Topics in

Multiframe Superresolution Restoration

A Dissertation

Submitted to the Graduate School

of the University of Notre Dame

in Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

by

Sean Borman

April 2, 2004

Roadmap

- Overview of the talk
- Background [Chap. 2]
- Simultaneous multiframe superresolution restoration [Chap. 7]
- Linear models for multiframe superresolution restoration [Chap. 8]
- Conclusions [Chap. 9]

Classical Single Frame Restoration

- The problem:

Given:	Degraded image	$\boldsymbol{g}$	$\in$	$\mathbb{R}^{N_r N_c \times 1}$
	Observation process	$\boldsymbol{A}$	$\in$	$\mathbb{R}^{N_r N_c \times N_r N_c}$
Find:	Restored image	$\boldsymbol{f}$	$\in$	$\mathbb{R}^{N_r N_c \times 1}$

- Graphical representation:

$$\begin{array}{cl} \boldsymbol{f} & \text{Unknown undegraded image} \\ \downarrow & \\ \boldsymbol{A} & \text{Observation model} \\ \downarrow & \\ \boldsymbol{g} & \text{Degraded observation} \end{array}$$

- Algebraically: $\boldsymbol{g} = \boldsymbol{A}\boldsymbol{f}$
- Single-Input Single-Output (SISO)
- Classical Linear Inverse Problem

Single Frame Restoration and Interpolation

- The problem:

Given:	Degraded image	$\boldsymbol{g}$	$\in$	$\mathbb{R}^{N_r N_c \times 1}$
	Observation process	$\boldsymbol{A}$	$\in$	$\mathbb{R}^{N_r N_c \times q^2 N_r N_c}$
Find:	Restored image	$\boldsymbol{f}$	$\in$	$\mathbb{R}^{q^2 N_r N_c \times 1}$

- Magnification factor $q \in \mathbb{N}$ to <u>increase resolution</u>
- <u>More unknowns than constraints</u>
- Well-posed problem (Hadamard) is one whose solution . . .
  1. exists
  2. is unique, and
  3. exhibits continuous dependence on the data
- This problem is ill-posed

Regularization and the Solution of Ill-Posed Problems

- Many methods for addressing ill-posed problems
- Generally referred to as regularized solution methods
- All are based on one idea
- Use prior knowledge to constrain the solution
- Schultz and Stevenson: Markov Random Field (MRF) image model

How About More Constraints?

- Use multiple observed frames $\boldsymbol{g}^{(1)}, \boldsymbol{g}^{(2)}, \ldots, \boldsymbol{g}^{(P)}$

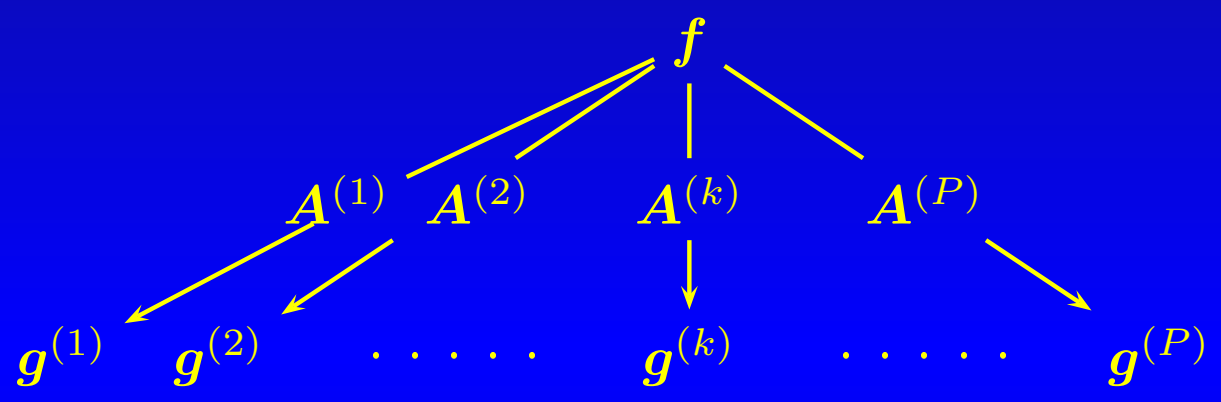

- Multi-Input Single-Output (MISO)
- $\boldsymbol{f}$ is constrained by all related observations
- or generalize futher . . .
- Multi-Input Multi-Output (MIMO)

Superresolution

- Increasing the number of samples (interpolation)
- This isn't superresolution!
- Superresolution = Bandwidth extrapolation:
  - beyond the diffraction limit
  - beyond the optical system cutoff
  - beyond the complete imaging system cutoff
- Is this really possible?
- Yes . . . (although some didn't think so)
  - Additional constraints from multiple observations (motion)
  - Inversion of aliased terms
  - Nonlinear a-priori information

Roadmap

- Overview of the talk
- Background [Chap. 2]
- Simultaneous multiframe superresolution restoration [Chap. 7]
- Linear models for multiframe superresolution restoration [Chap. 8]
- Conclusions [Chap. 9]

Simultaneous Multiframe Superresolution Restoration

- Given a noisy, under-sampled low-resolution image sequence, restore a superresolved sequence
- Use information from *multiple* observed images for restoration
- Subpixel motion contributes additional constraints
- Use *a-priori* knowledge
- Superresolution = bandwidth extrapolation
  - Optical system diffraction limit
  - Imaging system band limit

Existing Approaches

- Multi-input, single output

Proposed Approach

- Multi-input, multi-output

Proposed Approach

- Given observed image sequence, *simultaneously* solve for all restored images
- Allows incorporation of *temporal* constraints in addition to spatial constraints
- Include confidence parameters for subpixel motion estimates

Summary of the Problem

Given:	A short video sequence P observed LR frames N_r rows by N_c columns
Find:	P corresponding superresolved images qN_r rows by qN_c columns, $q \in \mathbb{N}$
Approach:	Linear inverse problem observation model Stochastic regularization (ill-posedness) Bayesian MAP estimate with MRF prior

Observation Model

- Relate observed LR sequence to unknown SR sequence
- Utilize linear model $\Longrightarrow$ linear inverse problem
- Observation model must include effects of:
  - motion
  - imaging degradation
- Identify two observation modes
  - temporally coincident
  - temporally non-coincident

Temporally Coincident Observation

- Observe LR frame $\boldsymbol{g}^{(l)}$ from SR frame $\boldsymbol{f}^{(l)}$
- Assume *known* SVPSF degradation $\boldsymbol{A}^{(l,l)}$

$$
\begin{array}{ccccccc}
\boldsymbol{f}^{(1)} & \boldsymbol{f}^{(2)} & \cdots\cdots & \boldsymbol{f}^{(l)} & \cdots\cdots & \boldsymbol{f}^{(P)} \\
| & | & & | & & | \\
\boldsymbol{A}^{(1,1)} & \boldsymbol{A}^{(2,2)} & & \boldsymbol{A}^{(l,l)} & & \boldsymbol{A}^{(P,P)} \\
\downarrow & \downarrow & & \downarrow & & \downarrow \\
\boldsymbol{g}^{(1)} & \boldsymbol{g}^{(2)} & \cdots\cdots & \boldsymbol{g}^{(l)} & \cdots\cdots & \boldsymbol{g}^{(P)}
\end{array}
$$

Temporally Coincident Observation

- Additive noise (measurement errors / errors in $\boldsymbol{A}^{(l,l)}$)

$$\boldsymbol{g}^{(l)} = \boldsymbol{A}^{(l,l)} \boldsymbol{f}^{(l)} + \boldsymbol{n}^{(l,l)}, \quad l \in \mathbb{P},$$

where,

$$\begin{aligned} \mathbb{P} &= \{1, 2, \ldots, P\} \\ \boldsymbol{g}^{(l)}, \boldsymbol{n}^{(l)} &\in \mathbb{R}^{N_r N_c \times 1} \\ \boldsymbol{f}^{(l)} &\in \mathbb{R}^{q^2 N_r N_c \times 1} \\ \boldsymbol{A}^{(l,l)} &\in \mathbb{R}^{N_r N_c \times q^2 N_r N_c} \end{aligned}$$

Temporally Coincident Observation

- Noise model — IID Gaussian, zero mean, per-pixel variances ${\sigma_j^{(l,l)}}^2$

$$\mathcal{P}\left(\boldsymbol{n}^{(l,l)}\right) = \frac{1}{(2\pi)^{\frac{N_r N_c}{2}} |\boldsymbol{K}^{(l,l)}|} \exp\left\{-\frac{1}{2}{\boldsymbol{n}^{(l,l)}}^\top {\boldsymbol{K}^{(l,l)}}^{-1} \boldsymbol{n}^{(l,l)}\right\}$$

$$\boldsymbol{K}^{(l,l)} = \mathrm{diag}\left({\sigma_1^{(l,l)}}^2, {\sigma_2^{(l,l)}}^2, \cdots, {\sigma_{N_r N_c}^{(l,l)}}^2\right)$$

- Observation noise assumed independent over $\{\boldsymbol{g}^{(l)}\}_1^P$

Temporally Non-Coincident Observation

- Motion compensated observation — more constraints
- Relate $\{\boldsymbol{f}^{(k)}\}_1^P$ to temporally non-coincident LR frame
- Find: *motion compensating* observation matrix $\boldsymbol{A}^{(l,k)}$ and vector $\boldsymbol{n}^{(l,k)}$ s.t.

$$\boldsymbol{g}^{(l)} = \boldsymbol{A}^{(l,k)}\boldsymbol{f}^{(k)} + \boldsymbol{n}^{(l,k)}\,, \quad l \in \mathbb{P} \setminus \{k\}$$

$$\mathcal{P}\left(\boldsymbol{n}^{(l,k)}\right) = \frac{1}{(2\pi)^{\frac{N_r N_c}{2}} |\boldsymbol{K}^{(l,k)}|} \exp\left\{-\frac{1}{2}\boldsymbol{n}^{(l,k)\top}\boldsymbol{K}^{(l,k)^{-1}}\boldsymbol{n}^{(l,k)}\right\}$$

Combined Observation Model

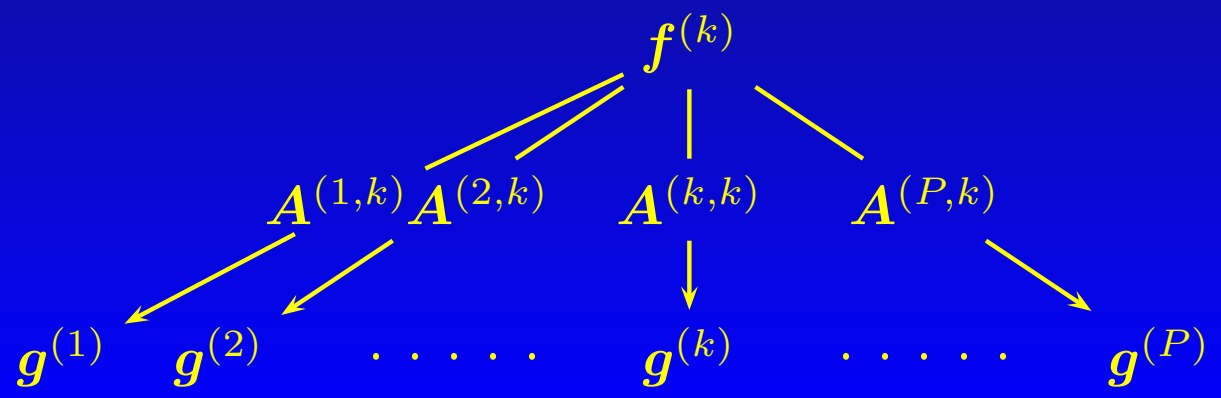

$$\begin{bmatrix} \boldsymbol{g}^{(1)} \\ \boldsymbol{g}^{(2)} \\ \vdots \\ \boldsymbol{g}^{(P)} \end{bmatrix} = \begin{bmatrix} \boldsymbol{A}^{(1,k)} \\ \boldsymbol{A}^{(2,k)} \\ \vdots \\ \boldsymbol{A}^{(P,k)} \end{bmatrix} \boldsymbol{f}^{(k)} + \begin{bmatrix} \boldsymbol{n}^{(1,k)} \\ \boldsymbol{n}^{(2,k)} \\ \vdots \\ \boldsymbol{n}^{(P,k)} \end{bmatrix}$$

$$\boldsymbol{g} = \boldsymbol{A}^{(*,k)} \boldsymbol{f}^{(k)} + \boldsymbol{n}^{(*,k)}$$

Simultaneous Multiframe Observation

Simultaneous observation model for $\{\boldsymbol{f}^{(k)}\}_1^P$

$$\boldsymbol{g}^{(l)} = \boldsymbol{A}^{(l,k)}\boldsymbol{f}^{(k)} + \boldsymbol{n}^{(l,k)}\ , \quad l,k \in \mathbb{P}$$

$$\begin{bmatrix} \boldsymbol{g} \\ \boldsymbol{g} \\ \vdots \\ \boldsymbol{g} \end{bmatrix} = \begin{bmatrix} \boldsymbol{A}^{(*,1)} & & & \\ & \boldsymbol{A}^{(*,2)} & & \\ & & \ddots & \\ & & & \boldsymbol{A}^{(*,P)} \end{bmatrix} \begin{bmatrix} \boldsymbol{f}^{(1)} \\ \boldsymbol{f}^{(2)} \\ \vdots \\ \boldsymbol{f}^{(P)} \end{bmatrix} + \begin{bmatrix} \boldsymbol{n}^{(*,1)} \\ \boldsymbol{n}^{(*,2)} \\ \vdots \\ \boldsymbol{n}^{(*,P)} \end{bmatrix}$$

This may be expressed in the classic form:

$$\boldsymbol{G} = \boldsymbol{A}\boldsymbol{f} + \boldsymbol{N}$$

Combined Noise Model

$$\mathcal{P}_{\boldsymbol{N}}(\boldsymbol{N}) = \frac{1}{(2\pi)^{\frac{P^2 N_r N_c}{2}} |\boldsymbol{K}|} \exp\left\{-\frac{1}{2}\boldsymbol{N}^\top \boldsymbol{K}^{-1}\boldsymbol{N}\right\}$$

$$\begin{aligned}\boldsymbol{K} \quad = \quad \mathrm{diag}\,\{ \quad & \boldsymbol{K}^{(1,1)}, \boldsymbol{K}^{(2,1)}, \ldots, \boldsymbol{K}^{(P,1)}, \\ & \boldsymbol{K}^{(1,2)}, \boldsymbol{K}^{(2,2)}, \ldots, \boldsymbol{K}^{(P,2)}, \\ & \vdots \\ & \boldsymbol{K}^{(1,P)}, \boldsymbol{K}^{(2,P)}, \ldots, \boldsymbol{K}^{(P,P)} \quad \}\end{aligned}$$

- $\boldsymbol{K}$ (diagonal) represents *a-priori* knowledge concerning observations and accuracy of motion information

Bayesian Estimation of SR Images

Maximize the a-posteriori probability:

$$\begin{aligned}
\hat{\boldsymbol{f}}_{\mathrm{MAP}} &= \arg\max_{\boldsymbol{f}} \left\{ \mathcal{P}\left(\boldsymbol{f}|\boldsymbol{G}\right) \right\} \\
&= \arg\max_{\boldsymbol{f}} \left\{ \frac{\mathcal{P}\left(\boldsymbol{G}|\boldsymbol{f}\right)\mathcal{P}\left(\boldsymbol{f}\right)}{\mathcal{P}\left(\boldsymbol{G}\right)} \right\} \\
&= \arg\max_{\boldsymbol{f}} \left\{ \log \mathcal{P}\left(\boldsymbol{G}|\boldsymbol{f}\right) + \log \mathcal{P}\left(\boldsymbol{f}\right) \right\}
\end{aligned}$$

$\mathcal{P}\left(\boldsymbol{G}|\boldsymbol{f}\right)$ — Likelihood function

$\mathcal{P}\left(\boldsymbol{f}\right)$ — Prior

The Likelihood Function

- Determined by noise pdf:

$$\begin{aligned}\mathcal{P}\left(\boldsymbol{G}|\boldsymbol{f}\right) &= \mathcal{P}_{\boldsymbol{N}}(\boldsymbol{G}-\boldsymbol{A}\boldsymbol{f}) \\ &= \frac{1}{(2\pi)^{\frac{q^2 N_r N_c}{2}}|\boldsymbol{K}|} \cdot \\ &\qquad \exp\left\{-\frac{1}{2}(\boldsymbol{G}-\boldsymbol{A}\boldsymbol{f})^\top \boldsymbol{K}^{-1}(\boldsymbol{G}-\boldsymbol{A}\boldsymbol{f})\right\}\end{aligned}$$

- Choose $\boldsymbol{K}$ according to confidence in observations
  - noise variance (temporally coincident)
  - motion estimator MAD measure (non-coincident)

Image Prior Model

- Markov random field (MRF)
- Local interaction $\Longleftrightarrow$ global pdf
- Pdf is Gibbsian (Hammersley-Clifford)

$$\mathcal{P}(\boldsymbol{f}) = \frac{1}{k_p} \exp\left\{ -\frac{1}{\beta} \sum_{c \in \mathcal{C}} \rho_\alpha \left(\boldsymbol{\partial}_c \boldsymbol{f} \right) \right\}$$

