BLIND SUPER-RESOLUTION FROM MULTIPLE UNDERSAMPLED IMAGES USING SAMPLING DIVERSITY

Faisal Al-Salem

Thesis Defense
May 20, 2010

Outline
PART I: Overview

- Introduction
  - Motivation.
  - Previous work.
  - Contribution.
  - Applications.
- A Novel Approach
  - LR images as basis signals
  - Sampling diversity
- The Expansion Coeff. of the Polyphase Comps (Overview)
- The Reference Polyphase Comp. (Overview)
- Results
- Summary

Outline
PART II: Details

- Solving for the Expansion Coeff. of the Polyphase Comps.
  - LS
  - TLS
  - PCA pre-denoising.
- Estimation of The Reference Polyphase Comp.
  - Minimizing the distance in the pixel domain.
  - Minimizing the distance in a decorrelated subspace.
  - An alternative to estimating the reference PPC.
- Future Work

INTRODUCTION (1)
MOTIVATION

Motivation: Why Super-resolution?

BLUR

- Diffraction limit.
- Loss of detail: soft images.
- Solution: optical super-resolution (beat the diffraction limit).

ALIASING

- Low –sensor density.
- Loss of detail: few pixels ⇒ blocky images.
- Solution: signal processing-based super-resolution.

Motivation: Cost vs. Physical Limits

COST REDUCTION

Need more pixels?

⇒ use larger imaging chip (large increase in cost).

<u>OR</u>: use smaller pixel size.

⇒ fewer photons/pixel

⇒ use very high quality photo sensors that can perform well under deprived light conditions.

⇒ Again, substantial increase in cost.

<u>OR</u> : use SR techniques.

BEYOND COST REDUCTION

Sensor is already diffraction limited?

⇒ zoom out to cover larger FOV (very important feature for many applications).

⇒ Higher density sensor is still required.

However, there is an optimal physical limit on pixel density (and chip/lens size). Also, particularly large pixel spacing is required in some applications (for thermal isolation in infrared imaging, for example).

⇒ SR is the only option when the optimal physical limits of sensor manufacturing are met.

INTRODUCTION (2)
PREVIOUS WORK

Previous Work: The General Setup

Multiframe SR methods generally assume:

- The original scene is static.
- Same downsampling in the horizontal and vertical directions.
- Additive noise corruption.
- Each measured LR image is the result of different *relative* scene motion (and blur).
- The motion vectors are usually assumed either known or reliably estimated using image registration methods. Some SR methods jointly estimate the HR image and motion parameters.

Previous Work: Solution of an Inverse Problem

- Matrix formulation

$$\begin{bmatrix} \underline{y}_1 \\ \vdots \\ \underline{y}_K \end{bmatrix} = \begin{bmatrix} \mathcal{D}H_1F_1 \\ \vdots \\ \mathcal{D}H_KF_K \end{bmatrix} \underline{u} + \begin{bmatrix} \underline{\upsilon}_1 \\ \vdots \\ \underline{\upsilon}_K \end{bmatrix} \Leftrightarrow \underline{Y} = S\underline{u} + \underline{\mathcal{V}}.$$

- Ill-posed, huge size problem
  - Example: if the size of the HR image is 500x500, then an (over)determined system matrix will have (at least) 250,000x250,000 = 62,500,000,000 elements.
  - Iterative, robust and stable solution is needed.
  - Farsiu et al. proposed an L1-norm data fitting term for robustness and an edge-preserving BTV for regularization.
  - Using their SR software, we implemented their method “Iterative L1” for the purpose of comparison.

S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar, “Fast and robust multiframe super resolution,” *IEEE Trans. IP*, 2004.

INTRODUCTION (3)
CONTRIBUTION

Contribution

Our proposed multiframe super-resolution method has the following characteristics:

- It can make use of (global/local) motion and/or blur to super-resolve images.
- Blind reconstruction: it requires no knowledge or estimation of the degradation process.
- Very fast.

INTRODUCTION (4)
APPLICATIONS

Applications (1): Global Translational Motion

- Classical SR problem.
- Example: a camera recording a sequence of a static scene while moving with slight translations.

TYPICAL SR METHODS

- Reliable image registration is required.
- Given the global motion vectors, and depending on the SR method, reconstruction can be *relatively* fast.

PROPOSED METHOD

- No need for registration.
- Very fast.

Applications (2): Super-resolution from Random Vibrations

EXAMPLES

Airborne/ground reconnaissance and machine vision systems.

- Vibrations are inevitable during imaging.
- Despite the best mechanical stabilization systems, images are still distorted by random motion blur.

TYPICAL SR METHODS

- Mis-registration due to randomness of blur from frame to frame.
- Poor SR performance.

PROPOSED METHOD

- No need for registration or blur estimation.
- Can make use of both motion and random blur.

Applications (3): Atmospheric Distortions

- Ground-based astronomical imaging and satellite imaging of the Earth
- Time-variant, shift-variant PSF.
- Severity of distortions increases with time of exposure and far-field imaging.
- Stacking (aim: deblurring)
  - High rate of frames per second $\Rightarrow$ decreases the severity of distortions but decreases the SNR too.
  - Process hundreds of frames to combine some of them without increasing the blur while increasing the SNR ratio.
  - Without distortions, stacking is needless while super-resolution is impossible.
- Satellite surveillance $\Rightarrow$ much less severe distortions.

A NOVEL APPROACH (1)
LR BASIS

LR Basis: LSI vs. LSV

Q: Can the polyphase components* (PPCs) of a HR image be written as a linear combination of LR images?

A1: Obviously, if the number of (lin. indept.) LR images equals to the number of pixels in a PPC, then this is always true.

A2: However, it can be shown that if each LR image corresponds to downsampling a *differently* distorted (warped/blurred) HR image, then the set of LR images can form a ***complete*** basis if:

1- Each distortion process can be modeled as a LSV transform, with a transform kernel that can be approx. as a set of r LSI kernels of the same finite size, acting on r different areas of the HR image.

2- And: $K \geq rL_1L_2$

K: # LR images.

