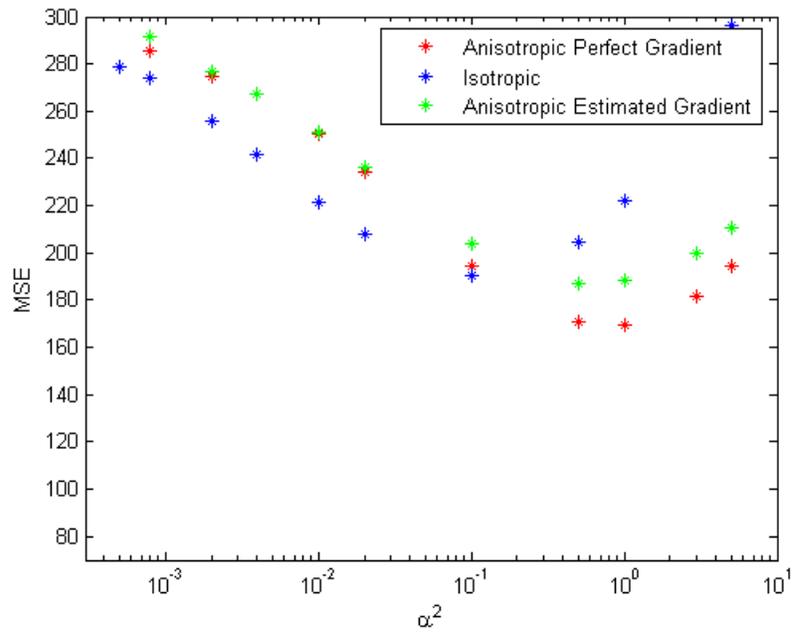


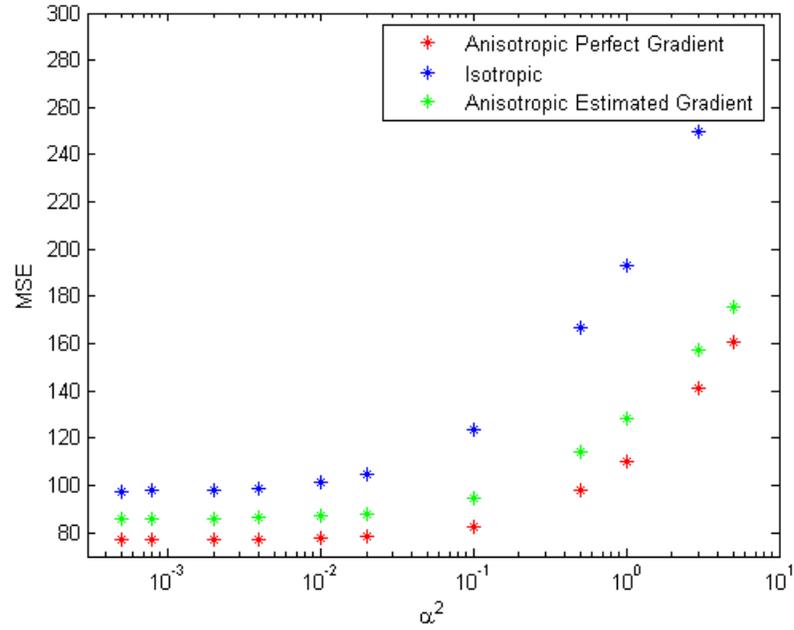
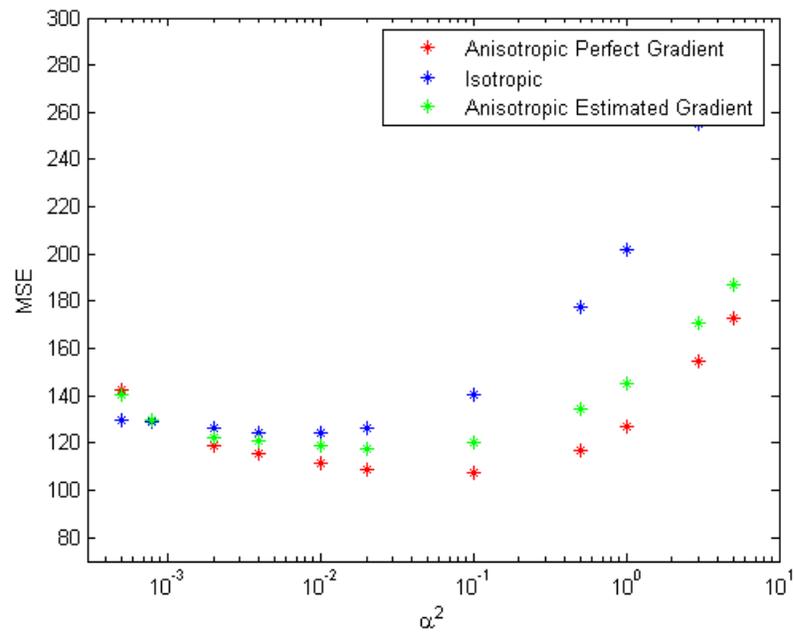
(c) Results for *Brick* image: Gaussian pre-filter and perfect motion.