- Spatiotemporal smoothness
- Convex nonlinear activity penalization function $\rho_\alpha(\cdot)$

Huber Penalty Function

$$\rho_\alpha(x) = \begin{cases} x^2 & |x| \le \alpha \\ 2\alpha|x| - \alpha^2 & |x| > \alpha \end{cases}$$

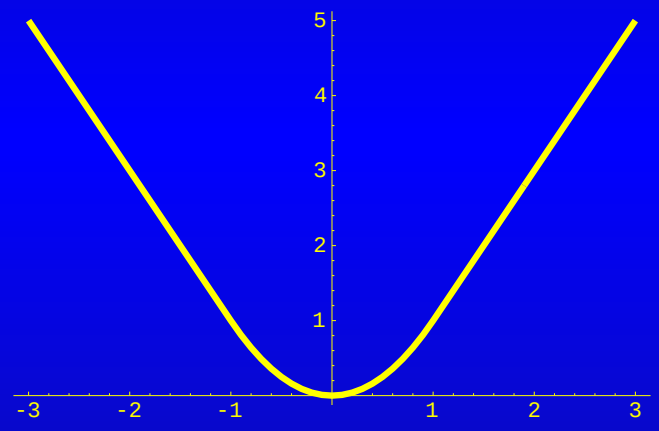

Spatiotemporal Clique Structure

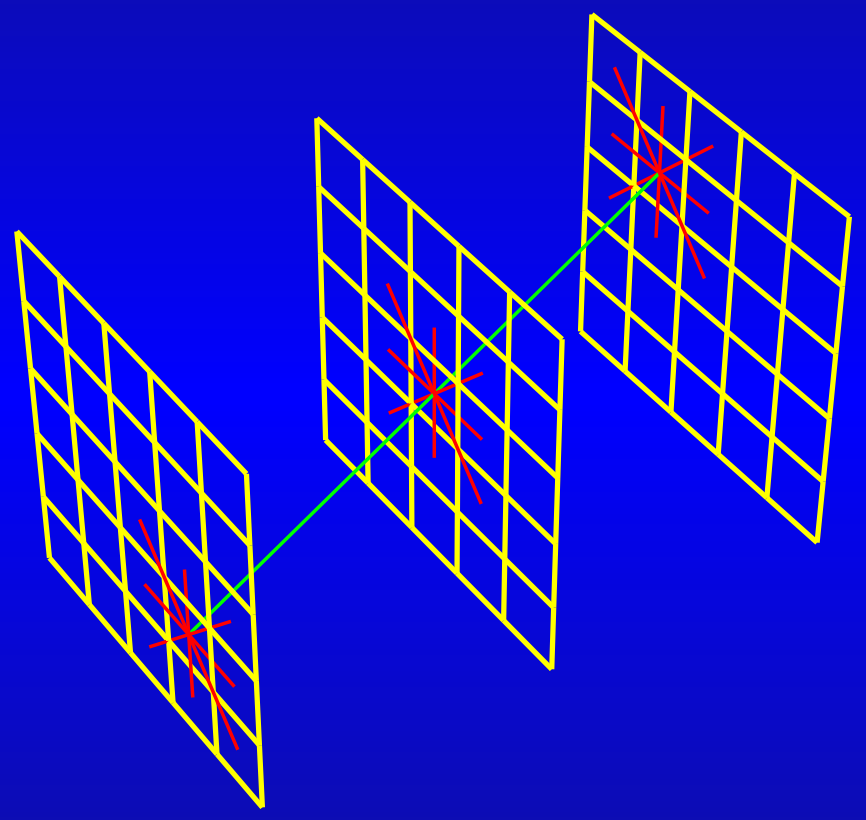

Spatiotemporal Clique Structure

Spatial activity measure — spatial derivatives

$$\boldsymbol{\partial}_1 \boldsymbol{f}^{(k)} = \quad f^{(k)}[n_1-1,n_2] \quad -2f^{(k)}[n_1,n_2] + \quad f^{(k)}[n_1+1,n_2]$$

$$\boldsymbol{\partial}_2 \boldsymbol{f}^{(k)} = \quad f^{(k)}[n_1,n_2-1] \quad -2f^{(k)}[n_1,n_2] + \quad f^{(k)}[n_1,n_2+1]$$

$$\boldsymbol{\partial}_3 \boldsymbol{f}^{(k)} = \tfrac{1}{2} f^{(k)}[n_1+1,n_2-1] - \quad f^{(k)}[n_1,n_2] + \tfrac{1}{2} f^{(k)}[n_1-1,n_2+1]$$

$$\boldsymbol{\partial}_4 \boldsymbol{f}^{(k)} = \tfrac{1}{2} f^{(k)}[n_1-1,n_2-1] - \quad f^{(k)}[n_1,n_2] + \tfrac{1}{2} f^{(k)}[n_1+1,n_2+1]$$

Temporal smoothness — time derivative along motion trajectory

$$\boldsymbol{\partial}_5 \boldsymbol{f}^{(k)} = f^{(k-1)}[n_1+\delta_1,n_2+\delta_2] - 2f^{(k)}[n_1,n_2] + f^{(k+1)}[n_1+\Delta_1,n_2+\Delta_2]$$

Objective Function

- Combine likelihood and prior terms

$$\hat{\boldsymbol{f}}_{\mathrm{MAP}} = \arg\max_{\boldsymbol{f}} \left\{ -\frac{1}{2}(\boldsymbol{G}-\boldsymbol{A}\boldsymbol{f})^{\top}\boldsymbol{K}^{-1}(\boldsymbol{G}-\boldsymbol{A}\boldsymbol{f}) - \frac{1}{\beta}\sum_{c\in\mathcal{C}}\rho_{\alpha}(\boldsymbol{\partial}_c\boldsymbol{f}) \right\}$$

$$\hat{\boldsymbol{f}}_{\mathrm{MAP}} = \arg\min_{\boldsymbol{f}} \left\{ \frac{1}{2}(\boldsymbol{G}-\boldsymbol{A}\boldsymbol{f})^{\top}\boldsymbol{K}^{-1}(\boldsymbol{G}-\boldsymbol{A}\boldsymbol{f}) + \frac{1}{\beta}\sum_{c\in\mathcal{C}}\rho_{\alpha}(\boldsymbol{\partial}_c\boldsymbol{f}) \right\}$$

Optimization

- Convex prior ensures convexity
- $\Longrightarrow$ Existence, uniqueness
- Convex, non-quadratic optimization
- Gradient descent optimization applicable
- Tractable even though dimensions are high

Original Low Resolution Frames

Zero Order Hold

Cubic Spline Interpolation

Superresolution Restoration

Superresolution vs Cubic Spline

Roadmap

- Overview of the talk
- Background [Chap. 2]
- Simultaneous multiframe superresolution restoration [Chap. 7]
- Linear models for multiframe superresolution restoration [Chap. 8]
- Conclusions [Chap. 9]

Introduction

- Multiframe superresolution (SR) restoration
  - Image restoration from an image sequence
  - Exceeds resolving ability of classical single-frame methods
- Objectives
  - Generalize linear multiframe observation model to . . .
    1. Nonaffine image registration (e.g. projectivity)
    2. Spatially-varying PSF
  - Must be compatible with existing restoration framework
- Approach
  - Use result from image resampling/warping (computer graphics)
  - Propose algorithm for computing generalized observation model
  - Demonstrate applicaton to multiframe superresolution experiment

Conceptual Image Resampling Pipeline (Heckbert)

$f(\boldsymbol{u})$ Discrete texture $\boldsymbol{u} = [u\ v]^T \in \mathbb{Z}^2$

$\downarrow$

$r(\boldsymbol{u})$ Reconstruction filter

$f_c(\boldsymbol{u})$

$\downarrow$

$H(\boldsymbol{u})$ Geometric transform (warp) $H : \boldsymbol{u} \longmapsto \boldsymbol{x}$

$g_c(\boldsymbol{x})$

$\downarrow$

$p(\boldsymbol{x})$ Anti-alias prefilter

$g'_c(\boldsymbol{x})$

$\downarrow$

Δ Sample

$\downarrow$

$g(\boldsymbol{x})$ Discrete resampled image $\boldsymbol{x} = [x\ y]^T \in \mathbb{Z}^2$

Realizable Image Resampling Pipeline (Heckbert)

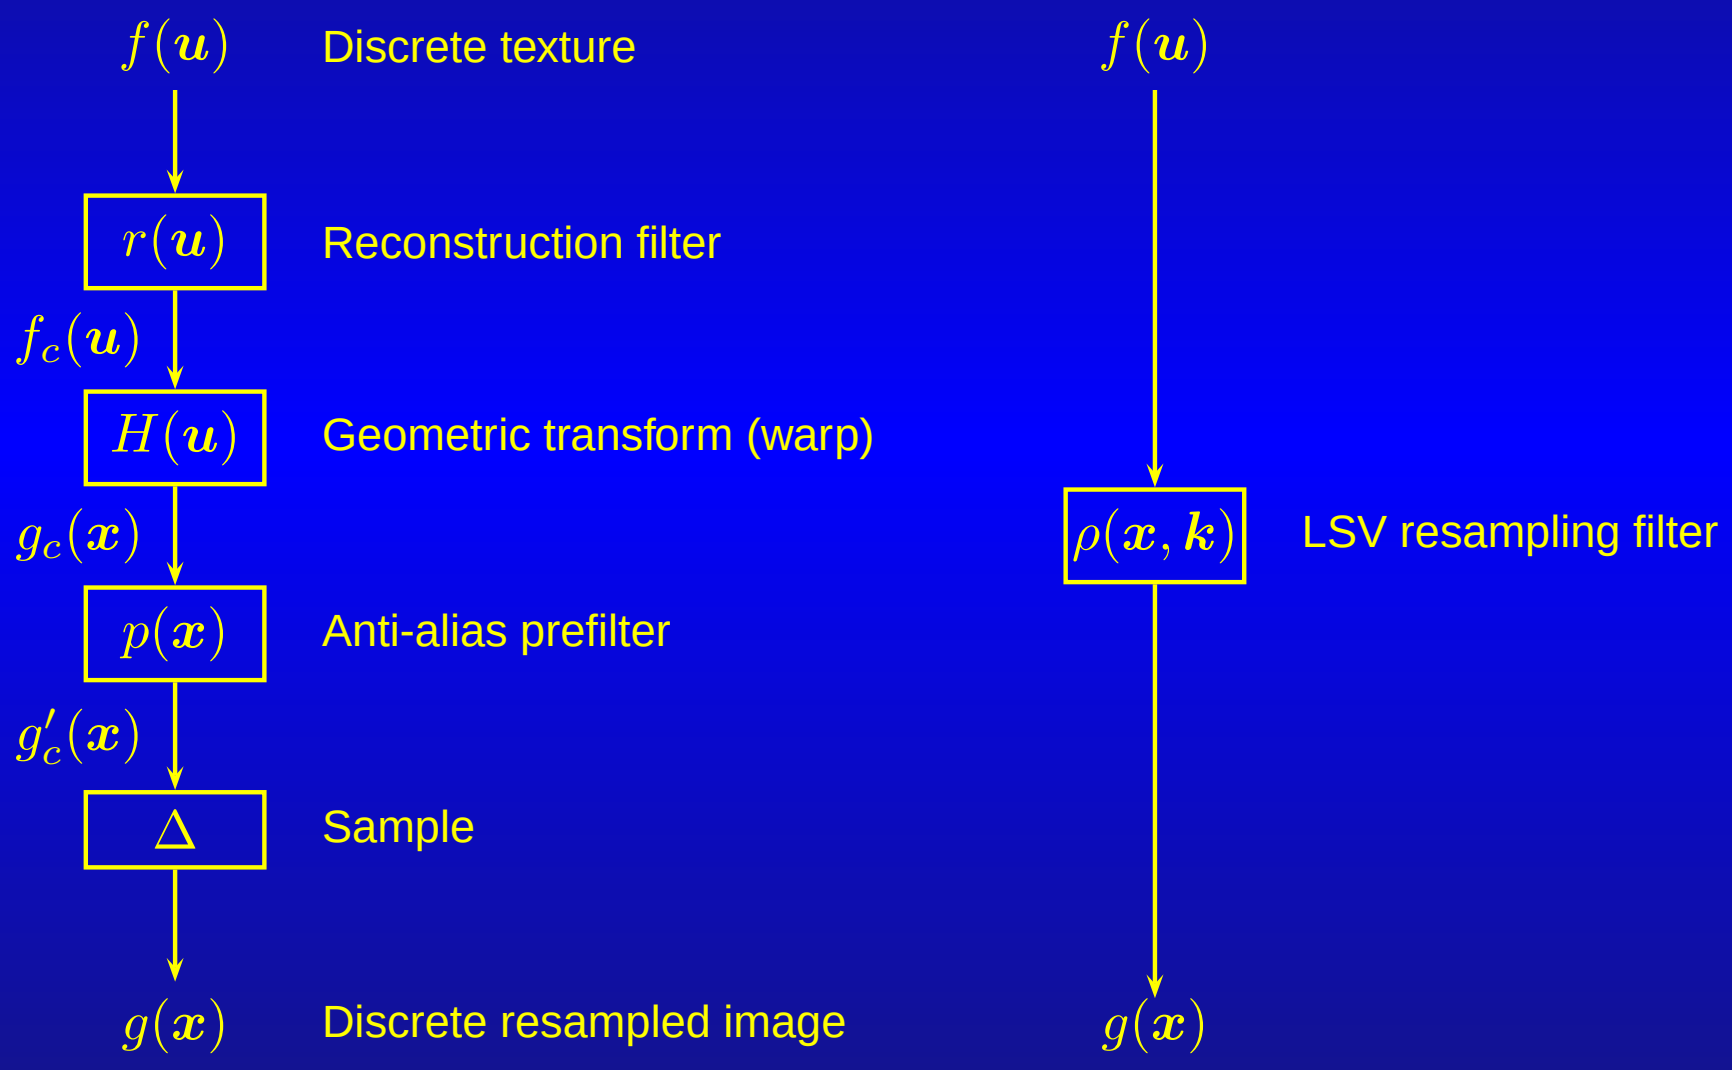

LSV Image Resampling Filter

- Compute warped image from texture using

$$g(\boldsymbol{x}) = \sum_{\boldsymbol{k} \in \mathbb{Z}^2} f(\boldsymbol{k}) \cdot \rho(\boldsymbol{x}, \boldsymbol{k}) \quad \text{for} \quad \boldsymbol{x}, \boldsymbol{k} \in \mathbb{Z}^2$$

- $\rho(\boldsymbol{x}, \boldsymbol{k})$ is a discrete linear, spatially varying (LSV) resampling filter

$$\rho(\boldsymbol{x}, \boldsymbol{k}) = \int p(\boldsymbol{x} - H(\boldsymbol{u})) \cdot r(\boldsymbol{u} - \boldsymbol{k}) \left| \frac{\partial H}{\partial \boldsymbol{u}} \right| \mathrm{d}\boldsymbol{u}$$

- The LSV resampling filter . . .
  - depends on the warp H
  - depends on the reconstruction filter r
  - is expressed in terms of a *warped* prefilter p
  - involves integration in texture space ($\boldsymbol{u}$)

Multiframe Observation Model

- Given images $g^{(i)}(\boldsymbol{x}),\ i \in \{1, 2, \ldots, P\}$ related to a continuous scene $f_c(\boldsymbol{u})$ via
  - coordinate transformations $H^{(i)} : \boldsymbol{u} \mapsto \boldsymbol{x}$ (scene/camera motion)
  - spatially varying PSF's $h^{(i)}$ (lens/sensor PSF, defocus, motion blur...)
  - spatial sampling
- Seek discretized approximation of $f_c(\boldsymbol{u})$ on high-resolution sampling lattice
- Using an interpolation kernel h_r we approximate $f_c(\boldsymbol{u})$ as

$$f_c(\boldsymbol{u}) \approx \sum_{\boldsymbol{k} \in \mathbb{Z}^2} f(\boldsymbol{V}\boldsymbol{k}) \cdot h_r\left(\boldsymbol{u} - \boldsymbol{V}\boldsymbol{k}\right) \quad \boldsymbol{k} \in \mathbb{Z}^2$$

$\boldsymbol{V} = \begin{bmatrix} 1/q_x & 0 \\ 0 & 1/q_y \end{bmatrix}$ is a sampling matrix