$L_1 \times L_2$: the size of an LSI kernel (assumed to be equal to or larger than the downsampling factor, which is assumed to be the same in the vertical/horizontal direction).

LR Basis: Example

<u>Example:</u>

If each distortion is a different LSV PSF that can be approximated as a set of ***2 4x4*** LSI PSFs ⇒ we need ***32*** LR images to form a LR basis that can represent the ***9*** PPCs of the HR image.

<u>Note:</u>

- if each distortion is represented by a ***single 4x4*** LSI kernel, then we only need ***16*** LR images for a complete basis.
- if the distortions are LSV, but the number of available LR images is too small then we can super-resolve *subregions* of the HR image.

A NOVEL APPROACH (2)
SAMPLING DIVERSITY

The Property of Sampling Diversity

Define:

- The primary PPCs of the HR image are the I^2 PPCs corresponding to $\downarrow I\text{x}I$.
- The secondary PPCs are the J^2 PPCs corresponding to $\downarrow J\text{x}J$.

If:

I and J are relatively prime (e.g. $J = I+1$)

Then:

ANY primary PPC shares a *sub* PPC with ANY secondary PPC.

Usefulness?

If we know ONLY ONE of the secondary PPCs, then we already know a portion (a sub PPC) of EVERY primary PPC. We refer to this single known secondary PPC as the *reference PPC*.

Sampling Diversity & LR Basis Working Together

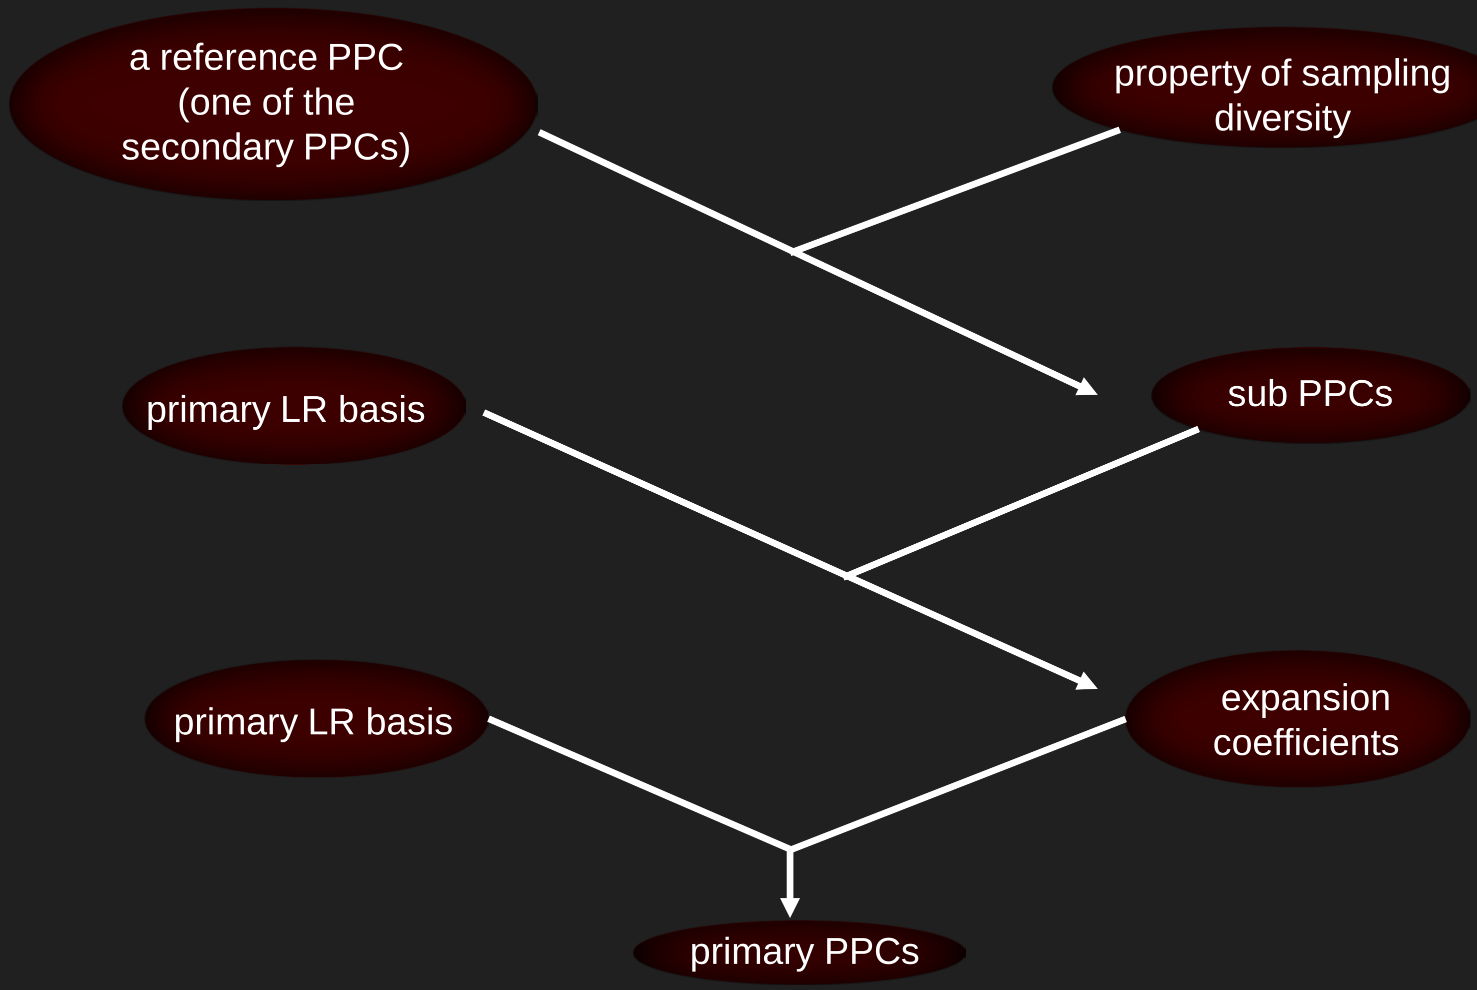

Sampling Diversity: Hardware Requirements

* A similar hardware setup is used for phase diversity (which is very different from sampling diversity).