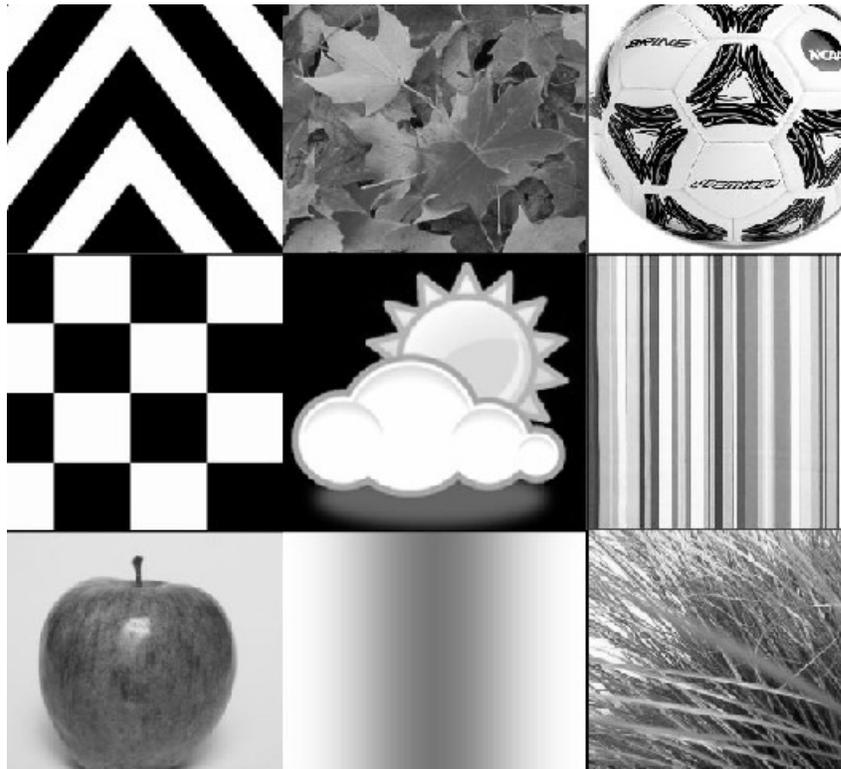
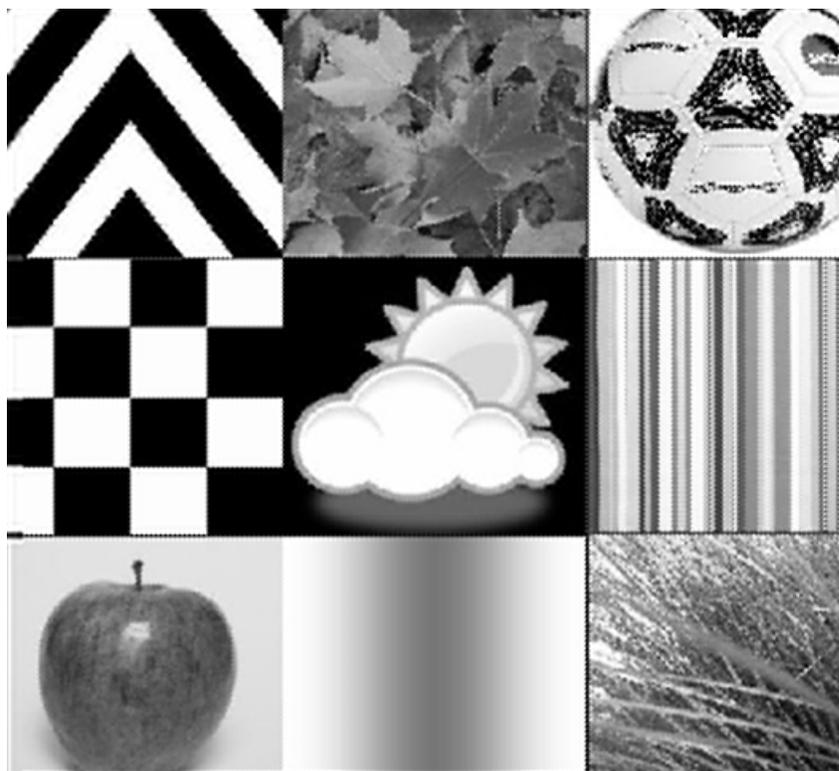


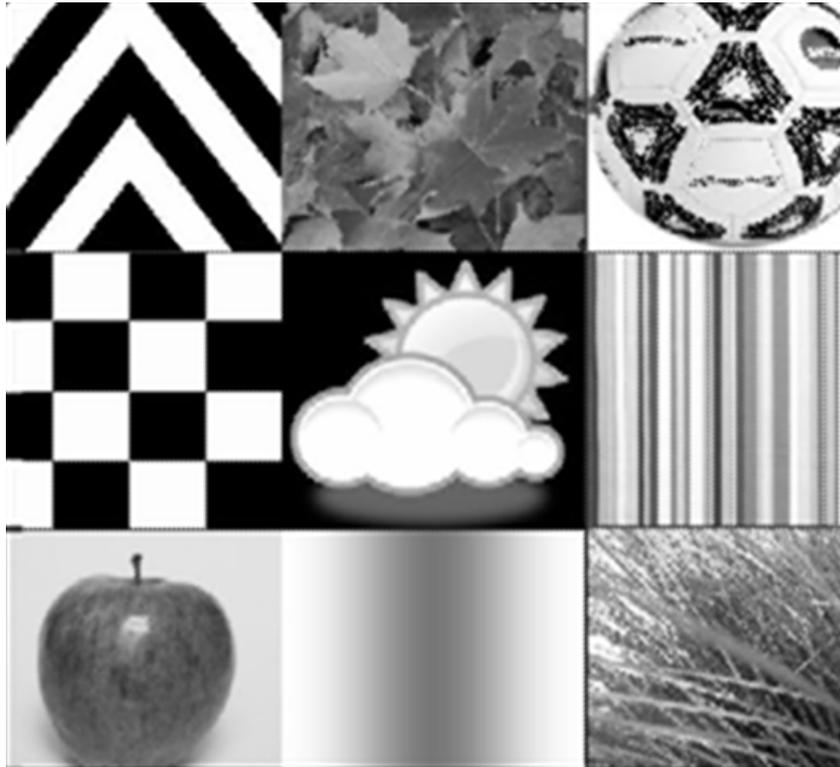
(d) Results for *Brick* image: Gaussian pre-filter and optical flow motion estimation.