$q_x, q_y \in \mathbb{N}$ are the horizontal and vertical magnification factors

Discrete-Discrete Multiframe Observation Model

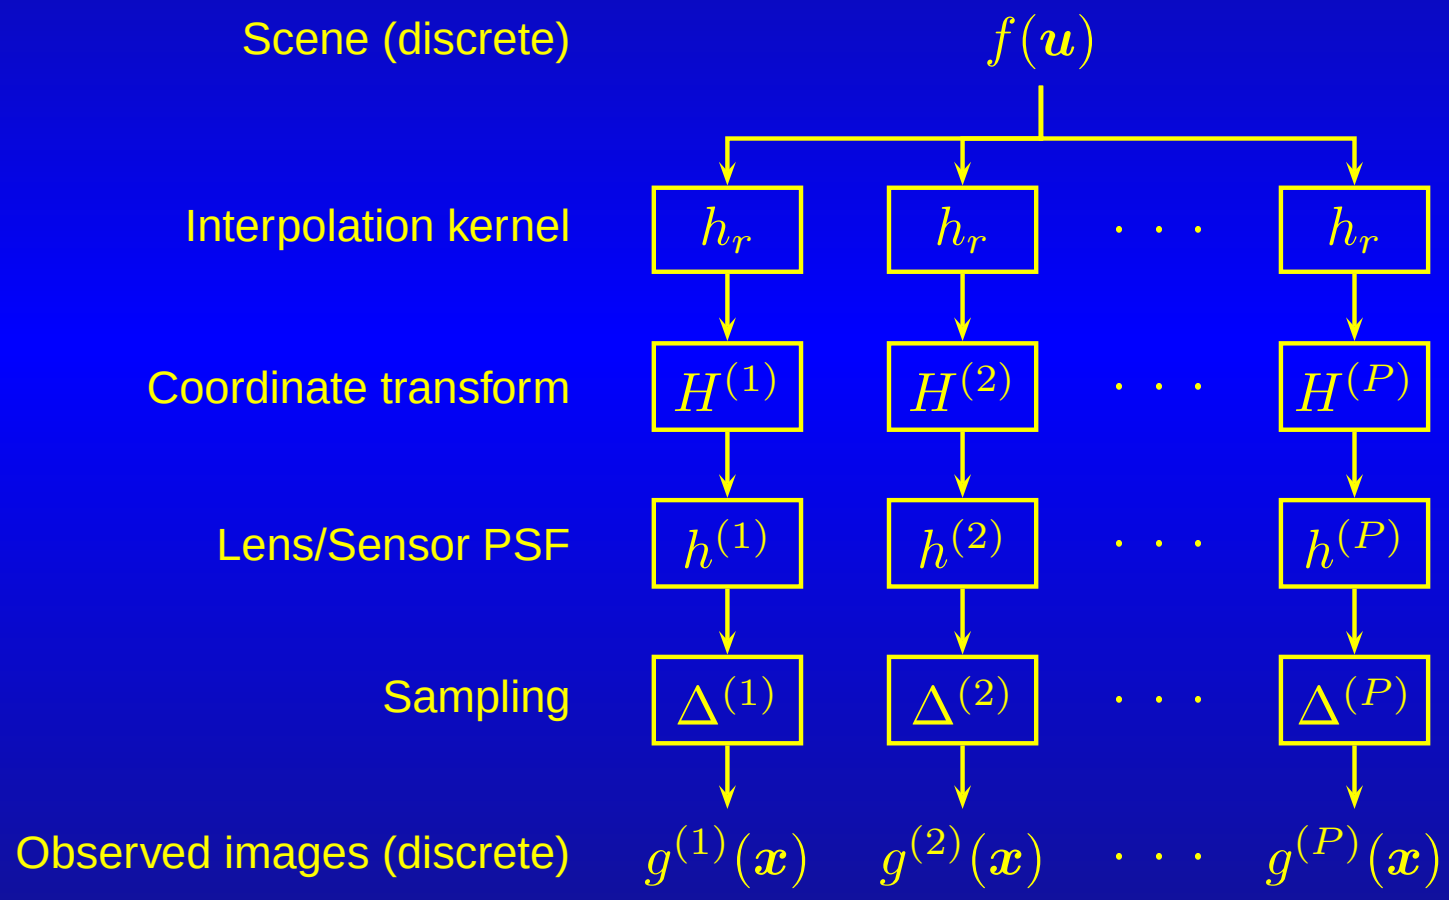

Realizable Discrete-Discrete Multiframe Observation Model

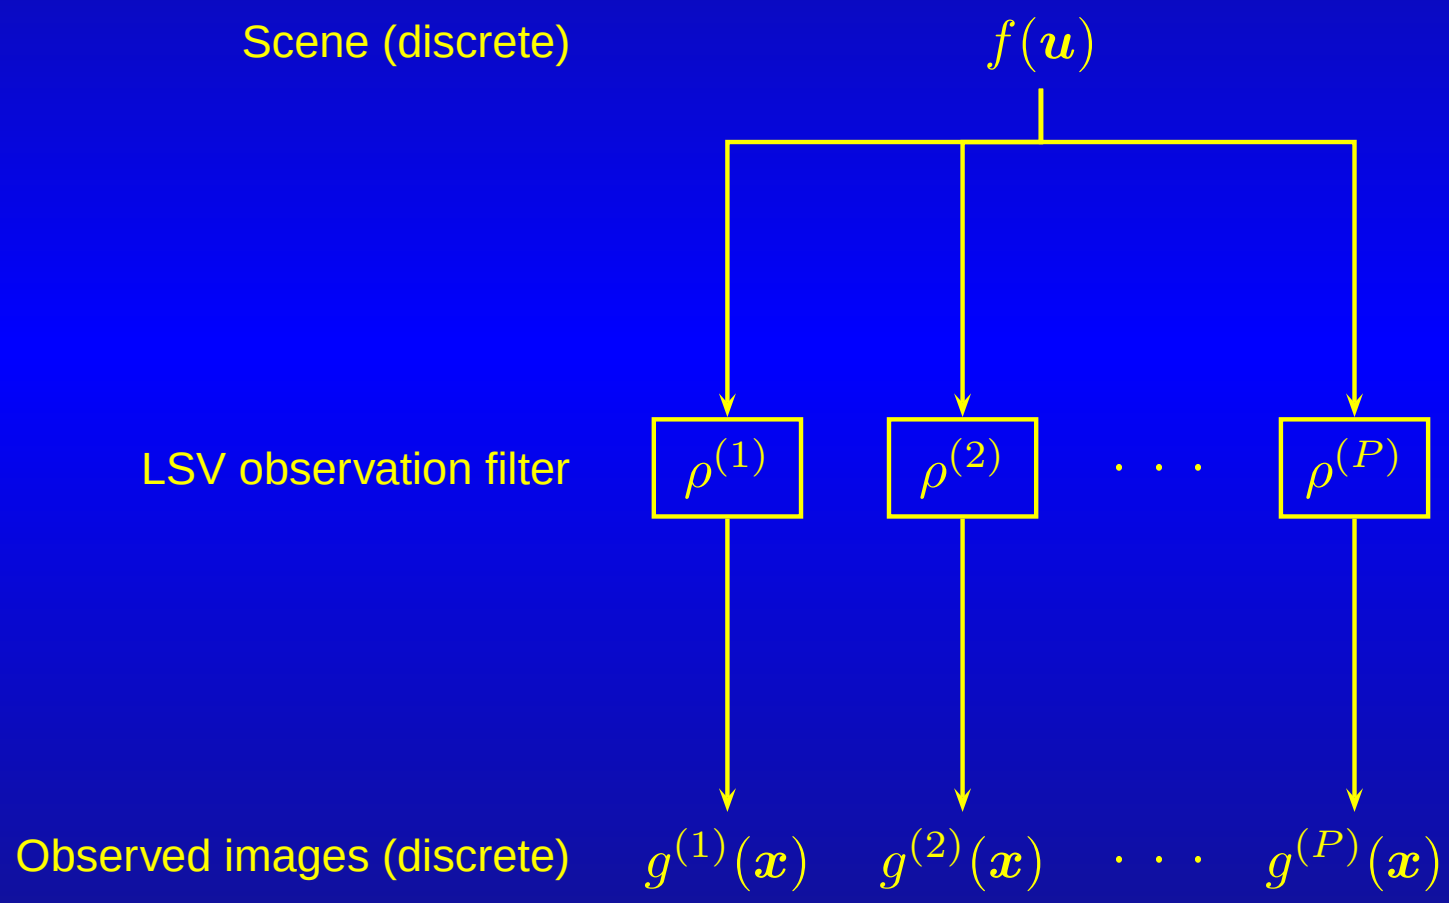

Realizable Discrete-Discrete Multiframe Observation Model

- We can relate $g^{(i)}(\boldsymbol{x})$ to $f(\boldsymbol{k})$ via LSV equations

$$g^{(i)}(\boldsymbol{x}) = \sum_{\boldsymbol{k}\in\mathbb{Z}^2} f(\boldsymbol{V}\boldsymbol{k}) \cdot \rho^{(i)}(\boldsymbol{x},\boldsymbol{k}) \quad \text{for} \quad \boldsymbol{x},\boldsymbol{k} \in \mathbb{Z}^2$$

- $\rho^{(i)}(\boldsymbol{x},\boldsymbol{k})$ are discrete, linear spatially varying (LSV) observation filters

$$\rho^{(i)}(\boldsymbol{x},\boldsymbol{k}) = \int h^{(i)}\left(\boldsymbol{x}, H^{(i)}(\boldsymbol{u})\right) \cdot h_r\left(\boldsymbol{u} - \boldsymbol{V}\boldsymbol{k}\right) \left|\frac{\partial H^{(i)}}{\partial \boldsymbol{u}}\right| \mathrm{d}\boldsymbol{u}$$

- The LSV observation filters . . .
  - depend on each coordinate transforms $H^{(i)}$
  - depend on the interpolation kernel h_r
  - are expressed in terms of the warped PSF's $h^{(i)}$
  - involve integration in the restoration space ($\boldsymbol{u}$)

Determining the Warped Pixel Response

1. Backproject PSF $h(\boldsymbol{x}, \boldsymbol{\alpha})$ from $g(\boldsymbol{x})$ to restored image using H^{-1} (red)
2.
3.
4.

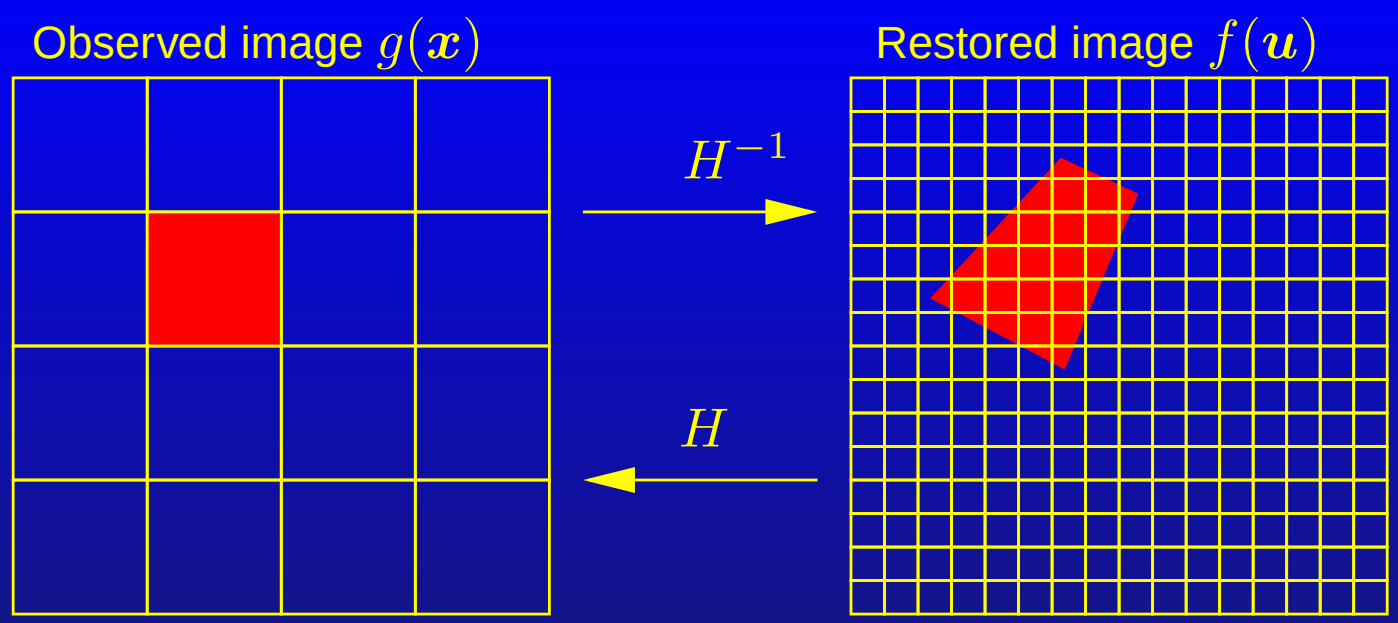

Determining the Warped Pixel Response

1. Backproject PSF $h(\boldsymbol{x}, \boldsymbol{\alpha})$ from $g(\boldsymbol{x})$ to restored image using H^{-1} (red)
2. Determine bounding region for image of $h(\boldsymbol{x}, \boldsymbol{\alpha})$ under backprojection (cyan)
3.
4.

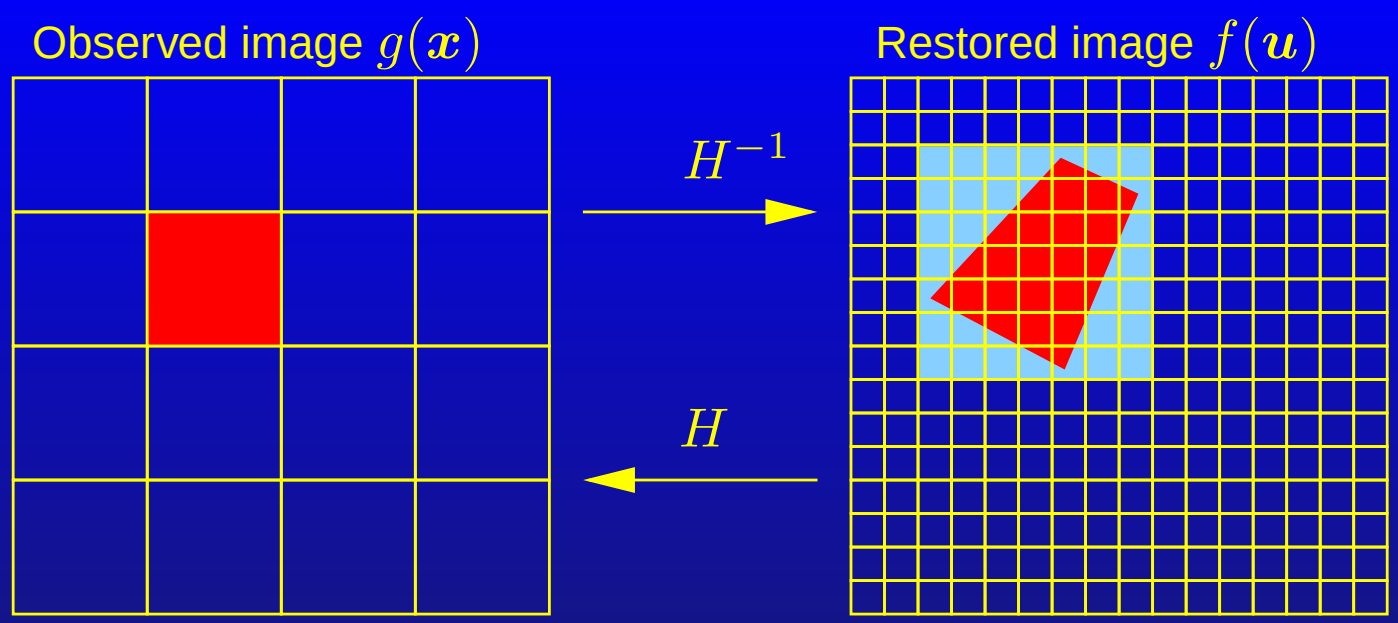

Determining the Warped Pixel Response

1. Backproject PSF $h(\boldsymbol{x}, \boldsymbol{\alpha})$ from $g(\boldsymbol{x})$ to restored image using H^{-1} (red)
2. Determine bounding region for image of $h(\boldsymbol{x}, \boldsymbol{\alpha})$ under backprojection (cyan)
3. $\forall$ S-R pixels $\boldsymbol{u}$ in region, project via H and find $h(\boldsymbol{x}, H(\boldsymbol{u}))$ (green)
4.