For a HR image of size 1920×2560.

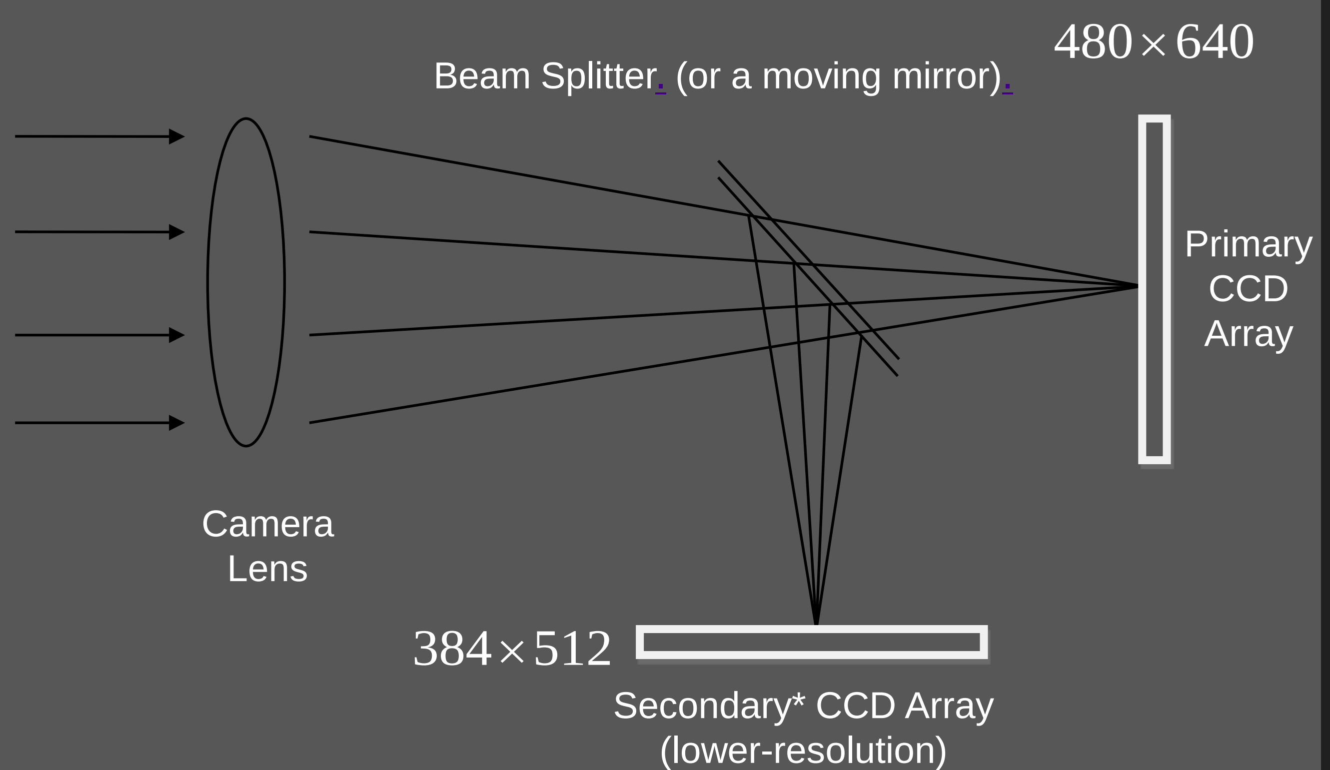

SOLVING FOR THE EXPANSION COEFFICIENTS

Solving for the Expansion Coefficients (1)

Solve *small* equations of the form:

$$A\underline{x}=\underline{b} \qquad (1)$$

$\underline{b}$ is a *sub* PPC (a portion) of a primary PPC.

A is the data submatrix (a submatrix of the LR basis matrix). # of columns = # of LR images.

$\underline{x}$ are the expansion coefficients.

Solving for the Expansion Coefficients (2)

Limitations?

- The LR basis is *noisy*? ⇒ *biased* solution.
  - Use PCA to pre-denoise the data.
- The LR basis is *incomplete*? ⇒ *biased* solution. (the PPCs are *partially* reconstructed).
- The best we can estimate is a HR image blurred (*biased*) by the CCD PSF.
- The estimated reference PPC is *noisy and biased*? ⇒ error in the RHS ⇒ noisy and *biased* solution.

Color Images

- Estimate the expansion coefficients of the green primary PPCs, using the green primary LR images as a basis set and an estimated green reference PPC.
- These are the same expansion coefficients of the red and blue PPCs in terms of the red and blue LR images, respectively.
- This means we only need green secondary LR images.

Post-processing

- TV denoising
  - Edge preserving.
  - Accounts for leftover noise.
- UM deblurring
  - Very simple/generic.
  - The best we can estimate is a HR image blurred by the CCD PSF.
  - Blur due to biased estimation of the reference PPC.
  - Blur due to noisy, incomplete LR basis.
- MD filtering
  - Impulsive noise after deblurring.
  - We estimate the HR image by estimating its PPCs separately and then interlacing, which might cause some subtle irregularities in pixel intensity levels that become more pronounced after sharpening.

ESTIMATION OF THE REFERENCE PPC

Estimation of the Reference PPC (1)