Figure 5.7: MSE results for *Brick* image.

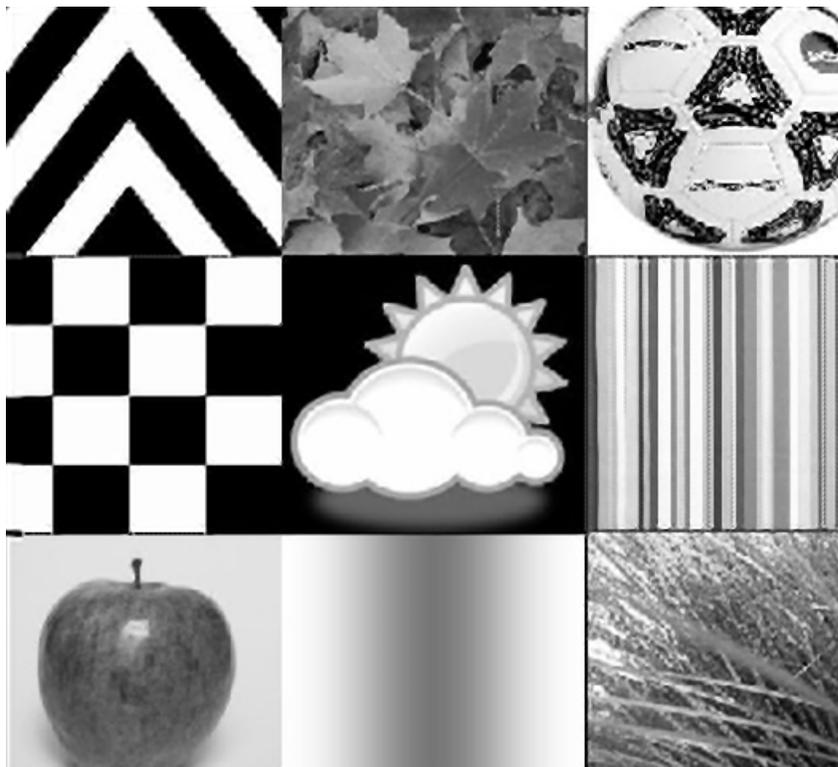
(a) Results for *Jeep* image: no pre-filter and perfect motion.(b) Results for *Jeep* image: no pre-filter and optical flow motion estimation.

(c) Results for *Jeep* image: Gaussian pre-filter and perfect motion.(d) Results for *Jeep* image: Gaussian pre-filter and optical flow motion estimation.**Figure 5-8:** MSE results for *Jeep* image.

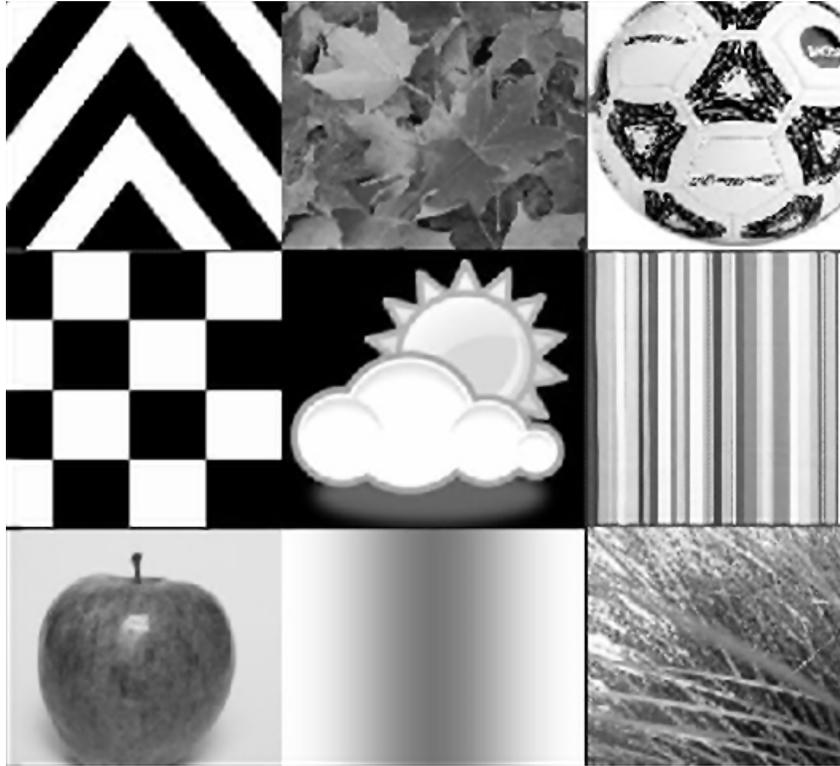
(a) Original HR *Squares* Image(b) Lowest-MSE SR reconstruction of *Squares*: no pre-filter and isotropic diffusion.