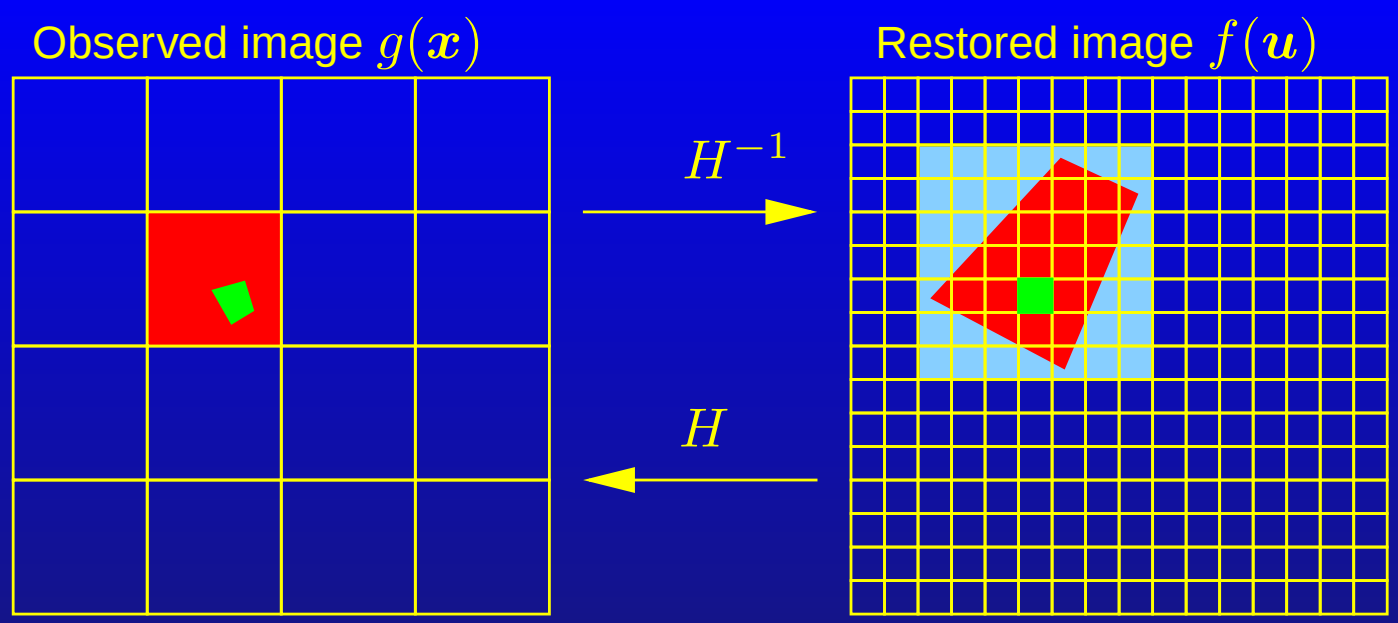

Determining the Warped Pixel Response

1. Backproject PSF $h(\boldsymbol{x}, \boldsymbol{\alpha})$ from $g(\boldsymbol{x})$ to restored image using H^{-1} (red)
2. Determine bounding region for image of $h(\boldsymbol{x}, \boldsymbol{\alpha})$ under backprojection (cyan)
3. $\forall$ S-R pixels $\boldsymbol{u}$ in region, project via H and find $h(\boldsymbol{x}, H(\boldsymbol{u}))$ (green)
4. Scale according to Jacobian and interpolation kernel h_r then integrate over $\boldsymbol{u}$

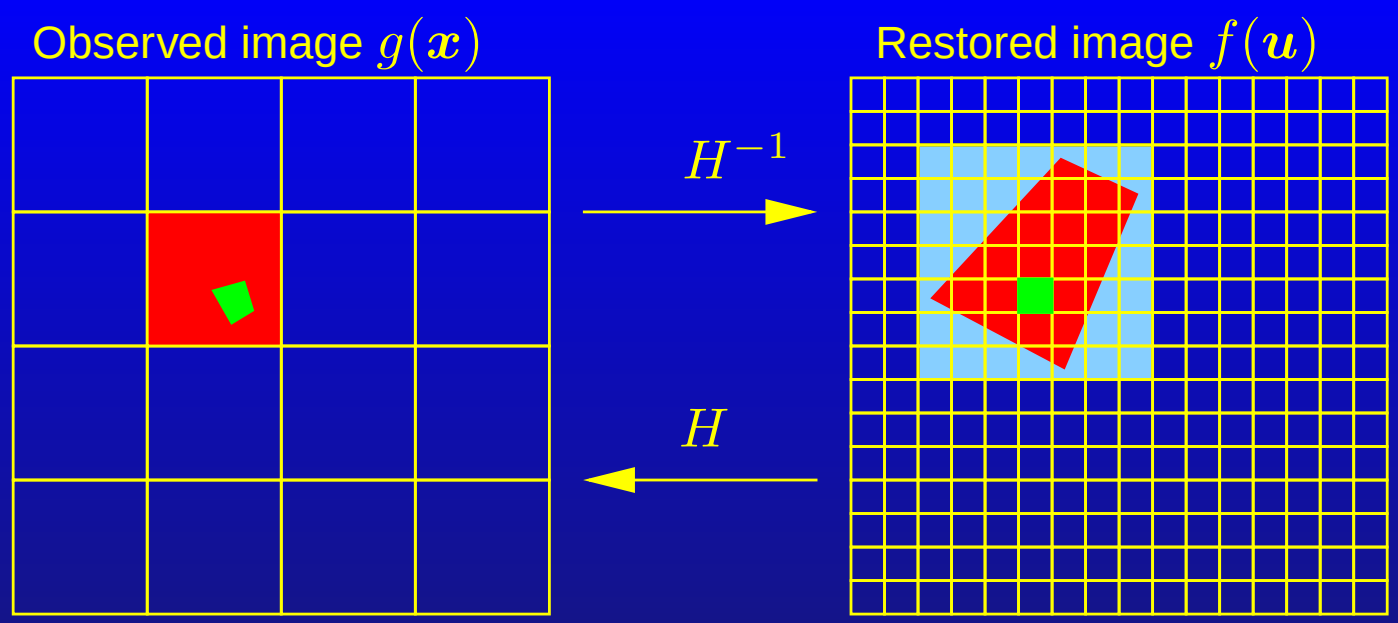

Algorithm to Determine the Observation Filter

for each observed image $\boldsymbol{g}^{(i)}$ $\{$

  for each pixel $\boldsymbol{x}$ $\{$

    back-project the boundary of $h^{(i)}(\boldsymbol{x}, \boldsymbol{\alpha})$ from $g^{(i)}(\boldsymbol{x})$

    to the restored image space using $H^{(i)^{-1}}$

    determine a bounding region for the image of $h(\boldsymbol{x}, \boldsymbol{\alpha})$ under $H^{(i)^{-1}}$

    for each pixel indexed by $\boldsymbol{k}$ in the bounding region $\{$

      set $\rho^{(i)}(\boldsymbol{x}, \boldsymbol{k}) = h^{(i)}\left(\boldsymbol{x}, H^{(i)}(\boldsymbol{u})\right) \cdot \left|\frac{\partial H^{(i)}}{\partial \boldsymbol{u}}\right|$ with $\boldsymbol{u} = \boldsymbol{V}\boldsymbol{k}$

    $\}$

    normalize $\rho^{(i)}(\boldsymbol{x}, \boldsymbol{k})$ so that $\sum_k \rho^{(i)}(\boldsymbol{x}, \boldsymbol{k}) = 1$

  $\}$

$\}$

Example of Observation Filter

- 100% fill-factor pixel
- Diffraction-limited optics
- Projective spatial transformation
- $q_x = q_y = 4$

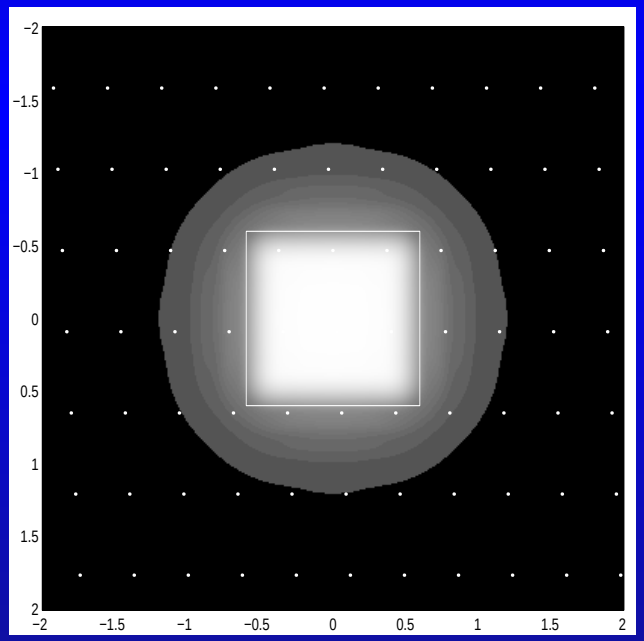

Observed Pixel PSF

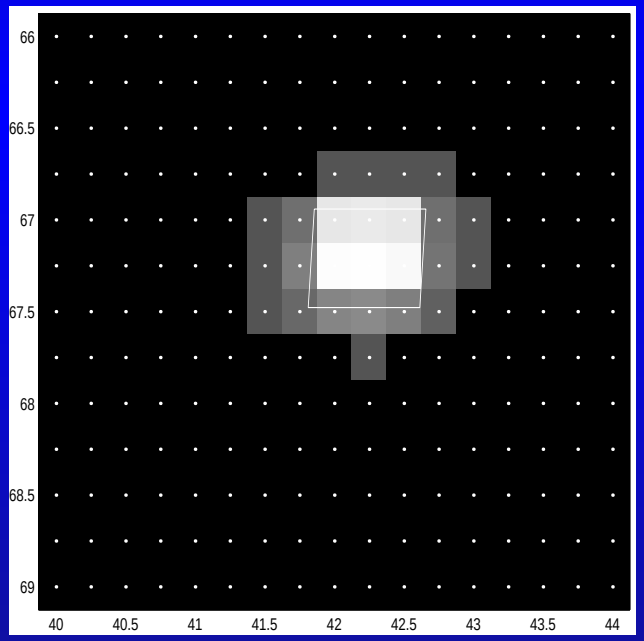

Observation filter $\rho^{(i)}(\boldsymbol{x}, \boldsymbol{k})$

Matrix-Vector Linear Multiframe Observation Model

- Represent images as lexicographically ordered vectors $\boldsymbol{g}^{(i)}$, $\boldsymbol{f}$ (finite case)

  $\Rightarrow$ single pixel observation may then be written as an inner product.

$$g^{(i)}(\boldsymbol{x}) = \sum_{\boldsymbol{k}\in\mathbb{Z}^2} \rho^{(i)}(\boldsymbol{x},\boldsymbol{k}) \cdot f(\boldsymbol{V}\boldsymbol{k}) \quad \text{or equivalently} \quad \boldsymbol{g}_j^{(i)} = \left\langle A_j^{(i)}, \boldsymbol{f} \right\rangle$$

- Stack inner product equations to get single image matrix-vector equation

$$\boldsymbol{g}^{(i)} = \boldsymbol{A}^{(i)}\boldsymbol{f}$$

- Stack matrices $\boldsymbol{A}^{(i)}$ and observations $\boldsymbol{g}^{(i)}$ to get

$$\boldsymbol{g} \doteq \begin{bmatrix} \boldsymbol{g}^{(1)} \\ \boldsymbol{g}^{(2)} \\ \vdots \\ \boldsymbol{g}^{(P)} \end{bmatrix} \quad \text{and} \quad \boldsymbol{A} \doteq \begin{bmatrix} \boldsymbol{A}^{(1)} \\ \boldsymbol{A}^{(2)} \\ \vdots \\ \boldsymbol{A}^{(P)} \end{bmatrix} \quad \text{so that we have} \quad \boldsymbol{g} = \boldsymbol{A}\boldsymbol{f}$$

A Bayesian Framework for Restoration

- Classic linear inverse problem
- Ill-posed, so use regularized solution method with *a-priori* information
- Use augmented observation model which includes noise

$$\boldsymbol{g} = \boldsymbol{A}\boldsymbol{f} + \boldsymbol{n}$$

- $\hat{\boldsymbol{f}}_{\mathrm{MAP}}$ maximizes the *a-posteriori* probability $\mathcal{P}\left(\boldsymbol{f}|\boldsymbol{g}\right)$

$$\begin{aligned}
\hat{\boldsymbol{f}}_{\mathrm{MAP}} &= \arg\max_{\boldsymbol{f}} \left\{\mathcal{P}\left(\boldsymbol{f}|\boldsymbol{g}\right)\right\} \\
&= \arg\max_{\boldsymbol{f}} \left\{\frac{\mathcal{P}\left(\boldsymbol{g}|\boldsymbol{f}\right)\mathcal{P}\left(\boldsymbol{f}\right)}{\mathcal{P}\left(\boldsymbol{g}\right)}\right\} \\
&= \arg\max_{\boldsymbol{f}} \left\{\log\mathcal{P}\left(\boldsymbol{g}|\boldsymbol{f}\right) + \log\mathcal{P}\left(\boldsymbol{f}\right)\right\}
\end{aligned}$$

A Bayesian Framework for Restoration

- Likelihood term: Assume noise is zero-mean Gaussian

$$\begin{aligned}\mathcal{P}(\boldsymbol{g}|\boldsymbol{f}) &= \mathcal{P}_{\boldsymbol{N}}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f}) \\ &\propto \exp\left\{-\frac{1}{2}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f})^T\boldsymbol{K}^{-1}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f})\right\}\end{aligned}$$

- Prior term: Markov random field (Gibbs density)

$$\mathcal{P}(\boldsymbol{f}) \propto \exp\left\{-\frac{1}{\beta}\sum_{c\in\mathcal{C}}\rho_T(\boldsymbol{\partial}_c\boldsymbol{f})\right\}$$

  - Huber penalty function $\rho_T(x)$ (edge preserving)
  - Local interactions $\boldsymbol{\partial}_c$ approximate 2nd spatial derivates in 4 orientations

A Bayesian Framework for Restoration

- Combined objective function

$$\begin{aligned}\hat{\boldsymbol{f}}_{\mathrm{MAP}} &= \arg\max_{\boldsymbol{f}} \left\{ -\frac{1}{2}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f})^T \boldsymbol{K}^{-1}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f}) - \frac{1}{\beta}\sum_{c\in\mathcal{C}} \rho_T\left(\boldsymbol{\partial}_c \boldsymbol{f}\right) \right\} \\ &= \arg\min_{\boldsymbol{f}} \left\{ \frac{1}{2}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f})^T \boldsymbol{K}^{-1}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f}) + \gamma\sum_{c\in\mathcal{C}} \rho_T\left(\boldsymbol{\partial}_c \boldsymbol{f}\right) \right\}\end{aligned}$$

- Use your favorite optimization technique to find $\hat{\boldsymbol{f}}_{\mathrm{MAP}}$
- Unique solution under very mild conditions