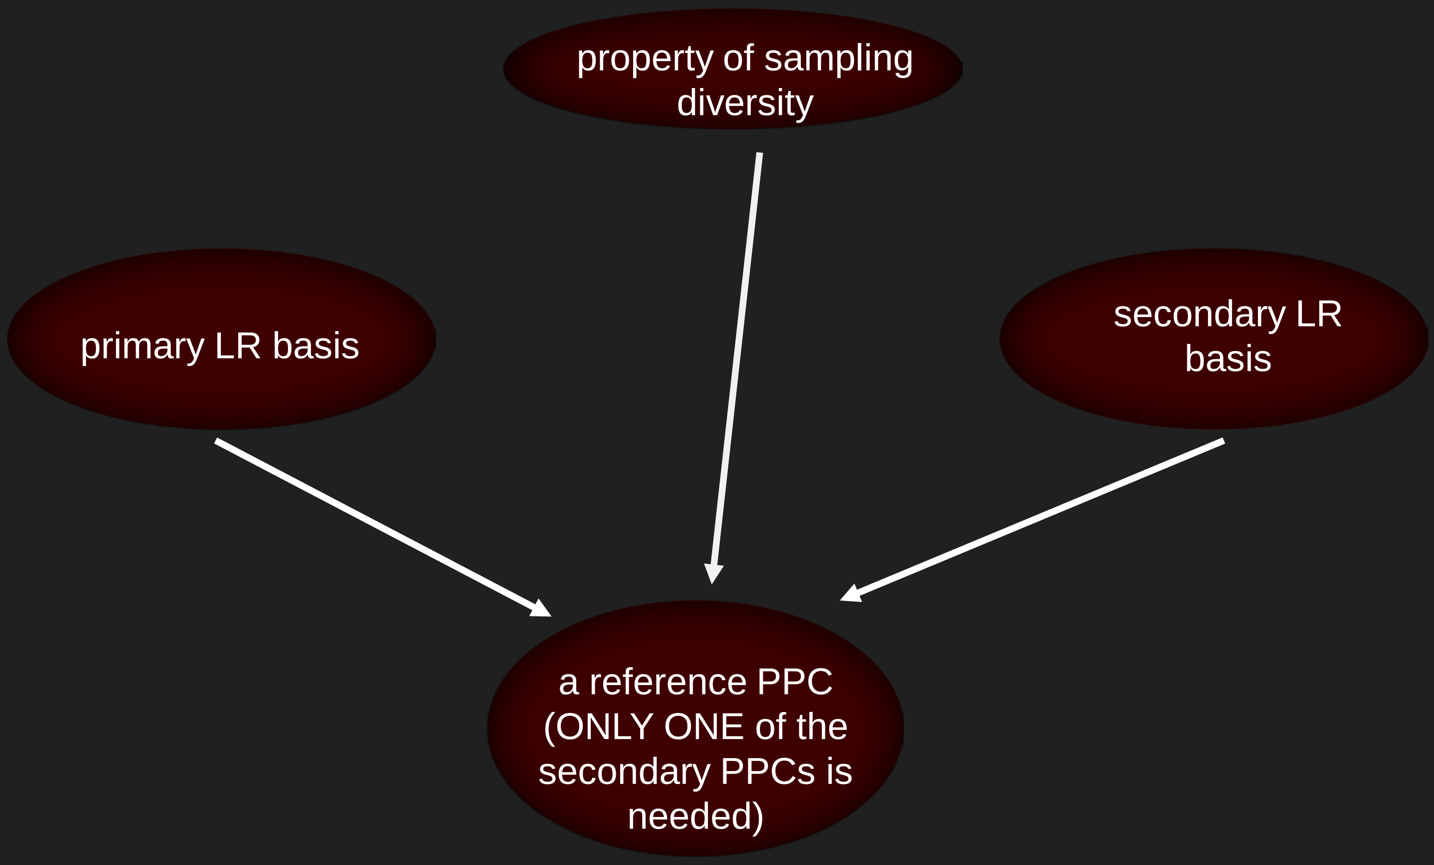

Estimation of the Reference PPC (2)

Solve ONLY ONE equation of the form:

$$A_1 \underline{x}_1 = A_2 \underline{x}_2 \qquad (2)$$

$\underline{x}_2$ are the expansion coefficients of the reference PPC in terms of the secondary LR Basis.

A_2 is a secondary data submatrix (a submatrix of the secondary LR basis matrix).

- We can solve Eq. (2) by minimizing the Euclidean distance in the pixel domain → very noisy/biased estimate.
- Minimize the Euclidean distance in decorrelated subspace → stable/less biased solution.

OR: Choose the "best" secondary PPC as a reference PPC.

EXPERIMENTAL RESULTS

Experiment: LSI Distortions (Synthetic Data: 16 frames)

Distorted HR sequence
(16 frames @ 5 frames/sec)

$\downarrow 4\times 4$

$\downarrow 5\times 5$

+ noise @ SNR = 30dB

+ noise @ SNR = 30dB

Primary LR sequence
(16 frames @ 5 frames/sec.)

Secondary LR sequence
(green color-band only)
(16 frames @ 5 frames/sec.)

Experiment: LSI Distortions (Synthetic Data: 16 frames)

Bicubic interp. 2.89 s.

Experiment: LSI Distortions (Synthetic Data: 16 frames)

Blind SR + (TV+UM+MD). 14.5 s.

Experiment: LSV Distortions (Synthetic Data: 100 frames)

Distorted HR sequence
(100 frames @ 5 frames/sec)

$\downarrow 4\times 4$

$\downarrow 5\times 5$

\+ noise @ SNR = 30dB

\+ noise @ SNR = 30dB

Primary LR sequence
(100 frames @ 5 frames/sec.)

Secondary LR sequence
(green color-band only)
(100 frames @ 5 frames/sec.)

Experiment: LSV Distortions (Synthetic Data: 100 frames)

Bicubic interp. 2.83 s.

Experiment: LSV Distortions (Synthetic Data: 100 frames)

Blind SR + (TV+UM+MD). 28.1 s.

Experiment: Approx. Pure Translations (Real Data: 35 frames)

2nd part of the HR sequence
(54 frames @ 5 frames/sec.)

$\downarrow 4 \times 4$

Primary LR sequence
(54 frames @ 5 frames/sec.)

Choose the 35 frames closest to the mean.

1st part of the HR sequence
(54 frames @ 5 frames/sec)

$\downarrow 5 \times 5$

Secondary LR sequence
(54 frames @ 5 frames/sec.)

Choose the best secondary PPC as the reference PPC.

Experiment: Approx. Pure Translations (Real Data: 35 frames)

Bicubic interp. 1.03 s.