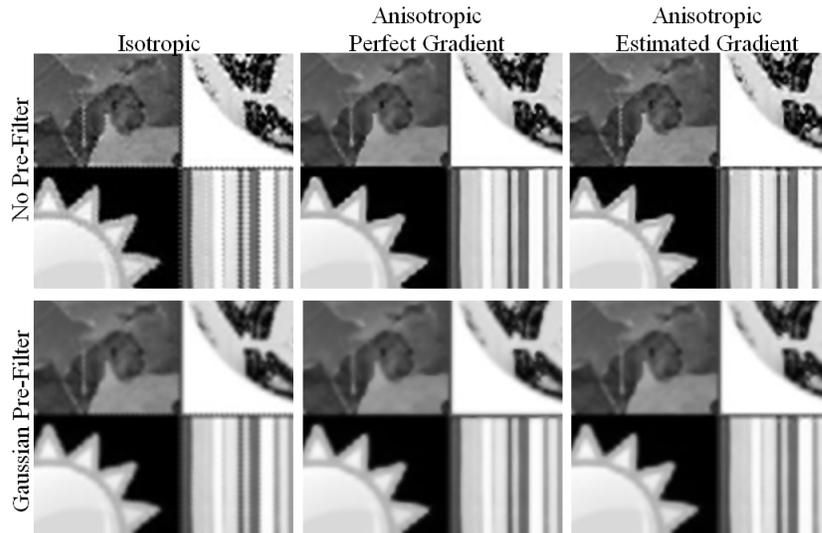
(c) Visually-optimal SR reconstruction of *Squares*: no pre-filter and isotropic diffusion.



(d) Lowest-MSE SR reconstruction of *Squares*: no pre-filter, anisotropic diffusion, perfect motion and perfect gradient.



(e) Lowest-MSE SR reconstruction of *Squares*: no pre-filter, anisotropic diffusion, perfect motion and estimated gradient.



(f) Zoomed-in comparison for *Squares* image.

Figure 5-9: SR reconstruction results for *Squares* HR image for synthetic-motion data and no pre-filter.



(a) Original HR *Brick* Image



(b) Lowest-MSE SR reconstruction of *Brick*: no pre-filter and isotropic diffusion.



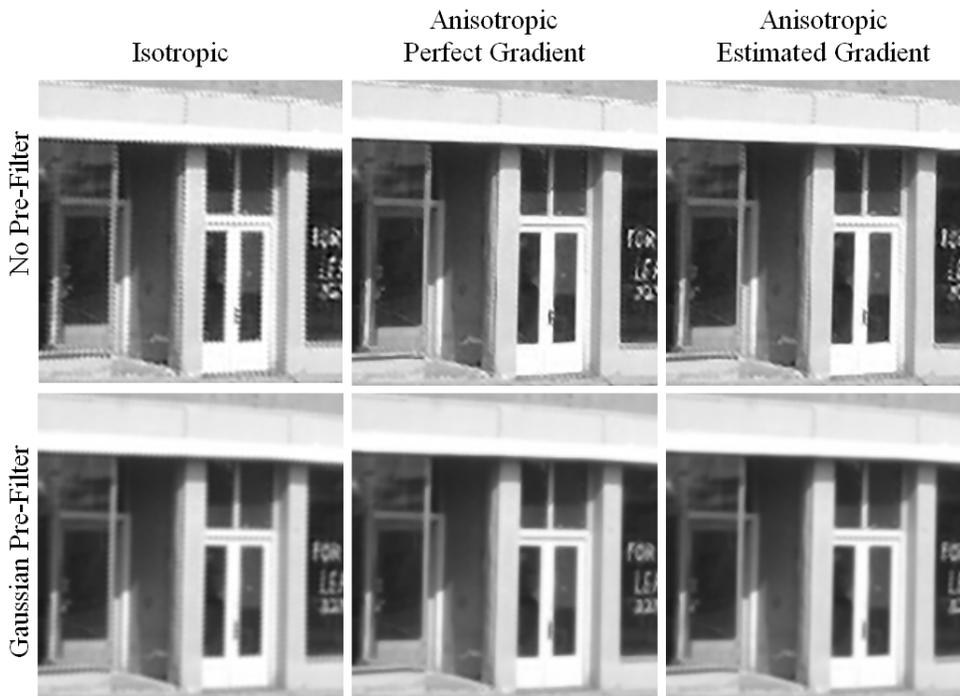
(c) Visually-optimal SR reconstruction of *Brick*: no pre-filter and isotropic diffusion.



(d) Lowest-MSE SR reconstruction of *Brick*: no pre-filter, anisotropic diffusion, perfect motion and perfect gradient.



(e) Lowest-MSE SR reconstruction of *Brick*: no pre-filter, anisotropic diffusion, perfect motion and estimated gradient.



(f) Zoomed-in comparison for *Brick* image.

Figure 5.10: SR reconstruction results for *Brick* HR image for synthetic-motion data and no pre-filter.



(a) Original HR *Jeep* Image



(b) Lowest-MSE SR reconstruction of *Jeep*: no pre-filter and isotropic diffusion.