Example

- Simulated imaging environment

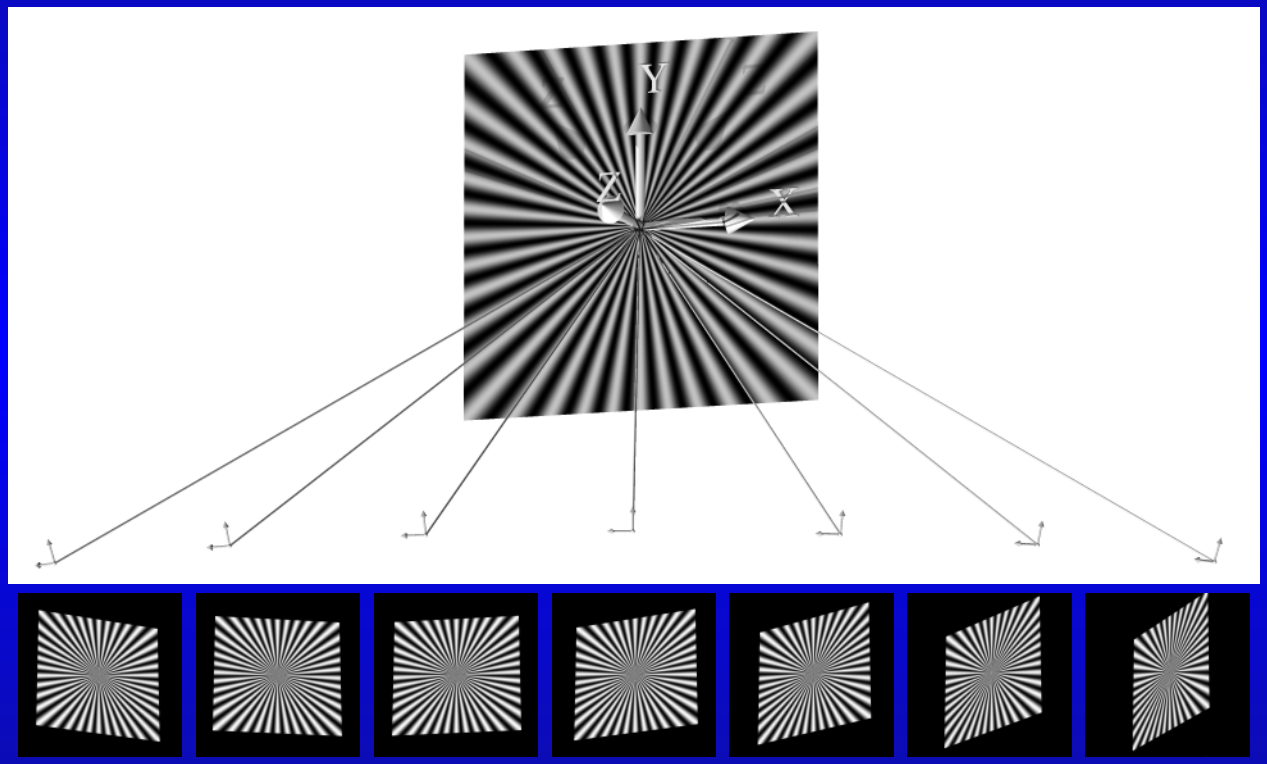

Example

- Superresolution Restoration

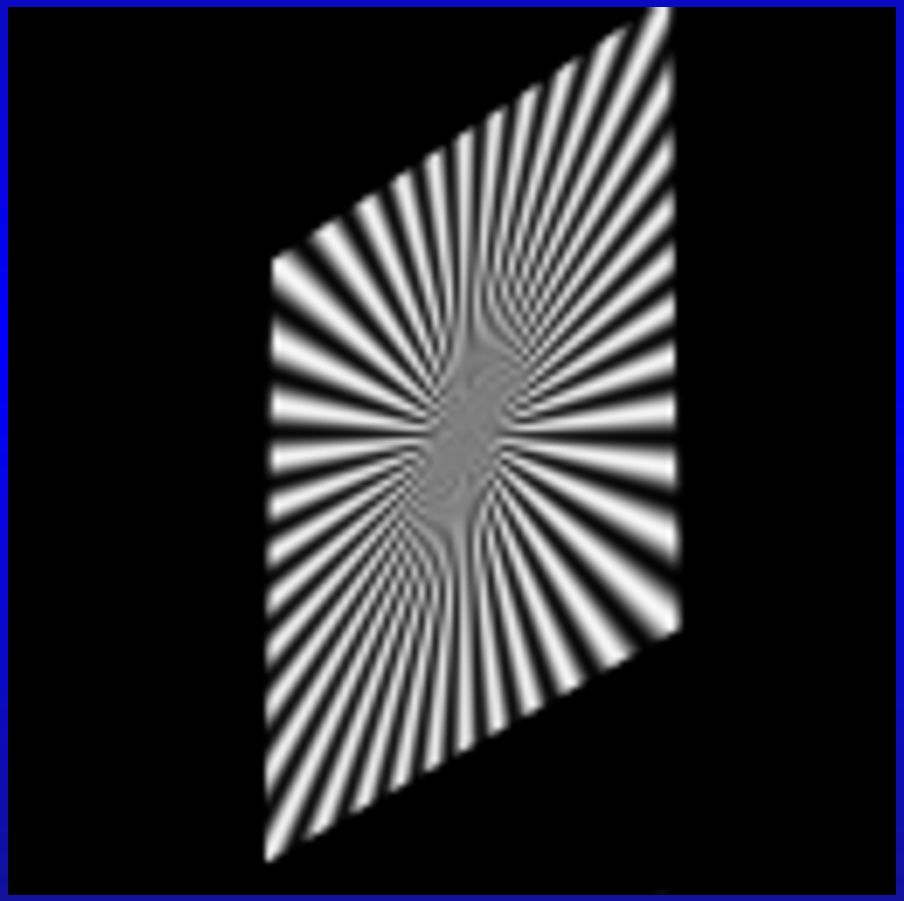

Cubic spline interpolation

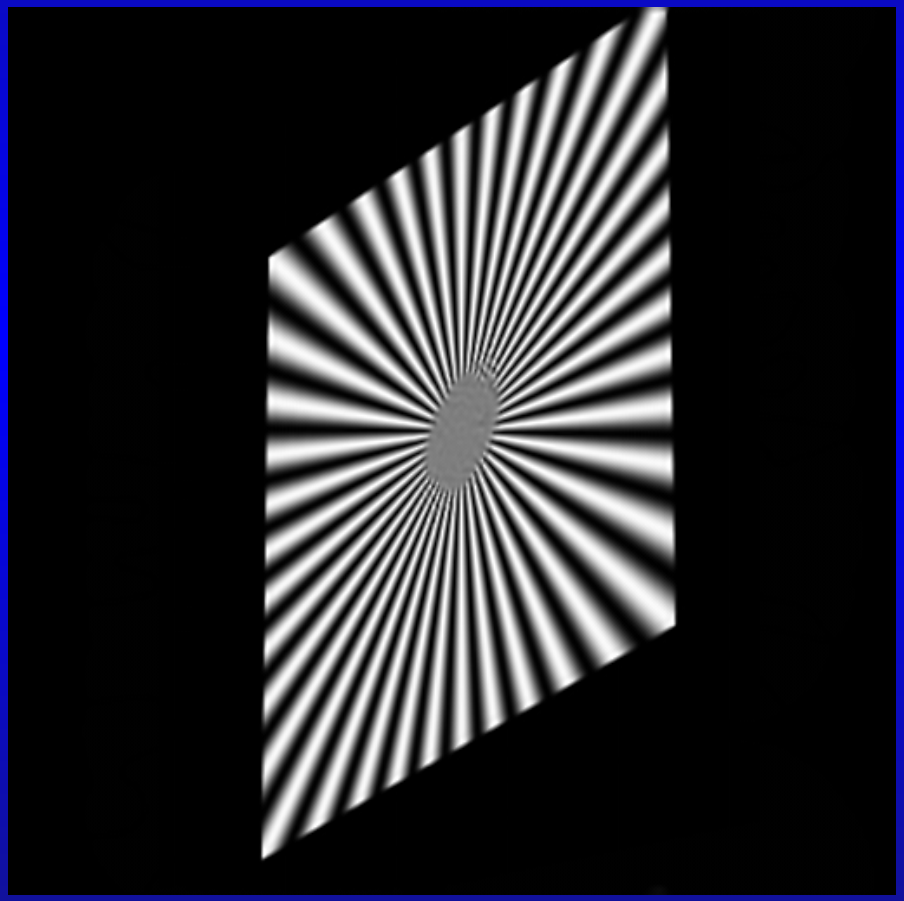

Multiframe restoration

Summary

- Generalized linear observation model used in multiframe superresolution restoration
  - Nonaffine image registration
  - Easy to accommodate spatially-varying PSFs
- Algorithm to find linear, spatially varying observation filter
- Leads to sparse observation matrix (construct only once)
- Well-suited to iterative restoration methods
- No changes to restoration framework necessary
- Demonstrate application

Roadmap

- Overview of the talk
- Background [Chap. 2]
- Simultaneous multiframe superresolution restoration [Chap. 7]
- Linear models for multiframe superresolution restoration [Chap. 8]
- Conclusions [Chap. 9]

Main Research Contributions

- Evaluating the performance of motion estimators
  - Provides a model for a general evaluation framework
- Proved equivalence of descriptions of pinhole camera models in two fields
- Introduced new simultaneous multiframe superresolution restoration framework
  - Spatiotemporal constraints
  - Motion estimation confidence parameters
- Generalized linear observation model to nonaffine motion and SVPSF
  - Almost arbitrary motion models and SVPSFs can be accommodated
  - No change needed in solution framework

Additional Contributions

- General background
  - multiframe superresolution restoration
  - ill-posed inverse problems and their solution
- Exhaustive survey and categorization of the literature
- Motion estimation (critical component)

We shall not cease from exploration
And the end of all our exploring
Will be to arrive where we started
and know the place for the first time.

T. S. Elliot

Roadmap

- Overview of the dissertation and talk [Chap. 1]
- Brief overview of the introductory material
  - Background [Chap. 2]
  - Review of multiframe superresolution restoration methods [Chap. 3]
  - Motion estimation [Chap. 4]
- Short discussion on the performance of block motion estimation methods [Chap. 5]
- Concentrate on selected topics
  - Raytracing and the camera matrix [Chap. 6]
  - Simultaneous multiframe superresolution restoration [Chap. 7]
  - Linear models for multiframe superresolution restoration [Chap. 8]
- Conclusions and directions for future work [Chap. 9]

Review of Multiframe Superresolution Restoration Techniques

- Exhaustive survey of the literature
- Categorize work according to methods used
- Two main classes:
  1. Frequency domain
  2. Spatial domain
- Early work used frequency domain approach
- All modern methods are spatial domain

Spatial and Frequency-Domain Methods

	Frequency-Domain	**Spatial-Domain**
Observation model	Frequency-domain	Spatial-domain
Motion models	Global translation	Almost unlimited
Degradation model	Limited	LSI or LSV
Noise model	Limited	Very flexible Even spatially varying

Spatial-Domain and Frequency-Domain Methods

	Frequency-Domain	**Spatial-Domain**
SR mechanism	Dealiasing	Dealiasing BW extrapolation with *a-priori* constraints
Simplicity	Very simple	Generally complex
Computational cost	Low	High
A-priori constraints	Limited	Almost unlimited
Regularization	Limited	Excellent
Extensibility	Poor	Excellent
Performance	Good for specific applications	Good

MAP vs POCS Restoration Methods

	Bayesian (MAP)	**POCS**
Theoretical framework	Rich	Limited
A-priori constraints	Prior probability density	Convex sets
	Projected constraints possible	Project onto constraints
SR solution	Unique	Nonunique
	MAP estimate	Volume of $\bigcap$ of sets
Optimization	Iterative	Iterative
	Various descent methods	Projection methods
	Good convergence	Often slow convergence
Computational cost	High	High
Complications	Optimization difficult for nonconvex priors	Projection operators potentially problematic

Roadmap

- Overview of the dissertation and talk [Chap. 1]
- Brief overview of the introductory material
  - Background [Chap. 2]
  - Review of multiframe superresolution restoration methods [Chap. 3]
  - Motion estimation [Chap. 4]
- Short discussion on the performance of block motion estimation methods [Chap. 5]
- Concentrate on selected topics
  - Raytracing and the camera matrix [Chap. 6]
  - Simultaneous multiframe superresolution restoration [Chap. 7]
  - Linear models for multiframe superresolution restoration [Chap. 8]
- Conclusions and directions for future work [Chap. 9]

Motion Estimation

- Central to multiframe restoration — register images to common reference
- Difficult, ill-posed problem
- Review main issues
  - basics of 3-D motion, projected motion, optical flow
  - the aperture problem
  - aggregation of motion constraints
  - brightness constancy
  - the optical flow equation
  - motion field estimation and segmentation
- Discuss 3 representative techniques
- All applicable to the multiframe restoration problem

Roadmap

- Overview of the dissertation and talk [Chap. 1]
- Brief overview of the introductory material
  - Background [Chap. 2]
  - Review of multiframe superresolution restoration methods [Chap. 3]
  - Motion estimation [Chap. 4]
- Short discussion on the performance of block motion estimation methods [Chap. 5]
- Concentrate on selected topics
  - Raytracing and the camera matrix [Chap. 6]
  - Simultaneous multiframe superresolution restoration [Chap. 7]
  - Linear models for multiframe superresolution restoration [Chap. 8]
- Conclusions and directions for future work [Chap. 9]

Introduction

- PINHOLE CAMERA models used in:
  1. Computer graphics (CG) raytracing
  2. Computer vision (CV)
- Same underlying model, different purposes
  1. Cast rays from eye into scene — Backprojection (CG)
  2. Project scene onto focal plane — Projection (CV)
- Very different mathematical representations
- Explore the relationship between them

The Pinhole Camera

- Image formation in a pinhole camera

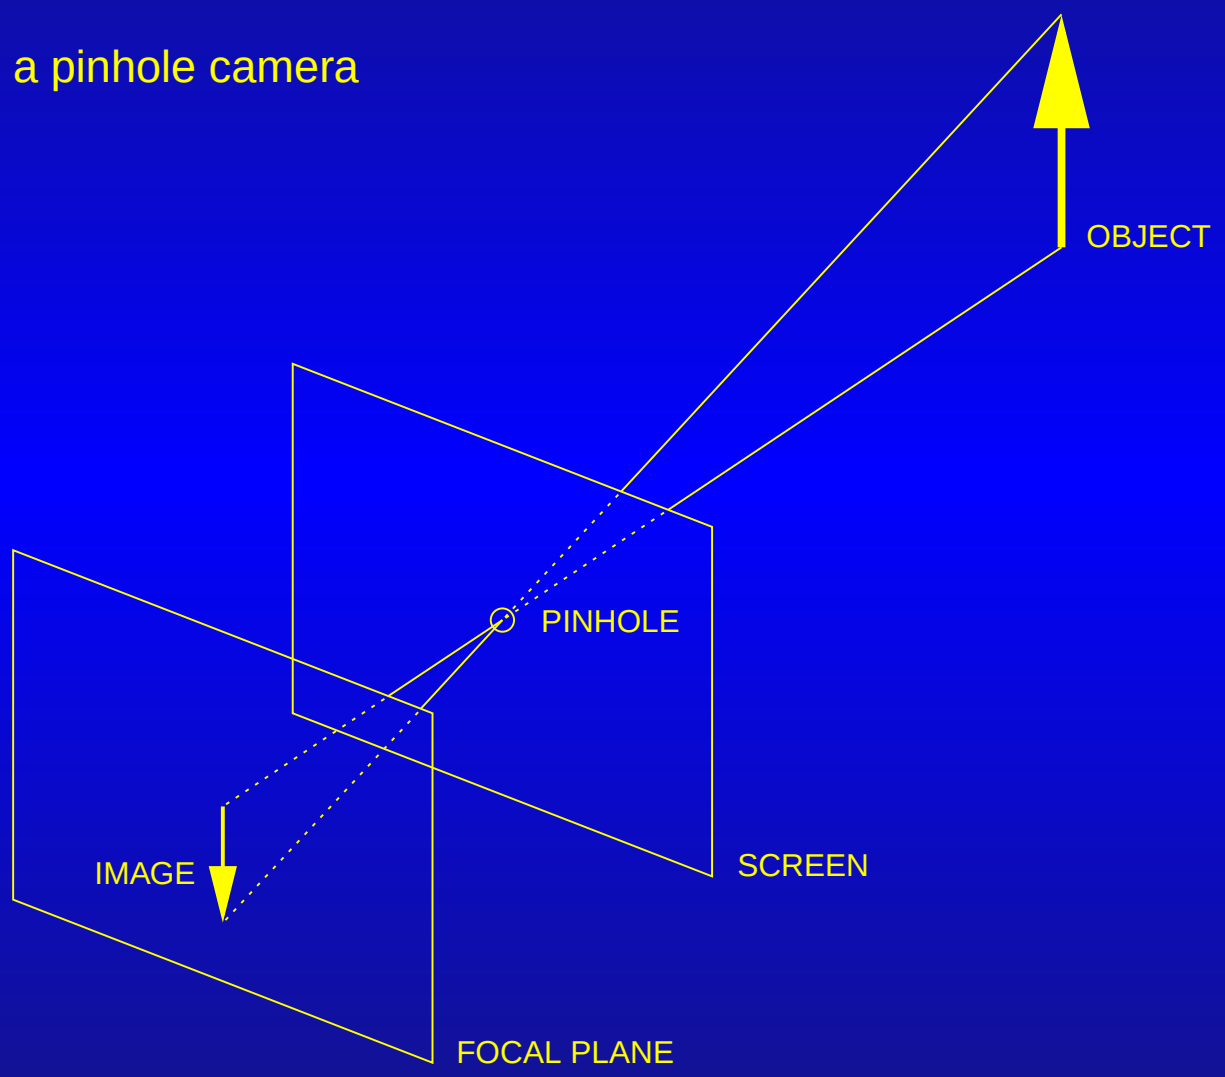

The Pinhole Camera Redux

- Pinhole camera model with virtual focal plane and camera center

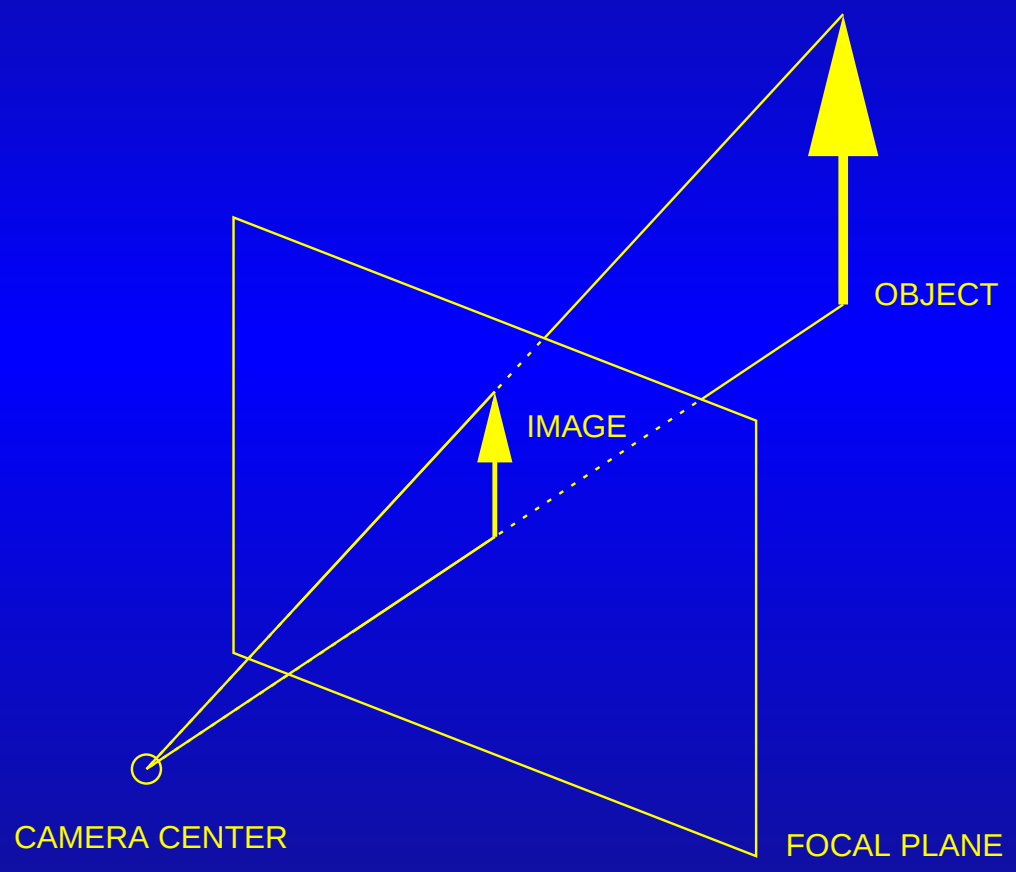

Raytracing Camera Model

- Cast rays from the camera center $\boldsymbol{c}$ . . . through focal plane . . . into 3-D scene
- Camera may be positioned and oriented arbitrarily
- Ray origin $\boldsymbol{c}$ and direction $\boldsymbol{d}$

$$
\begin{aligned}
\boldsymbol{r} &= \boldsymbol{c} + \gamma \boldsymbol{d} \\
&= \boldsymbol{c} + \gamma \left[u\hat{\boldsymbol{u}} + v\hat{\boldsymbol{v}} + f\hat{\boldsymbol{p}} \right] \\
&= \boldsymbol{c} + \alpha \left[\frac{u}{f}\hat{\boldsymbol{u}} + \frac{v}{f}\hat{\boldsymbol{v}} + \hat{\boldsymbol{p}} \right]
\end{aligned}
$$