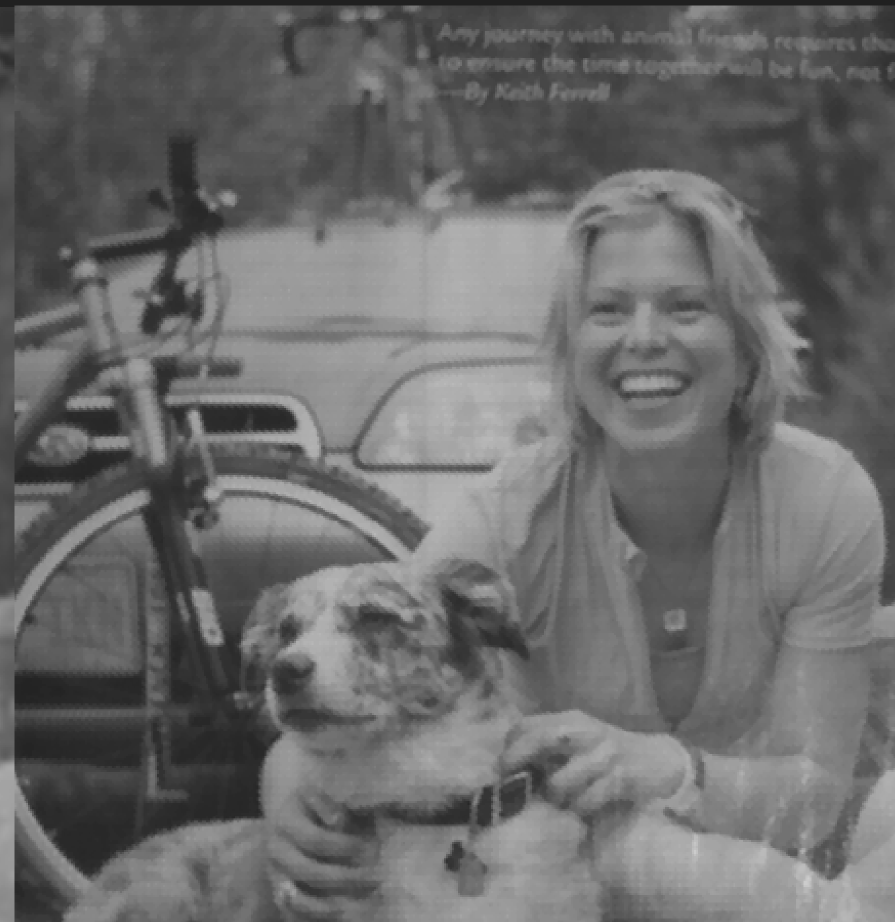

Iterative L1. ~ 4 mins.

Experiment: Approx. Pure Translations (Real Data: 35 frames)

Bicubic interp. 1.03 s.

Blind SR (w/best sec. LR)+ (UM+MD). 10.88 s.

Experiment: Approx. Pure Translations (Real Data: 35 frames)

Detail

Bicubic interp.

Iterative L1.

Experiment: Approx. Pure Translations (Real Data: 35 frames)

Detail

Bicubic interp.

Blind SR.

Experiment: Approx. Pure Translations (Real Data: 35 frames)

Detail

Bicubic interp.

Iterative L1.

Blind SR

Experiment: Approx. Pure Translations—Video (Real Data: 30 frames)

HR sequence
(100 frames @ 30 frames/sec.)

$\downarrow 4 \times 4$

Primary LR sequence
(100 frames @ 5 frames/sec.)

Discard every other frame.

Choose the 30 frames closest to the mean.

First HR frame (green color-band only).

$\downarrow 5 \times 5$

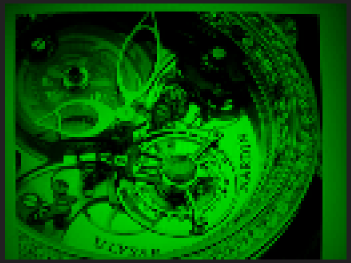

Experiment: Approx. Pure Translations—Video (Real Data: 30 frames)

Bicubic interp. 3 s.

Experiment: Approx. Pure Translations—Video (Real Data: 30 frames)

Iterative L1

Experiment: Approx. Pure Translations—Video (Real Data: 30 frames)

Blind SR (w/a single sec. LR)+ (TV+UM+MD). 21.3s.

Experiment: Random Vibrations (Real Data: 35 frames)