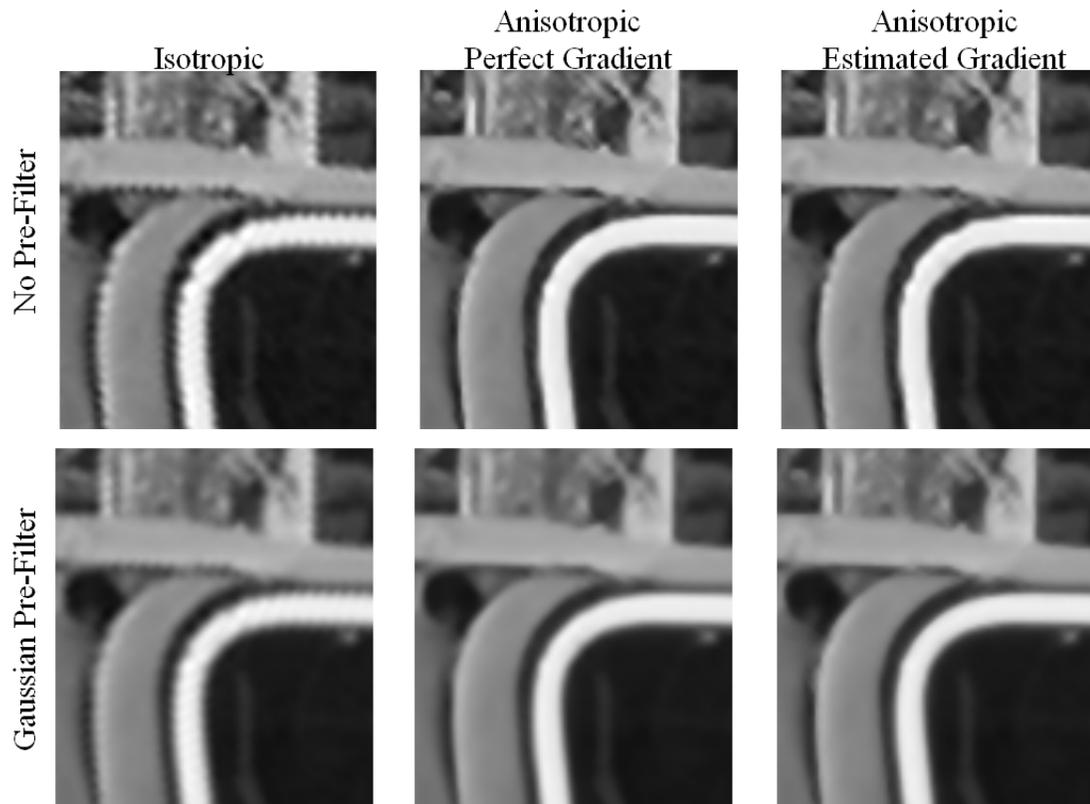
(c) Visually-optimal SR reconstruction of *Jeep*: no pre-filter and isotropic diffusion.



(d) Lowest-MSE SR reconstruction of *Jeep*: no pre-filter, anisotropic diffusion, perfect motion and estimated gradient.



(e) Lowest-MSE SR reconstruction of *Jeep*: no pre-filter, anisotropic diffusion, perfect motion and estimated gradient.



(f) Zoomed-in comparison for *Jeep* image.

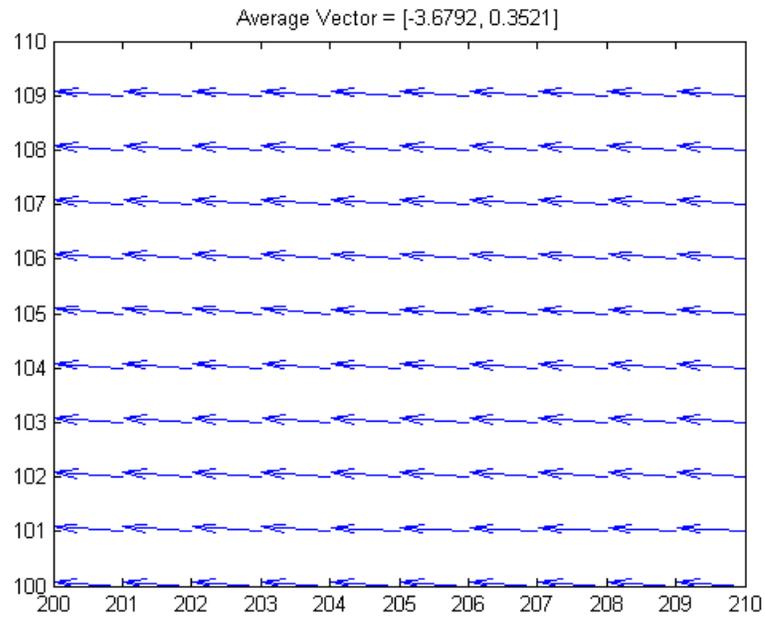
Figure 5-11: SR reconstruction results for *Jeep* HR image for synthetic-motion data and no pre-filter.