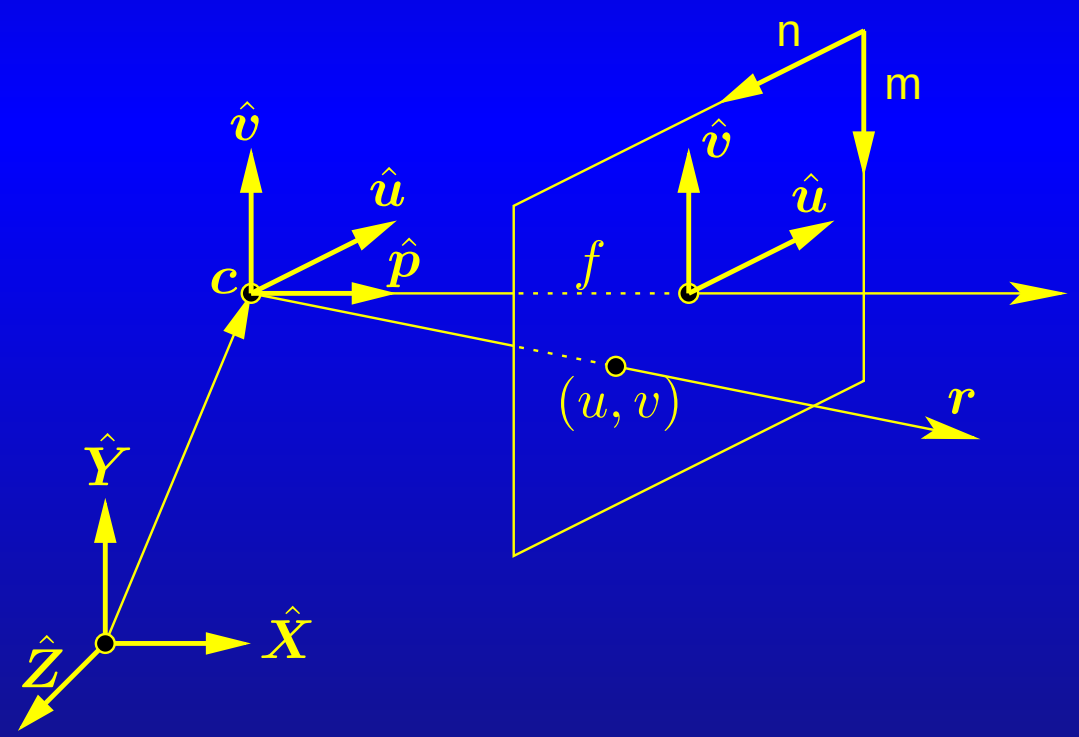

Raytracing Camera Model

- Focal plane array geometry

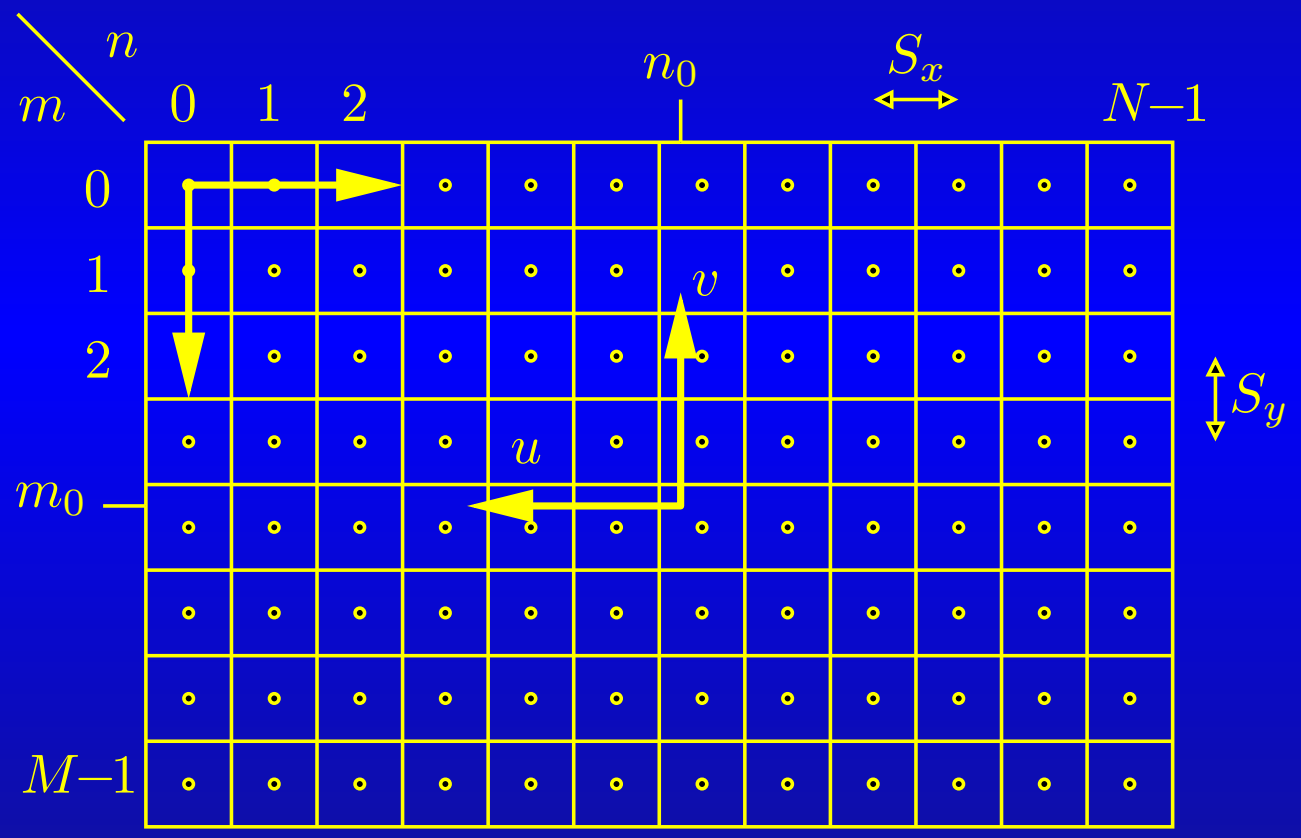

Raytracing Camera Model with FPA

- Transformation between u, v and m, n is given by:

$$u = (n_0 - n)S_x \quad \text{and} \quad v = (m_0 - m)S_y,$$

- So ray is then:

$$\begin{aligned}
\boldsymbol{r} &= \boldsymbol{c} + \alpha \left[\frac{u}{f}\hat{\boldsymbol{u}} + \frac{v}{f}\hat{\boldsymbol{v}} + \hat{\boldsymbol{p}} \right] \\
&= \boldsymbol{c} + \alpha \left[\frac{n_0 - n}{f} S_x \hat{\boldsymbol{u}} + \frac{m_0 - m}{f} S_y \hat{\boldsymbol{v}} + \hat{\boldsymbol{p}} \right] \\
&= \boldsymbol{c} + \alpha \left[-n\frac{S_x}{f}\hat{\boldsymbol{u}} - m\frac{S_y}{f}\hat{\boldsymbol{v}} + \left\{ n_0 \frac{S_x}{f}\hat{\boldsymbol{u}} + m_0 \frac{S_y}{f}\hat{\boldsymbol{v}} + \hat{\boldsymbol{p}} \right\} \right]
\end{aligned}$$

- Very efficiently computed using increments $\frac{S_x}{f}\hat{\boldsymbol{u}}$ and $\frac{S_y}{f}\hat{\boldsymbol{v}}$

Raytracing Equation in Matrix Form

- For given α, position of ray in world-coordinates is $\boldsymbol{q} = [q_x \; q_y \; q_z]^\top$

$$\begin{bmatrix} q_x \\ q_y \\ q_z \end{bmatrix} = \begin{bmatrix} c_x \\ c_y \\ c_z \end{bmatrix} + \alpha \begin{bmatrix} | & | & | \\ \hat{\boldsymbol{u}} & \hat{\boldsymbol{v}} & \hat{\boldsymbol{p}} \\ | & | & | \end{bmatrix} \begin{bmatrix} u/f \\ v/f \\ 1 \end{bmatrix}$$

- In focal plane coordinates:

$$\begin{bmatrix} q_x \\ q_y \\ q_z \end{bmatrix} = \begin{bmatrix} c_x \\ c_y \\ c_z \end{bmatrix} + \begin{bmatrix} | & | & | \\ \hat{\boldsymbol{u}} & \hat{\boldsymbol{v}} & \hat{\boldsymbol{p}} \\ | & | & | \end{bmatrix} \begin{bmatrix} -S_x/f & 0 & n_0 S_x/f \\ 0 & -S_y/f & m_0 S_y/f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \alpha\, n \\ \alpha\, m \\ \alpha \end{bmatrix}$$

Computer Vision Camera Matrix Model

- Camera in standard position

$$\frac{q_x}{q_z} = \frac{u}{f}$$
$$\frac{q_y}{q_z} = \frac{v}{f}$$

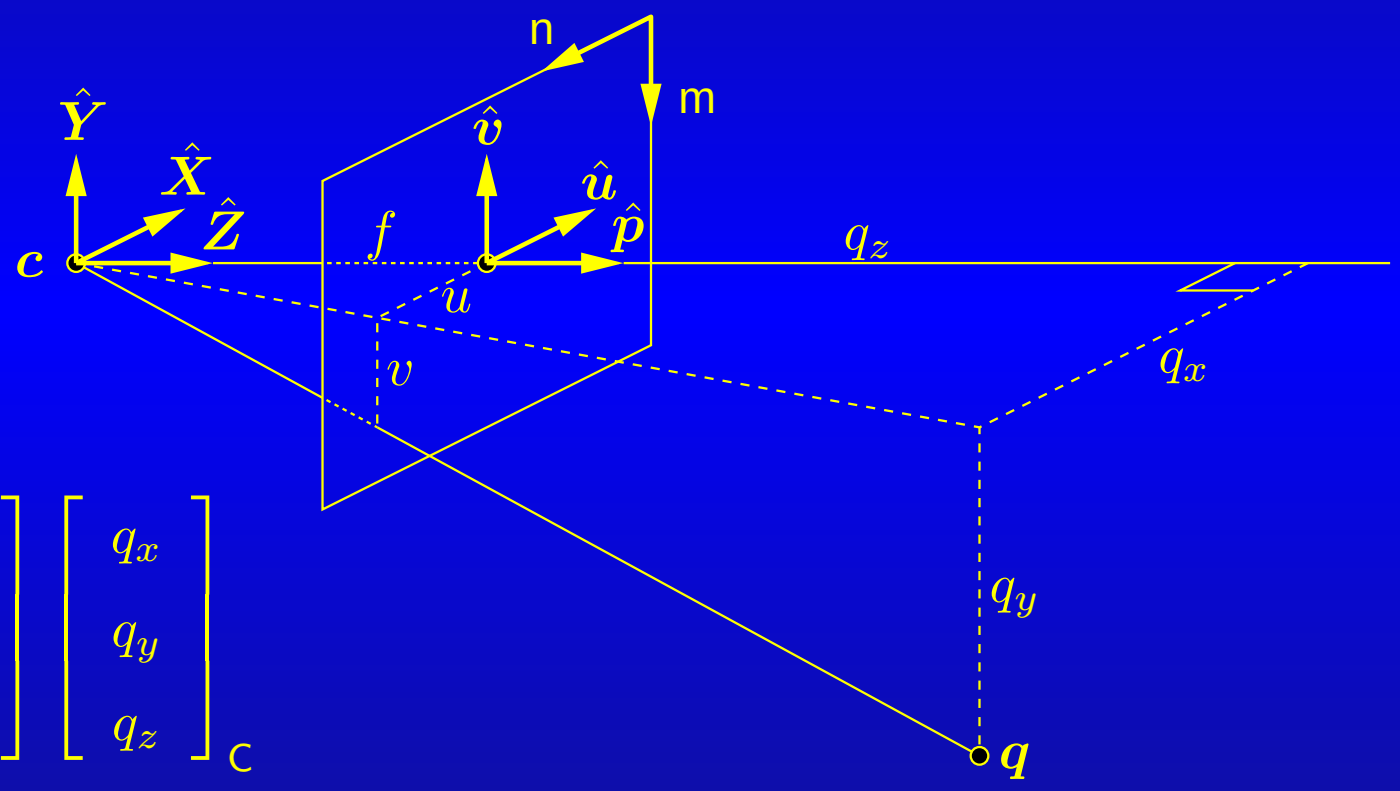

$$\begin{bmatrix} q_z u \\ q_z v \\ q_z \end{bmatrix} = \begin{bmatrix} f & 0 & 0 \\ 0 & f & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} q_x \\ q_y \\ q_z \end{bmatrix}_{\mathrm{C}}$$

Computer Vision Camera Matrix Model

- Using homogeneous coordinates:

$$\begin{bmatrix} q_z u \\ q_z v \\ q_z \end{bmatrix} = \begin{bmatrix} f & 0 & 0 \\ 0 & f & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} q_x \\ q_y \\ q_z \end{bmatrix}_{\mathsf{C}}$$

- Consider focal plane array coordinate transformation

$$\begin{bmatrix} q_z n \\ q_z m \\ q_z \end{bmatrix} = \begin{bmatrix} -f/S_x & 0 & n_0 \\ 0 & -f/S_y & m_0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} q_x \\ q_y \\ q_z \end{bmatrix}_{\mathsf{C}}$$

Computer Vision Camera Matrix Model

- Generalize to arbitrary camera position $\boldsymbol{c}$ and orientation $\boldsymbol{R} \triangleq [\, \hat{\boldsymbol{u}}\ \hat{\boldsymbol{v}}\ \hat{\boldsymbol{p}} \,]$
- Transformation of a point between world and camera coordinates:

$$\boldsymbol{q}_{\mathsf{C}} = \boldsymbol{R}^{-1}(\boldsymbol{q}_{\mathsf{W}} - \boldsymbol{c})$$

- Using homogeneous coordinates

$$\begin{bmatrix} q_x \\ q_y \\ q_z \end{bmatrix}_{\mathsf{C}} = \begin{bmatrix} - \hat{\boldsymbol{u}}^T - \\ - \hat{\boldsymbol{v}}^T - \\ - \hat{\boldsymbol{p}}^T - \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & -c_x \\ 0 & 1 & 0 & -c_y \\ 0 & 0 & 1 & -c_z \end{bmatrix} \begin{bmatrix} q_x \\ q_y \\ q_z \\ 1 \end{bmatrix}_{\mathsf{W}}$$

Computer Vision Camera Matrix Model

- Complete model is:

$$\begin{bmatrix} q_z n \\ q_z m \\ q_z \end{bmatrix} = \begin{bmatrix} -f/S_x & 0 & n_0 \\ 0 & -f/S_y & m_0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} - \hat{\boldsymbol{u}}^T - \\ - \hat{\boldsymbol{v}}^T - \\ - \hat{\boldsymbol{p}}^T - \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & -c_x \\ 0 & 1 & 0 & -c_y \\ 0 & 0 & 1 & -c_z \end{bmatrix} \begin{bmatrix} q_x \\ q_y \\ q_z \\ 1 \end{bmatrix}_{\mathrm{W}}$$

- Canonical camera matrix decomposition:

$$\boldsymbol{P} = \boldsymbol{K}\boldsymbol{R}^{-1} \left[\mathbf{I} \,|\, -\boldsymbol{c} \right]$$

Relating the Raytracing and Camera Matrix Representations

- Raytracing:

$$\begin{bmatrix} q_x \\ q_y \\ q_z \end{bmatrix} = \begin{bmatrix} c_x \\ c_y \\ c_z \end{bmatrix} + \begin{bmatrix} | & | & | \\ \hat{\boldsymbol{u}} & \hat{\boldsymbol{v}} & \hat{\boldsymbol{p}} \\ | & | & | \end{bmatrix} \begin{bmatrix} -S_x/f & 0 & n_0 S_x/f \\ 0 & -S_y/f & m_0 S_y/f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \alpha\, n \\ \alpha\, m \\ \alpha \end{bmatrix}$$