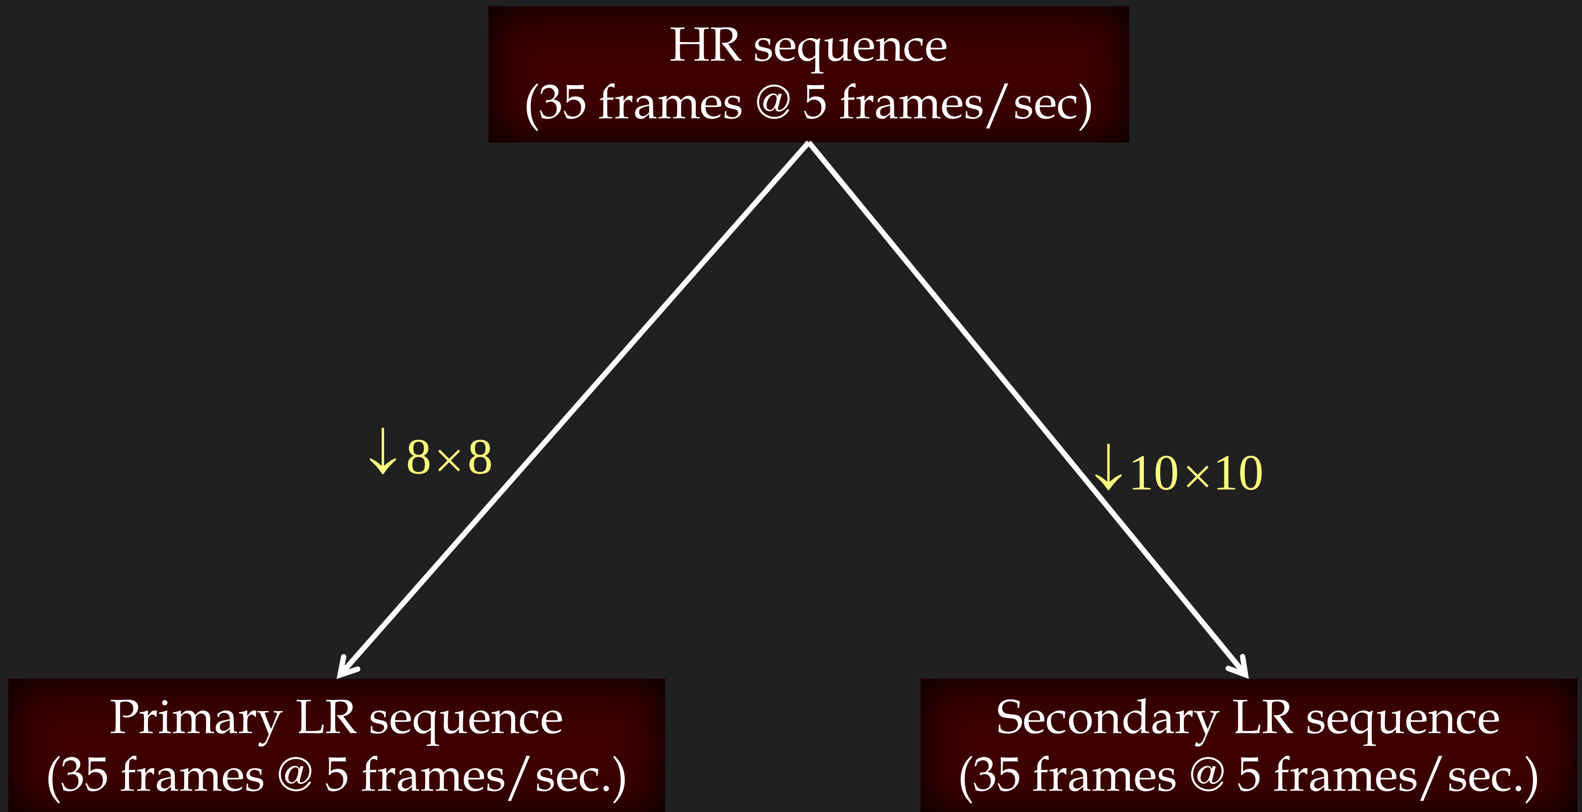

Experiment: Random Vibrations (Real Data: 35 frames)

Bicubic interp. 0.83 s.

Iterative L1 + UM.

Experiment: Random Vibrations (Real Data: 35 frames)

Blind SR(w/best sec. LR) + (TV+UM+MD). 7.22 s.

Blind SR + (TV+UM+MD). 6.9 s.

Experiment: Rhythmic Vibrations—Video (Real Data: 70 frames)

HR sequence
(70 frames @ 30 frames/sec)

↓5×5

Secondary LR sequence
(green color-band only)
(70 frames @ 5 frames/sec.)

Choose the best secondary LR frame as the reference PPC.

↓4×4

Primary LR sequence
(70 frames @ 5 frames/sec.)

Experiment: Rhythmic Vibrations—Video (Real Data: 70 frames)

Bicubic interp. + UM. 2.65 s.

Iterative L1 + UM. ~15+ min.

Experiment: Rhythmic Vibrations—Video (Real Data: 70 frames)

Bicubic interp. + UM. 2.65 s.

Blind SR (w/best sec. LR)+ (TV+UM). 13.57 s.

Experiment: Rhythmic Vibrations—Video (Real Data: 70 frames)