5.2 Real Motion Data Sets

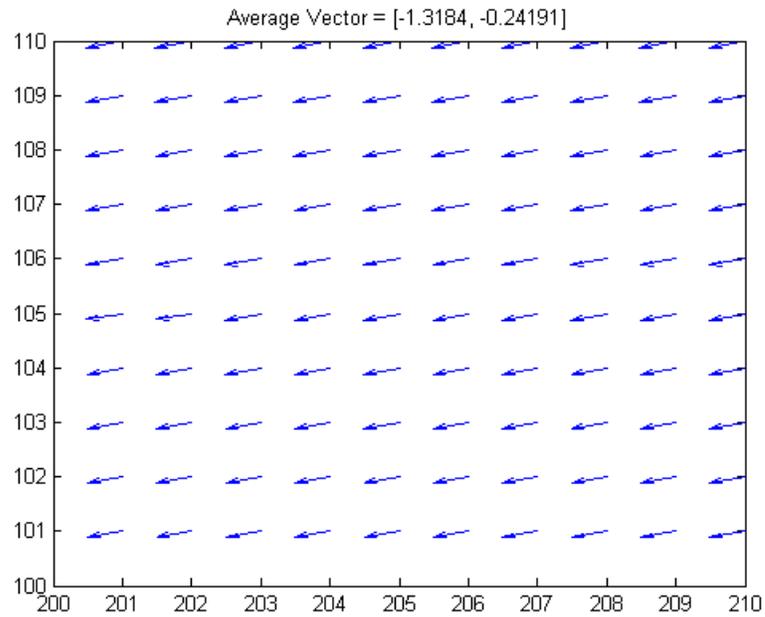
The results for the synthetic data sets demonstrate the edge-enhancing effects of the anisotropic SR reconstruction. They also demonstrate the need for precise motion estimation and the limitations introduced by blur. Any inaccuracies in the motion estimation result in artifacts in the SR image. Pixels which are misaligned and mapped to incorrect locations on the HR grid also result in errors in the gradient estimation which limits the effectiveness of the anisotropic diffusion. To further test the SR implementation, real video sequences were evaluated.

For the video sequence analyzed, three LR 240x352 frames were used to create a 960x1408 HR image, up-sampling by 4 in x and y . This sequence was captured by a camera subject to vibrations caused by wind load and nearby air conditioning equipment. The exact motion of the camera is unknown, however, it appears to be globally translational (far away objects, very small camera pan and tilt). Therefore, the optical flow estimation should yield a smooth vector field. Figure 5.12 displays a zoomed-in region of the optical flow field computed from for the first video sequence. As expected, the optical flow vector field is fairly uniform across the entire image.

In order to compare the best reconstructions for both methods, the regularization parameter was chosen based on the lowest MSE for synthetic data sets with optical flow estimation. For the synthetic data sets with optical flow estimation, a regularization weight of 0.1 for the anisotropic reconstruction and 0.004 for the isotropic reconstruction had the lowest MSE. The images in Figure 5.13 show the isotropic and anisotropic reconstruction results obtained from 3 video frames of the video sequences. The images are shown at a 25% scale, therefore it is very difficult to observe differences in the reconstructions. The improvements of the anisotropic method are more visible in the zoomed-in images shown in Figure 5.14.



(a) Zoomed-in region of normalized optical flow vectors between images 1 and 2 for real video sequence. Average (unnormalized) optical flow vector = $(-3.7692, 0.3521)$



(b) Zoomed-in region of normalized optical flow vectors between images 2 and 3 for real video sequence. Average (unnormalized) optical flow vector = $(-1.3184, -0.2419)$

Figure 5-12: Optical flow vectors for real video sequence.



(a) Results for isotropic reconstruction for $\alpha^2 = 0.004$.



(b) Results for anisotropic reconstruction for $\alpha^2 = 0.1$.

Figure 5-13: Isotropic and anisotropic SR reconstruction results for a real video sequence.

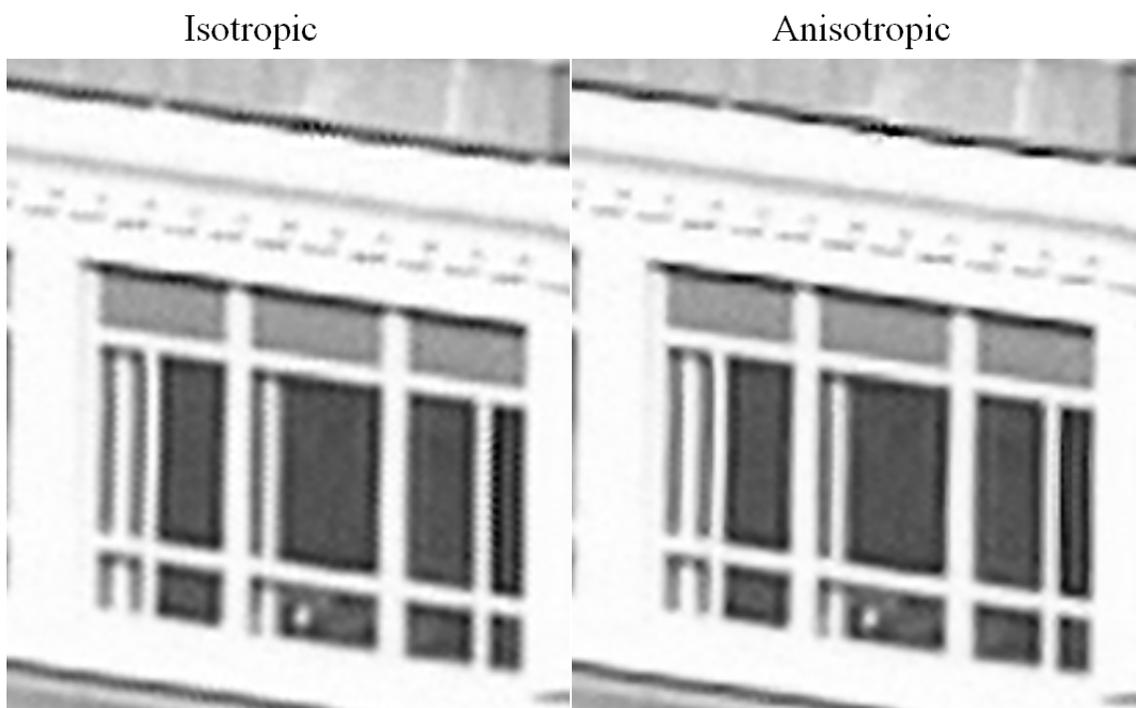


Figure 5.14: Zoomed-in region of isotropic and anisotropic SR reconstruction from Figure 5.13.