- Camera Matrix:

$$\begin{bmatrix} q_z n \\ q_z m \\ q_z \end{bmatrix} = \begin{bmatrix} -f/S_x & 0 & n_0 \\ 0 & -f/S_y & m_0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} - \hat{\boldsymbol{u}}^T - \\ - \hat{\boldsymbol{v}}^T - \\ - \hat{\boldsymbol{p}}^T - \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & -c_x \\ 0 & 1 & 0 & -c_y \\ 0 & 0 & 1 & -c_z \end{bmatrix} \begin{bmatrix} q_x \\ q_y \\ q_z \\ 1 \end{bmatrix}_{\mathrm{W}}$$

- Exact inverses of each other
- Equivalent decompositions, same underlying representation

Roadmap

- Overview of the dissertation and talk [Chap. 1]
- Brief overview of the introductory material
  - Background [Chap. 2]
  - Review of multiframe superresolution restoration methods [Chap. 3]
  - Motion estimation [Chap. 4]
- Short discussion on the performance of block motion estimation methods [Chap. 5]
- Concentrate on selected topics
  - Raytracing and the camera matrix [Chap. 6]
  - Simultaneous multiframe superresolution restoration [Chap. 7]
  - Linear models for multiframe superresolution restoration [Chap. 8]
- Conclusions and directions for future work [Chap. 9]

Objectives and Method

- Empirical performance evaluation of a block-matching, subpixel motion estimator
- Need a corpus of data with known motion (difficult)
- Solution:
  - Simulate a realistic imaging system
  - Use synthetic scene with known motion
  - Estimate motion, analyze errors
- Try to determine realistic best-case performance
- Useful general framework in context of superresolution restoration

Results

- $p \in \mathbb{N}$: The interpolation factor
- α : The % of errors within $\pm\frac{1}{2p}$ pixels of true motion
- Fix α (80%)
  - Compare actual performance with theoretical bounds
  - $\Rightarrow$ diminishing returns for increasing p
- Vary α
  - Find highest p yielding α-% of errors within $\pm\frac{1}{2p}$ pixels of true motion
  - $\Rightarrow$ accuracy / resolution trade-off

Results

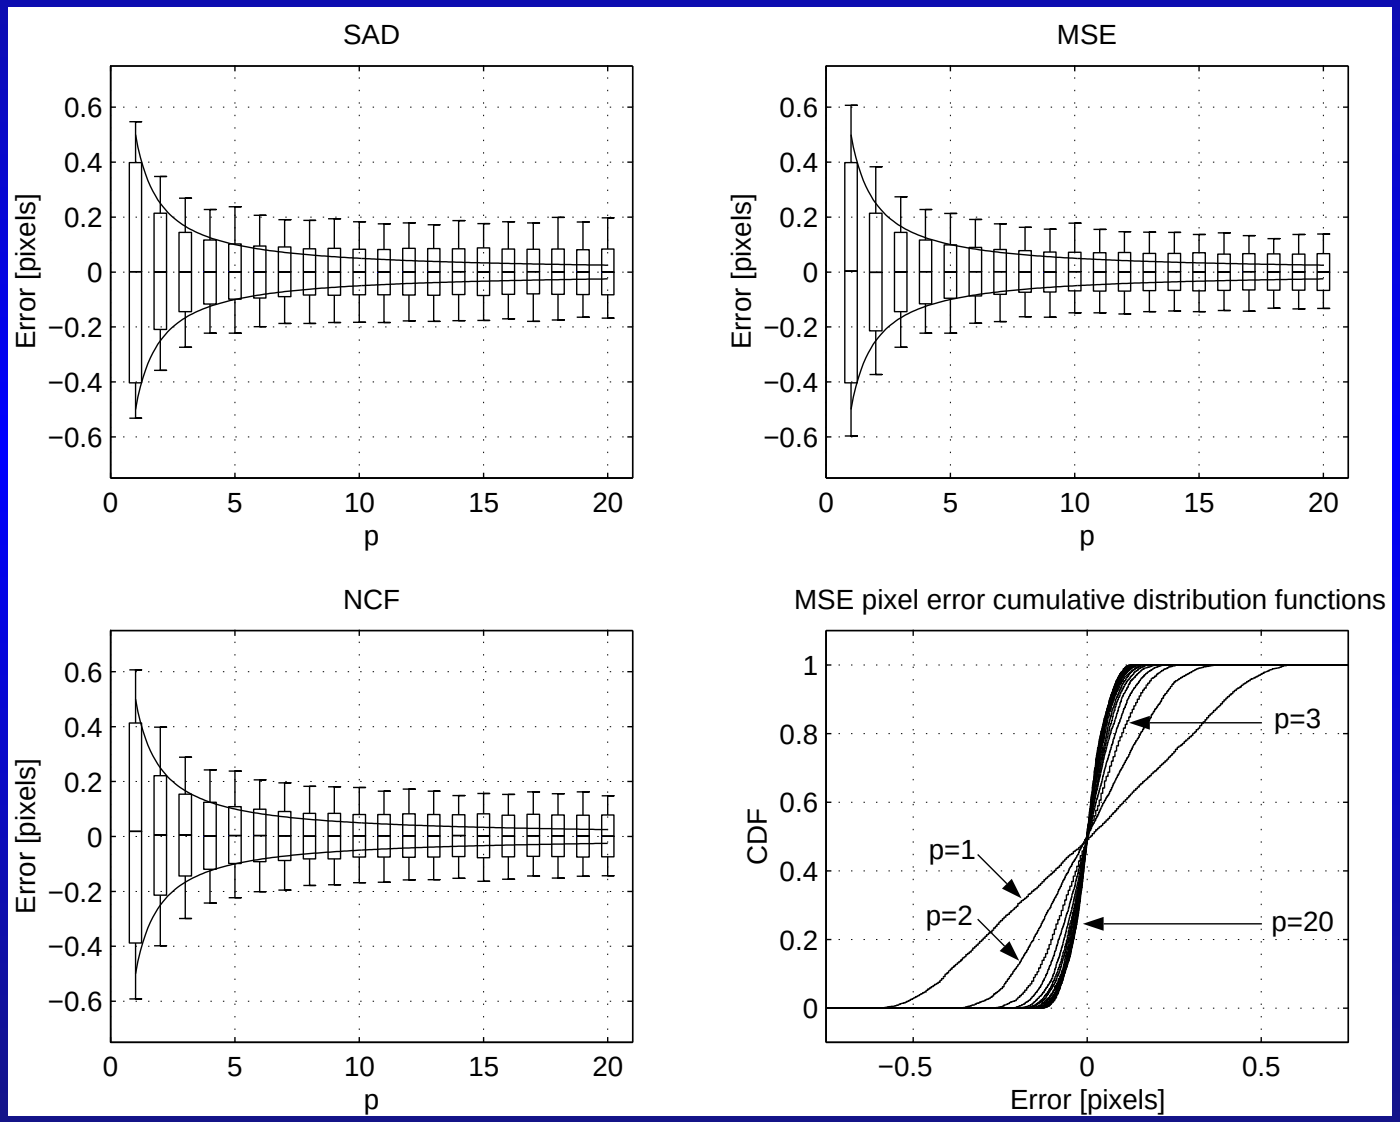
SAD
MSE
NCF
MSE pixel error cumulative distribution functions
Error [pixels]
p
CDF
Error [pixels]
p=1
p=2
p=3
p=20

Results

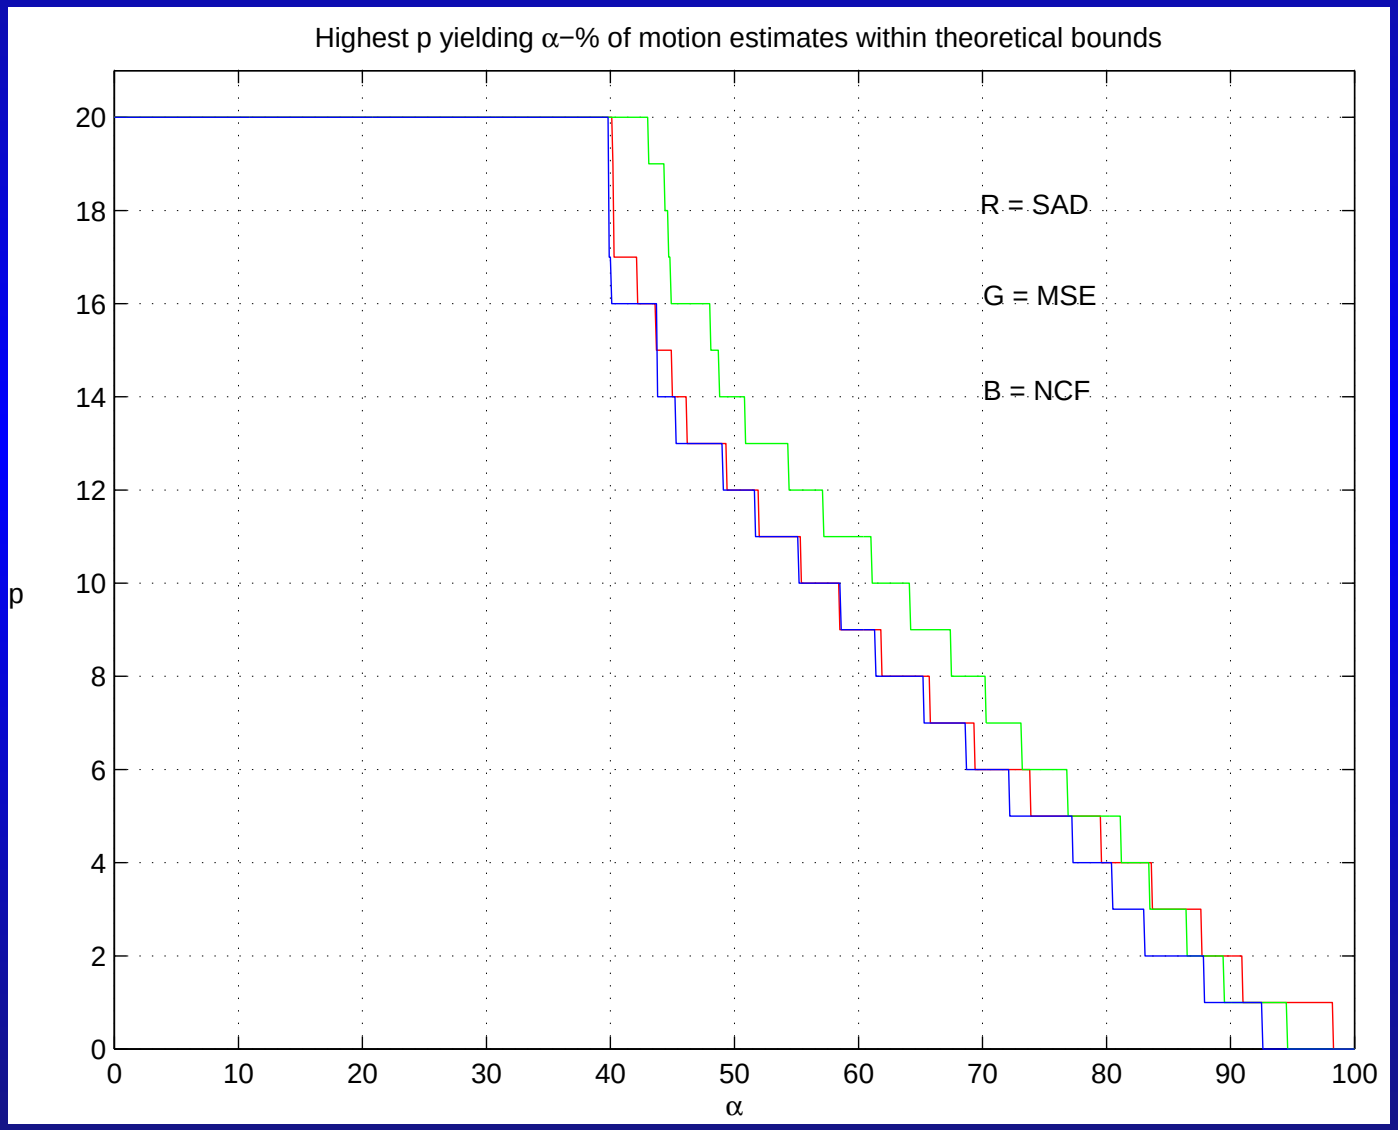
Highest p yielding α−% of motion estimates within theoretical bounds
R = SAD
G = MSE
B = NCF
p
α
0
2
4
6
8
10
12
14
16
18
20
0
10
20
30
40
50
60
70
80
90
100

Temporally Coincident Observation

y_1	y_2	y_3
y_4	y_5	y_6
y_7	y_8	y_9

z_1	z_2	z_3	z_4	z_5	z_6
z_7	z_8	z_9	z_{10}	z_{11}	z_{12}
z_{13}	z_{14}	z_{15}	z_{16}	z_{17}	z_{18}
z_{19}	z_{20}	z_{21}	z_{22}	z_{23}	z_{24}
z_{25}	z_{26}	z_{27}	z_{28}	z_{29}	z_{30}
z_{31}	z_{32}	z_{33}	z_{34}	z_{35}	z_{36}

Observed image : $y_5 = \frac{1}{4}(z_{15} + z_{16} + z_{21} + z_{22})$

Temporally Non-Coincident Observation

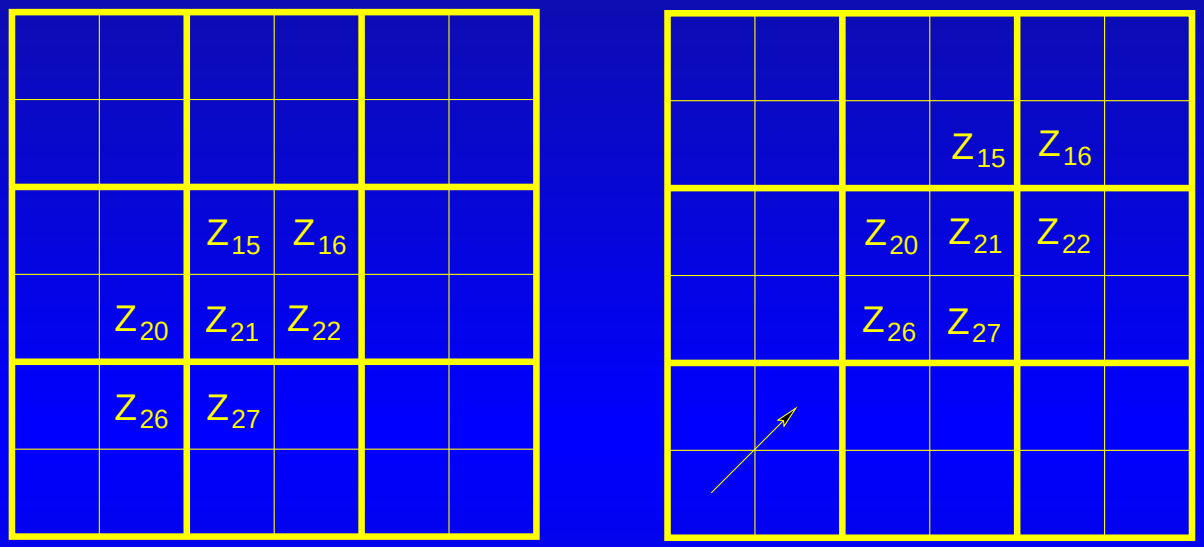

Original image : $y_{o_5} = \frac{1}{4}(z_{15} + z_{16} + z_{21} + z_{22})$

Translated image : $y_{t_5} = \frac{1}{4}(z_{20} + z_{21} + z_{26} + z_{27})$

Image Resampling

- Objective: Sampling of discrete image under coordinate transformation
- Discrete input image (texture): $f(\boldsymbol{u})$ with $\boldsymbol{u} = [u\ v]^T \in \mathbb{Z}^2$
- Discrete output image (warped): $g(\boldsymbol{x})$ with $\boldsymbol{x} = [x\ y]^T \in \mathbb{Z}^2$
- Forward mapping: $H : \boldsymbol{u} \longmapsto \boldsymbol{x}$
- Simplistic approach: $\forall \boldsymbol{x} \in \mathbb{Z}^2, \quad g(\boldsymbol{x}) = f(H^{-1}(\boldsymbol{x}))$
- Problems:
  1. $H^{-1}(\boldsymbol{x})$ need not fall on sample points (interpolation required)
  2. $H^{-1}(\boldsymbol{x})$ may undersample $f(\boldsymbol{u})$ resulting in aliasing

     (This occurs when the the mapping results in minification)

Conceptual Image Resampling Pipeline

1. Continuous reconstruction (interpolation) of input image (texture):

$$f_c(\boldsymbol{u}) = f(\boldsymbol{u}) \circledast r(\boldsymbol{u}) = \sum_{\boldsymbol{k} \in \mathbb{Z}^2} f(\boldsymbol{k}) \cdot r(\boldsymbol{u} - \boldsymbol{k})$$

2. Warp the continuous reconstruction:

$$g_c(\boldsymbol{x}) = f_c\left(H^{-1}(\boldsymbol{x})\right)$$

3. Apply the anti-alias prefilter $p(\boldsymbol{x})$:

$$g'_c(\boldsymbol{x}) = g_c(\boldsymbol{x}) \circledast p(\boldsymbol{x}) = \int g_c(\boldsymbol{\alpha}) \cdot p(\boldsymbol{x} - \boldsymbol{\alpha}) \, \mathrm{d}\boldsymbol{\alpha}$$

4. Sample to produce the discrete output image:

$$g(\boldsymbol{x}) = g'_c(\boldsymbol{x}) \quad \text{for} \quad \boldsymbol{x} \in \mathbb{Z}^2$$

Realizable Image Resampling Pipeline

- Never reconstruct continuous images:

$$
\begin{aligned}
g(\boldsymbol{x})\Big|_{\boldsymbol{x}\in\mathbb{Z}^2} &= g_c'(\boldsymbol{x})\Big|_{\boldsymbol{x}\in\mathbb{Z}^2} \\
&= \int f_c\left(H^{-1}(\boldsymbol{\alpha})\right)\cdot p(\boldsymbol{x}-\boldsymbol{\alpha})\,\mathrm{d}\boldsymbol{\alpha} \\
&= \int p(\boldsymbol{x}-\boldsymbol{\alpha})\sum_{\boldsymbol{k}\in\mathbb{Z}^2} f(\boldsymbol{k})\cdot r\left(H^{-1}(\boldsymbol{\alpha})-\boldsymbol{k}\right)\mathrm{d}\boldsymbol{\alpha} \\
&= \sum_{\boldsymbol{k}\in\mathbb{Z}^2} f(\boldsymbol{k})\rho(\boldsymbol{x},\boldsymbol{k})
\end{aligned}
$$

where

$$
\rho(\boldsymbol{x},\boldsymbol{k}) = \int p(\boldsymbol{x}-\boldsymbol{\alpha})\cdot r\left(H^{-1}(\boldsymbol{\alpha})-\boldsymbol{k}\right)\mathrm{d}\boldsymbol{\alpha}
$$

is a <u>spatially varying</u> resampling filter.