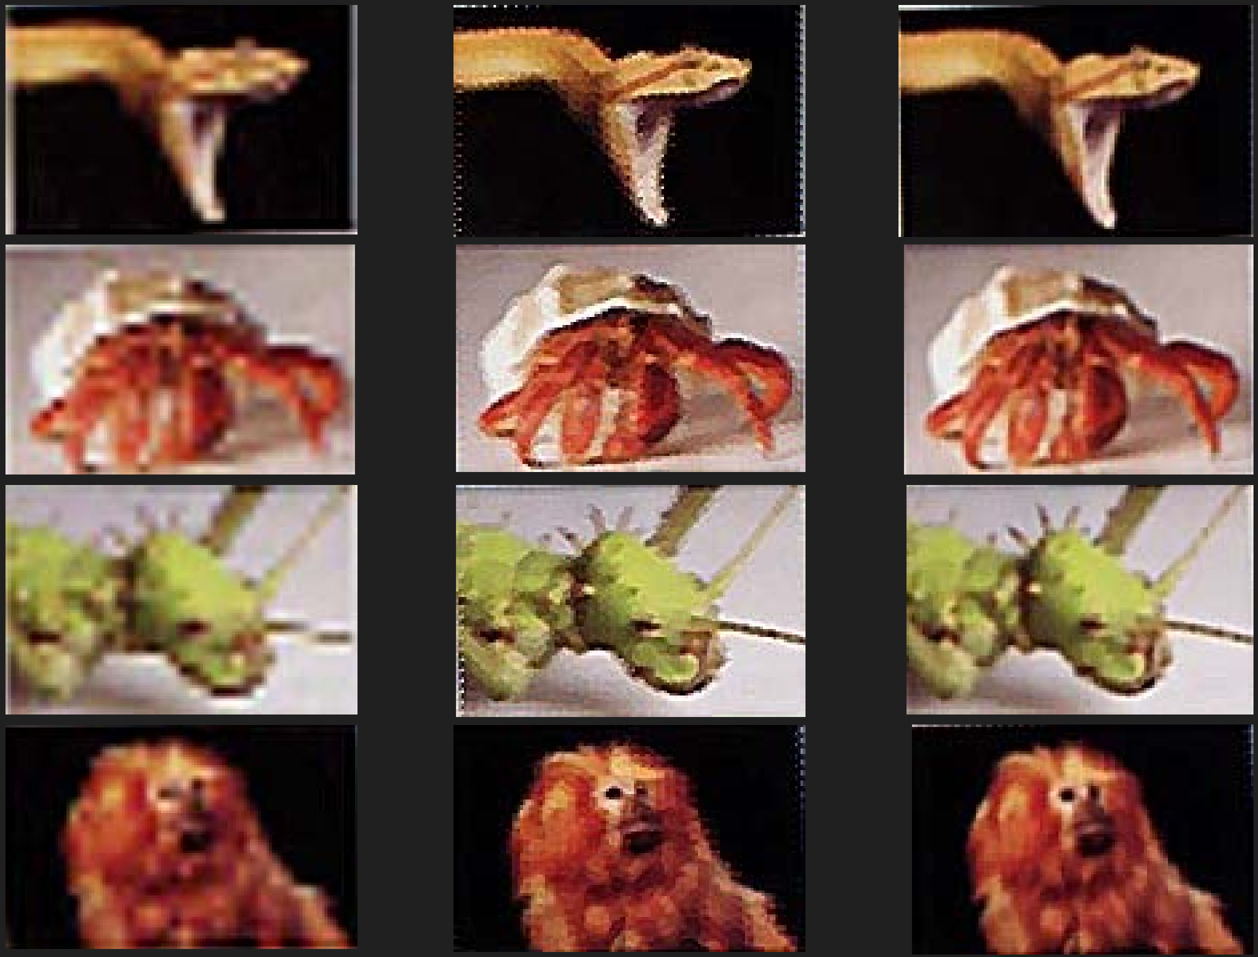

Bicubic interp.	Iterative L1	Blind SR

Experiment: Atmospheric Turbulence—Video (Real Data: 100 frames)

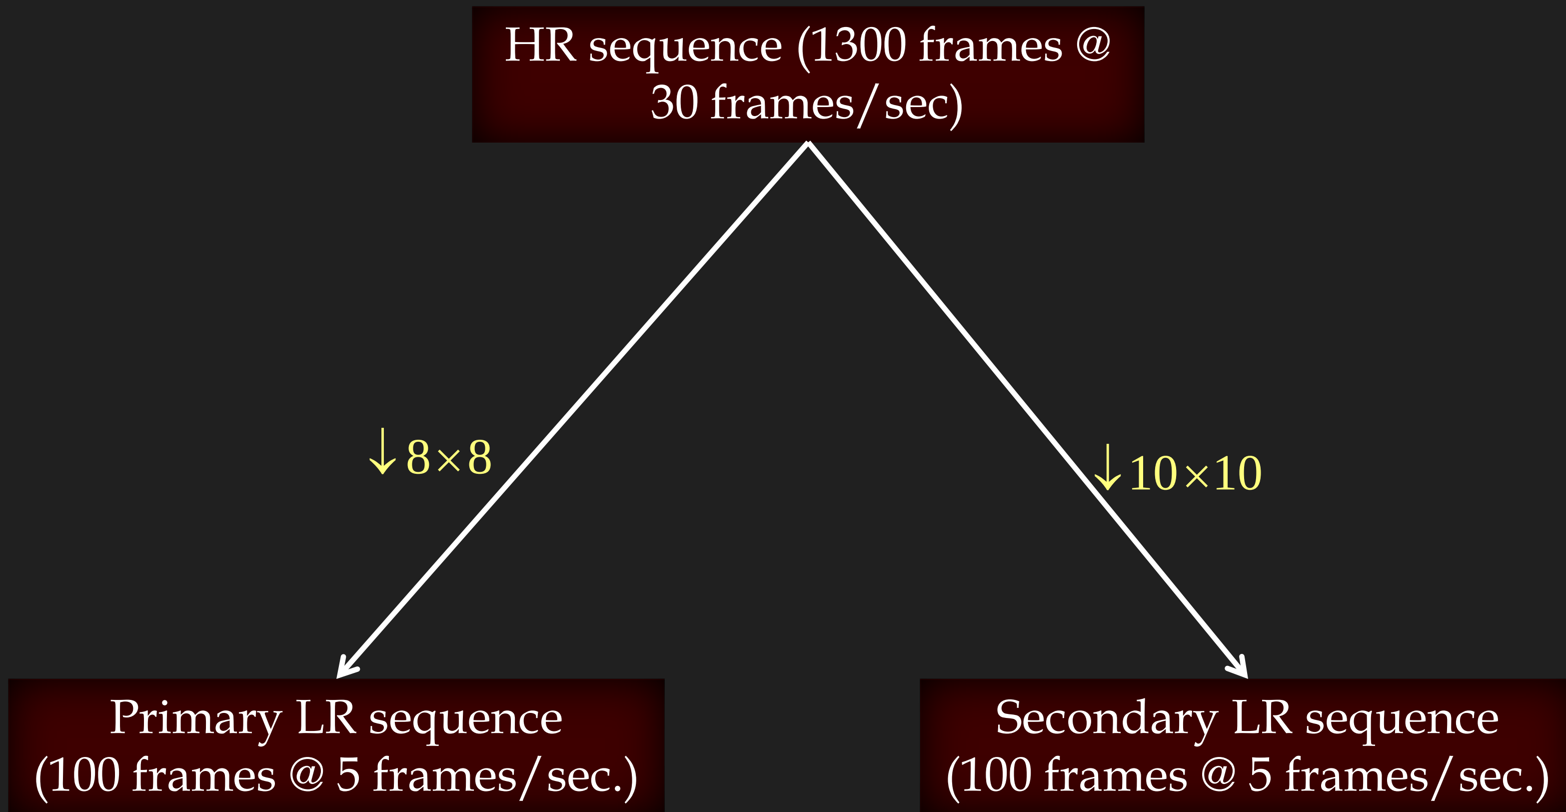

The original HR sequence is courtesy of Dr. Joseph Zawodny of NASA Langley Center.
1.7 pixels/Airy radius. Angular resolution = 0.34 arcseconds/pixel.

Experiment: Atmospheric Turbulence—Video (Real Data: 100 frames)

Bicubic interp. +UM. 0.83 s.