Chapter 6

Conclusions

6.1 Discussion of Results

In this thesis we have presented an anisotropic spline-based superresolution reconstruction method. Building on the work of Vazquez, the SR cost function proposed implements a discrete spline thin-plate model with edge-preserving anisotropic weighting. To construct the anisotropic weighting matrix, the gradient was estimated from the LR images. Although this estimate only provided a coarse estimate of the gradient, the anisotropic implementation still had a beneficial impact. In order to control the complexity of the experiments, the method was tested against several synthetic data sets with varying motion and blur parameters. The results were compared to Vazquez’s isotropic diffusion method visually and using MSE. The results of the synthetic data sets confirm the edge enhancing effects of the anisotropic implementation. However, the performance gain over the isotropic method is strongly dependent on accuracy of motion estimation. Under ideal conditions, when the motion is perfectly known, the anisotropic method significantly outperforms the isotropic approach.

We have also tested the impact of Gaussian pre-filtering (during the generation of synthetic-motion data sets) on the SR reconstruction results. Such a pre-filter minimizes aliasing in the LR images, however, it also limits the improvements of the HR reconstruction. The differences of pre-filter and no pre-filter case are not clearly visible through MSE, rather the effects of the Gaussian blur are more apparent through visual inspection. The blur limits the sharpness of the high-resolution reconstruction. The SR method proposed here models the underlying high-resolution image as a continuous spline function. Note,

however, that the observed LR images are modeled as perfect impulse samples of the continuous spline function. Therefore, there is no PSF introduced into the image formation model and thus our solution does not take such a PSF into account. As a result, the algorithm does not accurately represent the image formation model for LR images with blur. The inclusion of PSF would improve the results of the pre-filtered reconstruction and possibly make it better than the no pre-filter case.

Lastly, the anisotropic SR algorithm was tested against real video sequences. As expected, the anisotropic diffusion method shows only slight improvement over the isotropic implementation. Inaccuracies in motion estimation, unaccounted camera blur and compression artifacts hinder the improvements of the SR reconstruction.

6.2 Suggestions for Future Work

The SR algorithm proposed in this thesis assumes impulse sampling of the continuous spline function. To further improve this implementation, a PSF should be incorporated into the image formation model. As a result, the Φ matrix would be a function of both the 2-D cubic B-spline function and the PSF. Ideally, the PSF would exactly describe the imaging systems response to a point source. However, a rough approximation, such as a Gaussian blur kernel, may be sufficient to improve the anisotropic implementation. Including a PSF in the matrix Φ would more accurately model the image formation process.

This work may also be extended to applications with more complicated motion, for example moving objects. In this thesis, we evaluated data sets with a static scene and global translational motion. Most video sequences consist of more complicated motion, with moving objects, zoom, warping etc. However, prior to applying this implementation to other motion models, the motion estimation itself needs to be improved. The effects of inaccurate motion estimation have been demonstrated throughout this thesis.

Superresolution image restoration relies on aliasing in the LR images so that HR frequency content can be recovered. However, motion estimation does not perform as well when aliasing is present. These two steps of the superresolution process have conflicting re-

quirements. This SR implementation could be further improved with a motion estimation algorithm which can better handle aliased images. A possible solution may be to apply special filtering (based on some knowledge about aliasing present) to the LR images to smooth out the high-frequency aliasing prior to motion estimation.

References

- Borman, S. (2004). *Topics in Multiframe Superresolution Restoration*. PhD thesis, University of Notre Dame, Notre Dame, IN.
- Borman, S. and Stevenson, R. (1998). Super-resolution from image sequences, a review. In *Proceedings of the 1998 Midwest Symposium on Circuits and Systems*, pages 374–378, Notre Dame, IN, USA. IEEE.
- Borman, S. and Stevenson, R. (1999). Simultaneous multi-frame map super-resolution video enhancement using spatio-temporal priors. *International Conference on Image Processing*, pages 469–473.
- Chantas, G. K., Galatsanos, N. P., and Woods, N. A. (2007). Super-resolution based on fast registration and maximum *a Posteriori* reconstruction. *IEEE Transactions on Image Processing*, 16(7):1821–1830.
- Cheeseman, P., Kanefsky, B., Kraft, R., Stutz, J., and Hanson, R. (1994). Super-resolved surface reconstruction from multiple images. Technical Report FIA-94-12, NASA Ames Research Center, Moffett Field, CA.
- Dai, S., Han, M., Wu, Y., and Gong, Y. (2007). Bilateral back-projection for single image super resolution. *IEEE International Conference on Multimedia and Expo*, pages 1039–1042.
- Elad, M. and Feuer, A. (1997). Restoration of single super-resolution image from several blurred. *IEEE Transactions on Image Processing*, pages 1646–1658.
- Farsiu, S., Robinson, M. D., Elad, M., and Milanfar, P. (2004a). Advances and challenges in super-resolution. *International Journal of Imaging Systems and Technology*, 14(2):47–57.
- Farsiu, S., Robinson, M. D., Elad, M., and Milanfar, P. (2004b). Fast and robust multiframe super resolution. *IEEE Transactions on Image Processing*, 13(10):1327–1344.
- Hadamard, J. (1923). *Lectures on the Cauchy Problem in Linear Partial Differential Equations*. Yale University Press, New Haven, CT.
- Hardie, R., Barnard, K., Bognar, J., Armstrong, E., and Watson, E. (1998). High resolution image reconstruction from a sequence of rotated and translated frames and its application to an infrared imaging system. *Optical Engineering*, 37(1):247–260.