Realizable Image Resampling Pipeline

- Consider the resampling filter

$$\rho(\boldsymbol{x}, \boldsymbol{k}) = \int p(\boldsymbol{x} - \boldsymbol{\alpha}) \cdot r\left(H^{-1}(\boldsymbol{\alpha}) - \boldsymbol{k}\right) \mathrm{d}\boldsymbol{\alpha}$$

  - expressed i.t.o. warped reconstruction filter r
  - integration in $\boldsymbol{x}$-space (warped)

- Change variables $\boldsymbol{\alpha} = H(\boldsymbol{u})$

$$\rho(\boldsymbol{x}, \boldsymbol{k}) = \int p(\boldsymbol{x} - H(\boldsymbol{u})) \cdot r\left(\boldsymbol{u} - \boldsymbol{k}\right) \left|\frac{\partial H}{\partial \boldsymbol{u}}\right| \mathrm{d}\boldsymbol{u}$$

  - expressed i.t.o. the warped prefilter p
  - integrate in $\boldsymbol{u}$-space (texture)

Multiframe Observation Model

- Observe low resolution image sequence $g^{(i)}(\boldsymbol{x}),\ i \in \{1, 2, \ldots, P\}$
- Observations derive from a continuous scene $f_c(\boldsymbol{u})$
- Related via:
  - Coordinate transformations $H^{(i)} : \boldsymbol{u} \mapsto \boldsymbol{x}$ (scene/camera motion)
  - Spatially varying PSF's $h^{(i)}$ (lens/sensor PSF, defocus, motion blur...)
  - Spatial sampling

$$g^{(i)}(\boldsymbol{x})\Big|_{\boldsymbol{x}\in\mathbb{Z}^2} = \int h^{(i)}(\boldsymbol{x}, \boldsymbol{\alpha}) \cdot f_c\left(H^{(i)^{-1}}(\boldsymbol{\alpha})\right) \mathrm{d}\boldsymbol{\alpha}.$$

Discrete-Discrete Multiframe Observation Model

- Seek discretized approximation of $f_c(\boldsymbol{u})$ on high-resolution sampling lattice
- Interpolate samples $f(\boldsymbol{k})$, $\boldsymbol{k} \in \mathbb{Z}^2$ using kernel h_r

$$f_c(\boldsymbol{u}) \approx \sum_{\boldsymbol{k} \in \mathbb{Z}^2} f(\boldsymbol{V}\boldsymbol{k}) \cdot h_r\left(\boldsymbol{u} - \boldsymbol{V}\boldsymbol{k}\right)$$

  - $\boldsymbol{V}$ is the sampling matrix

$$\boldsymbol{V} = \begin{bmatrix} 1/q_x & 0 \\ 0 & 1/q_y \end{bmatrix}$$

  - q_x and q_y are horizontal and vertical sampling densities

Discrete-Discrete Multiframe Observation Model

- Combine with earlier result:

$$g^{(i)}(\boldsymbol{x})\Big|_{\boldsymbol{x}\in\mathbb{Z}^2} = \int h^{(i)}(\boldsymbol{x},\boldsymbol{\alpha}) \cdot f_c\left(H^{(i)^{-1}}(\boldsymbol{\alpha})\right) \mathrm{d}\boldsymbol{\alpha}$$

$$= \int h^{(i)}(\boldsymbol{x},\boldsymbol{\alpha}) \sum_{\boldsymbol{k}\in\mathbb{Z}^2} f(\boldsymbol{V}\boldsymbol{k}) \cdot h_r\left(H^{(i)^{-1}}(\boldsymbol{\alpha}) - \boldsymbol{V}\boldsymbol{k}\right) \mathrm{d}\boldsymbol{\alpha}$$

- Compare with resampling expression:

$$g(\boldsymbol{x})\Big|_{\boldsymbol{x}\in\mathbb{Z}^2} = \int p(\boldsymbol{x}-\boldsymbol{\alpha}) \sum_{\boldsymbol{k}\in\mathbb{Z}^2} f(\boldsymbol{k}) \cdot r\left(H^{-1}(\boldsymbol{\alpha}) - \boldsymbol{k}\right) \mathrm{d}\boldsymbol{\alpha}$$

- Identical in form to resampling expressions

Discrete-Discrete Multiframe Observation Model

$\Rightarrow$ Define spatially varying observation filter (c.f. resampling filter)

$$\rho^{(i)}(\boldsymbol{x},\boldsymbol{k}) = \int h^{(i)}(\boldsymbol{x},\boldsymbol{\alpha}) \cdot h_r\left(H^{(i)^{-1}}(\boldsymbol{\alpha}) - \boldsymbol{V}\boldsymbol{k}\right) \mathrm{d}\boldsymbol{\alpha}$$

- Relate $g^{(i)}(\boldsymbol{x})$ to $f(\boldsymbol{k})$ via LSV equations

$$g^{(i)}(\boldsymbol{x})\Big|_{\boldsymbol{x}\in\mathbb{Z}^2} = \sum_{\boldsymbol{k}\in\mathbb{Z}^2} f(\boldsymbol{V}\boldsymbol{k}) \cdot \rho^{(i)}(\boldsymbol{x},\boldsymbol{k})$$

- Change of variables $\boldsymbol{\alpha} = H^{(i)}(\boldsymbol{u})$:

$$\rho^{(i)}(\boldsymbol{x},\boldsymbol{k}) = \int h^{(i)}\left(\boldsymbol{x}, H^{(i)}(\boldsymbol{u})\right) \cdot h_r\left(\boldsymbol{u} - \boldsymbol{V}\boldsymbol{k}\right) \left|\frac{\partial H^{(i)}}{\partial \boldsymbol{u}}\right| \mathrm{d}\boldsymbol{u}$$

  - expressed i.t.o warped PSF $h^{(i)}(\boldsymbol{x},\boldsymbol{\alpha})$
  - integrate in $\boldsymbol{u}$-space (restoration)

Multiframe Observation Model & Image Resampling

Multiframe observation model:

$$g^{(i)}(\boldsymbol{x})\Big|_{\boldsymbol{x}\in\mathbb{Z}^2} = \int h^{(i)}(\boldsymbol{x},\boldsymbol{\alpha}) \sum_{\boldsymbol{k}\in\mathbb{Z}^2} f(\boldsymbol{V}\boldsymbol{k}) \cdot h_r\left(H^{(i)^{-1}}(\boldsymbol{\alpha}) - \boldsymbol{V}\boldsymbol{k}\right) \mathrm{d}\boldsymbol{\alpha}$$

Image resampling:

$$g(\boldsymbol{x})\Big|_{\boldsymbol{x}\in\mathbb{Z}^2} = \int p(\boldsymbol{x}-\boldsymbol{\alpha}) \sum_{\boldsymbol{k}\in\mathbb{Z}^2} f(\boldsymbol{k}) \cdot r\left(H^{-1}(\boldsymbol{\alpha}) - \boldsymbol{k}\right) \mathrm{d}\boldsymbol{\alpha}$$

Multiframe Observation Model & Image Resampling

Multiframe Observation model		Image resampling	
Discrete scene estimate	$f(\boldsymbol{u})$	Discrete texture	$f(\boldsymbol{u})$
Interpolation kernel	$h_r(\boldsymbol{u})$	Reconstruction filter	$r(\boldsymbol{u})$
Scene/camera motion	$H^{(i)}(\boldsymbol{u})$	Geometric transform	$H(\boldsymbol{u})$
Observation SVPSF	$h^{(i)}(\boldsymbol{x}, \boldsymbol{\alpha})$	Anti-alias pre-filter	$p(\boldsymbol{x})$
Observed images	$g^{(i)}(\boldsymbol{x})$	Warped output image	$g(\boldsymbol{x})$

Observation filter: $\rho^{(i)}(\boldsymbol{x}, \boldsymbol{k}) = \int h^{(i)}\left(\boldsymbol{x}, H^{(i)}(\boldsymbol{u})\right) \cdot h_r\left(\boldsymbol{u} - \boldsymbol{V}\boldsymbol{k}\right) \left| \frac{\partial H^{(i)}}{\partial \boldsymbol{u}} \right| \mathrm{d}\boldsymbol{u}$

Resampling filter: $\rho(\boldsymbol{x}, \boldsymbol{k}) = \int p\left(\boldsymbol{x} - H(\boldsymbol{u})\right) \cdot r\left(\boldsymbol{u} - \boldsymbol{k}\right) \left| \frac{\partial H}{\partial \boldsymbol{u}} \right| \mathrm{d}\boldsymbol{u}$

Linear Multiframe Observation Model

- Recall Linear Shift Varying observation equation

$$g^{(i)}(\boldsymbol{x})\Big|_{\boldsymbol{x}\in\mathbb{Z}^2} = \sum_{\boldsymbol{k}\in\mathbb{Z}^2} f(\boldsymbol{V}\boldsymbol{k}) \cdot \rho^{(i)}(\boldsymbol{x},\boldsymbol{k})$$

- Admits matrix-vector representation in finite case
- Single row of observation matrix: (observed images have N_r rows $\times N_c$ cols)

$$\boldsymbol{g}_j^{(i)} = \sum_{k=1}^{q_y N_r q_x N_c} A_{jk}^{(i)} \boldsymbol{f}_k$$

- Matrix-vector representation (single observed image)

$$\boldsymbol{g}^{(i)} = \boldsymbol{A}^{(i)} \boldsymbol{f}$$

Linear Multiframe Observation Model

- Matrix-vector representation (P images)

$$\boldsymbol{g} \doteq \begin{bmatrix} \boldsymbol{g}^{(1)} \\ \boldsymbol{g}^{(2)} \\ \vdots \\ \boldsymbol{g}^{(P)} \end{bmatrix} \quad \text{and} \quad \boldsymbol{A} \doteq \begin{bmatrix} \boldsymbol{A}^{(1)} \\ \boldsymbol{A}^{(2)} \\ \vdots \\ \boldsymbol{A}^{(P)} \end{bmatrix}$$

- Compact form

$$\boldsymbol{g} = \boldsymbol{A}\boldsymbol{f}$$

A Bayesian Framework for Restoration

- Classic linear inverse problem
- Usually underconstrained (too few observations)
- Ill-posed, so use regularized solution method with *a-priori* information
- Augment observation model to include noise

$$\boldsymbol{g} = \boldsymbol{A}^{(\boldsymbol{f})} + \boldsymbol{n}$$

- Noise model is zero-mean Gaussian

$$\mathcal{P}_{\boldsymbol{N}}(\boldsymbol{n}) = \frac{1}{(2\pi)^{\frac{PN_rN_c}{2}}|\boldsymbol{K}|}\exp\left\{-\frac{1}{2}\boldsymbol{n}^T\boldsymbol{K}^{-1}\boldsymbol{n}\right\}$$

$\boldsymbol{K}$ is the p.d. covariance matrix.

A Bayesian Framework for Restoration

- Noise model is zero-mean Gaussian

$$\mathcal{P}_{\boldsymbol{N}}(\boldsymbol{n}) = \frac{1}{(2\pi)^{\frac{PN_rN_c}{2}}|\boldsymbol{K}|}\exp\left\{-\frac{1}{2}\boldsymbol{n}^T\boldsymbol{K}^{-1}\boldsymbol{n}\right\}$$

$\boldsymbol{K}$ is the p.d. covariance matrix.

- Likelihood term

$$\begin{aligned}\mathcal{P}(\boldsymbol{g}|\boldsymbol{f}) &= \mathcal{P}_{\boldsymbol{N}}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f}) \\ &= \frac{1}{(2\pi)^{\frac{PN_rN_c}{2}}|\boldsymbol{K}|}\exp\left\{-\frac{1}{2}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f})^T\boldsymbol{K}^{-1}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f})\right\}\end{aligned}$$

A Bayesian Framework for Restoration

- Prior term (Markov Random Field)
- Density is Gibbsian (Hammersley-Clifford)

$$\mathcal{P}\left(\boldsymbol{f}\right) = \frac{1}{k_p} \exp\left\{ -\frac{1}{\beta} \sum_{c \in \mathcal{C}} \rho_T \left(\boldsymbol{\partial}_c \boldsymbol{f} \right) \right\}$$

- Huber penalty function $\rho_T(x)$ (edge preserving)
- Local interactions $\boldsymbol{\partial}_c$ approximate 2nd spatial derivates in 4 orientations

A Bayesian Framework for Restoration

- Combined objective function

$$\begin{aligned}\hat{\boldsymbol{f}}_{\mathrm{MAP}} &= \arg\max_{\boldsymbol{f}} \left\{ -\frac{1}{2}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f})^T \boldsymbol{K}^{-1}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f}) - \frac{1}{\beta}\sum_{c\in\mathcal{C}} \rho_T\left(\boldsymbol{\partial}_c \boldsymbol{f}\right) \right\} \\ &= \arg\min_{\boldsymbol{f}} \left\{ \frac{1}{2}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f})^T \boldsymbol{K}^{-1}(\boldsymbol{g}-\boldsymbol{A}\boldsymbol{f}) + \gamma\sum_{c\in\mathcal{C}} \rho_T\left(\boldsymbol{\partial}_c \boldsymbol{f}\right) \right\}\end{aligned}$$

- Use your favorite optimization technique to find $\hat{\boldsymbol{f}}_{\mathrm{MAP}}$
- Unique solution under very mild conditions

Simulation Details

Projection model	Ideal pinhole
Image array dimensions	128×128
Pixel dimensions	9μm $\times$ 9μm
Camera focal length	10 mm
Camera f/number	2.8
Illumination wavelength	550 nm
Diffraction limit cutoff	649.351 cycles/mm
Sampling rate	111.111 samples/mm
Folding frequency	55.5556 cycles/mm

Table 1: Camera intrinsic characteristics.

Simulation Details

Camera center			Camera gaze point		
x	y	z	x	y	z
-3.0902	-5.0000	9.5106	0.0100	0.0050	0.0000
-1.0453	-5.0000	9.9452	0.0033	0.0017	0.0000
1.0453	-5.0000	9.9452	-0.0033	-0.0017	0.0000
3.0902	-5.0000	9.5106	-0.0100	-0.0050	0.0000
5.0000	-5.0000	8.6603	-0.0167	-0.0083	0.0000
6.6913	-5.0000	7.4315	-0.0233	-0.0117	0.0000
8.0902	-5.0000	5.8779	-0.0300	-0.0150	0.0000

Table 2: Camera extrinsic parameters.

Directions for Future Research

1. Motion estimation (vis à vis superresolution restoration)
   - Simultaneous motion estimation and segmentation
   - Stochastic observation models and estimation methods
   - Sequential Monte Carlo methods (tracking, multimodal pdfs)
   - Robustness to undersampling, degradations in LR sequence
2. Observation model
   - Hyperspectral/multispectral data
   - Video compression models, sensor, acquisition degradations
   - Estimation of observation model from data
3. Prior model
   - “Recognition-based” priors and domain-specific knowledge
   - Temporal constaints

Directions for Future Research

4. Restoration frameworks
   - Bayesian methods with projected constraints
   - Simultaneous motion estimation and restoration
   - Sequential Monte Carlo methods (avoid batch mode processing)
   - Blind restoration
   - Optimization issues
5. Theoretical foundations
   - Virtually untouched!