- Hardie, R. C. and Droege, D. R. (2007). A map estimator for simultaneous superresolution and detector nonuniformity correction. *European Association for Signal Processing Journal on Applied Signal Process*, 2007(1):206–206.
- Horn, B. and Schunck, B. (1981). Determining optical flow. *Artificial Intelligence*, pages 185–203.
- Hsu, J. T., Yen, C. C., Li, C. C., Sun, M., Tian, B., and Kaygusuz, M. (2004). Application of wavelet-based pocs superresolution for cardiovascular mri image enhancement. *Proceedings of the Third International Conference on Image and Graphics*, pages 217–221.
- Irani, M. and Peleg, S. (1990). Super resolution from image sequences. *International Conference on Pattern Recognition*, 90:115–120.
- Irani, M. and Peleg, S. (1993). Motion analysis for image enhancement resolution, occlusion, and transparency. *Journal of Visual Communication and Image Representation*, 4:324–335.
- Karl, W. C. (2000). Regularization in image restoration and reconstruction. *Handbook of Image and Video Processing*, pages 141–160.
- Kim, H., Jang, J., and Hong, K. (2004). Edge-enhancing super-resolution using anisotropic diffusion. *International Conference on Image Processing*, 3:1767–1770.
- Kim, S., Bose, N., and Valenzuela, H. (1990). Recursive reconstruction of high resolution image from noisy undersampled multiframes. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38:1013–1027.
- Kim, S. and Su, W. (1993). Subpixel accuracy image registration by spectrum cancellation. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 5:153–156.
- Mann, S. and Picard, R. (1994). Virtual bellows: Constructing high quality stills from video. In *International Conference on Image Processing*, volume 1, pages 363–367.
- Mansouri, A., Mitiche, A., and Konrad, J. (1998). Selective image diffusion: Application to disparity estimation. *International Conference on Image Processing*, 3:284–288.
- Nguyen, N., Milanfar, P., and Golub, G. (2001). A computationally efficient superresolution image reconstruction algorithm. *IEEE Transactions on Image Processing*, 10(4):573–583.
- Ozkan, M., Tekalp, A., and Sezan, M. (2001). Pocs-based restoration of space-varying blurred images. *International Conference on Image Processing*, 3:138–141.

- Park, S., Park, M., and Kang, M. (2003). Super-resolution image reconstruction, a technical overview. *IEEE Signal Processing Magazine*, 20:21–36.
- Patti, A. and Altunbasak, Y. (1998). Artifact reduction for pocs-based super resolution with edge adaptive regularization and higher-order interpolants. *International Conference on Image Processing*, 3:217–221.
- Perona, P. and Malik, J. (1990). Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern and Analysis and Machine Intelligence*, 12(7):629–639.
- Prendergast, R. and Nguyen, T. (2005). Improving frequency domain super-resolution via undersampling model. *International Conference on Image Processing*, 1.
- Schoenberg, I. (1992). Contribution to the problem of approximation of equidistant data by analytic functions. *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 3:169–172.
- Schultz, R. and Stevenson, R. (1994). A bayesian approach to image expansion for improved definition. *IEEE Transactions on Image Processing*, 3(3):233–242.
- Schulz, T. (2005). Multi-frame image restoration. *The Image and Video Processing Handbook*, pages 175–190.
- Segall, C. A., Katsaggelos, A. K., Molina, R., and Mateos, J. (2004). Bayesian resolution enhancement of compressed video. *IEEE Transactions on Image Processing*, 13(7):898–911.
- Shewchuck, J. (1994). An introduction to the conjugate gradient method without the agonizing pain.
- Smith, J. O. (2007). *Mathematics of the Discrete Fourier Transform (DFT)*. W3K Publishing.
- Stark, H. and Oskoui, P. (1989). High-resolution image recovery from image-plane arrays, using convex projections. *Optical Society of America*, 6:1715–1726.
- Stasinski, R. and Konrad, J. (2006). Pocs reconstruction of irregularly-sampled images based on oversampling and linear space-variant filtering. *Sampling Theory in Signal and Image Processing*, 5:37–58.
- Tekalp, A., Ozkan, M., and Sezan, M. (1992). High-resolution image reconstruction from lower-resolution image sequences and space-varying image restoration. *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 3:169–172.
- Tsai, R. and Huang, T. (1984). Multiframe image restoration and registration. *Advances in Computer Vision and Image Processing*, 1:317–339.

- Unser, M. (1999). Splines, a perfect fit for signal and image processing. *IEEE Signal Processing Magazine*, pages 22–38.
- Vazquez, C. (2002). *Reconstruction D’Images a Partir D’Echantillons Irreguliere-ment Espaces*. PhD thesis, Universite du Quebec.
- Vazquez, C., Aly, H., Dubois, E., and Mitiche, A. (2004). Motion compensated super-resolution of video by level sets evolution. *International Conference on Image Processing*, 3:1767–1770.
- Vazquez, C., Dubois, E., and Konrad, J. (2003). Reconstruction of irregularly-sampled images in spline spaces. *IEEE Transactions on Image Processing*, 3:405–408.