Experiment: Atmospheric Turbulence—Video (Real Data: 100 frames)

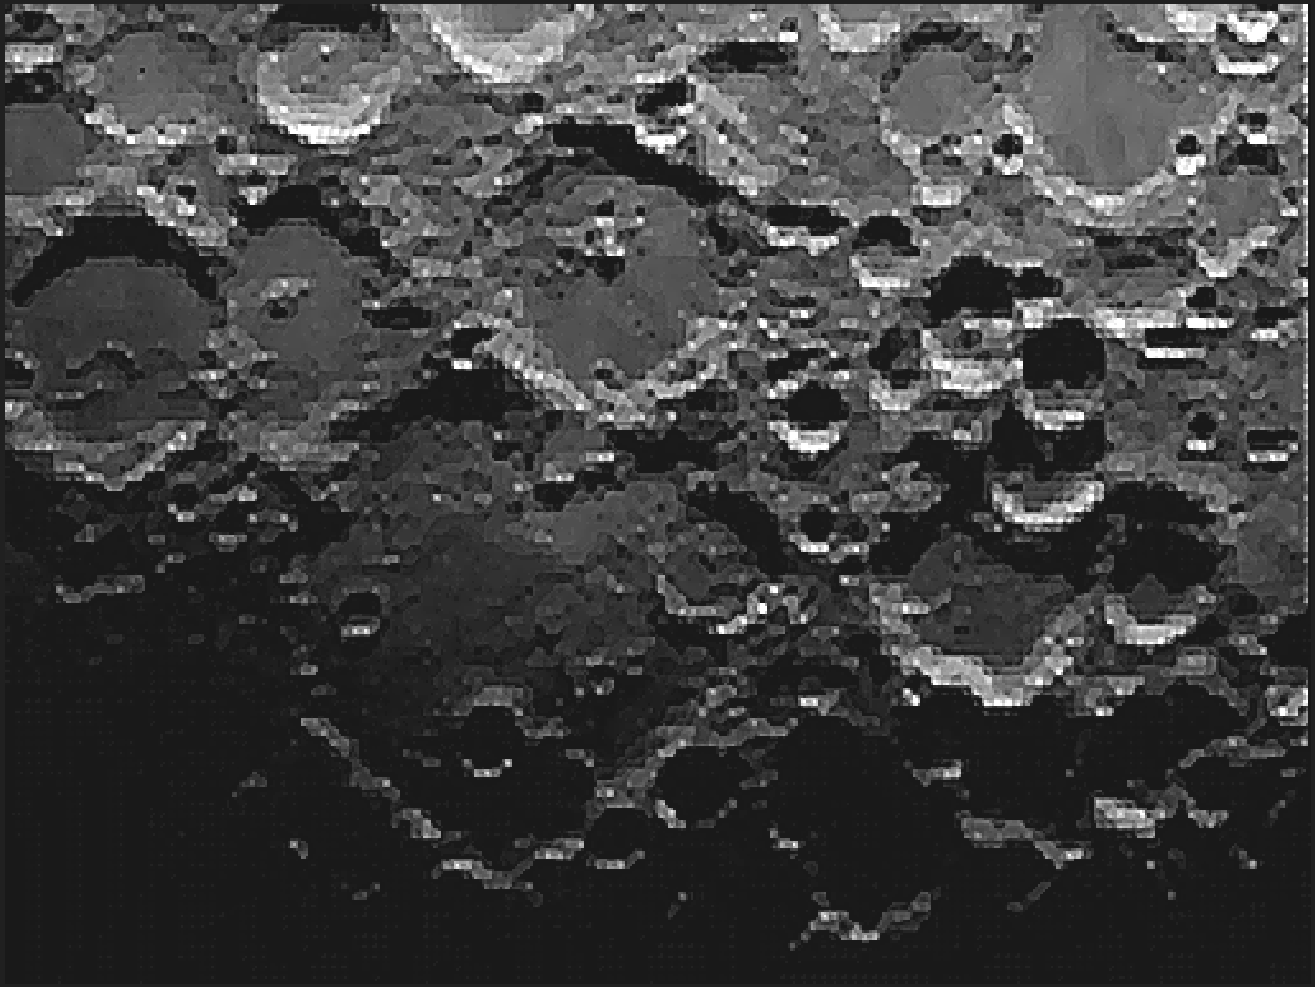

Iterative L1 + UM.

Experiment: Atmospheric Turbulence—Video (Real Data: 100 frames)

Blind SR (w/best sec. LR) + (UM+MD). 9.31 s.

Experiment: Atmospheric Turbulence—Video (Real Data: 100 frames)

Blind SR + (UM+MD). 10.9 s.

SUMMARY

Summary (1)

- Reformulate the problem as a change of basis.
- The ***completeness*** of the LR basis is dependent on the type (LSI vs. LSV) and extent (severity) of the distortion processes.
- Estimate the ***expansion coefficients*** of the PPCs in terms of the LR basis, using portions (sub PPCs) of the PPCs. These sub PPCs are estimated using the property of sampling diversity with a hardware requirement of adding an additional (lower resolution) sensor.

Summary (2)

- Our proposed method veers away from the major limitations associated with typical model-based solution of the SR problem. It is fast, does not require any estimation of the degradation process and is robust in the sense that we use no model . (no model $\Rightarrow$ no room for model errors).

- Besides requiring an additional *lower resolution* sensor, completeness of the LR basis is the only key assumption we make; the invalidity of which has only one consequence: the PPCs will be partially reconstructed.

- In certain applications where typical multiframe SR performs poorly (e.g. in the case of random vibrations), our method not only provides a much faster solution, it actually benefits from the random nature of distortions.

PART II
DETAILS

SOLVING FOR THE EXPANSION COEFFICIENTS

Solving for the Expansion Coefficients: An illustration of the Usefulness of the Sampling Diversity

For ease of illustration, suppose the primary and the secondary downsampling factors, I and J are equal to 4 and 5, respectively.

- ⇒ There are 16 primary PPCs. Each one of these has 25 sub PPCs.
- ⇒ There are 25 secondary PPCs. Each one of these has 16 sub PPCs.

Assume we *know*, say, the 13^{th} secondary PPC.

- ⇒ This is our reference PPC, and any sub PPC out of the 16 sub PPCs of this ref. PPC is also a sub PPC (a portion) of one of the 16 primary PPCs.
- ⇒ Example: the 10^{th} (out of 16) sub PPC of the ref. PPC is the 17^{th} (out of 25) sub PPC of the 4^{th} primary PPC.

Solving for the Expansion Coefficients: An illustration of the Usefulness of the Sampling Diversity

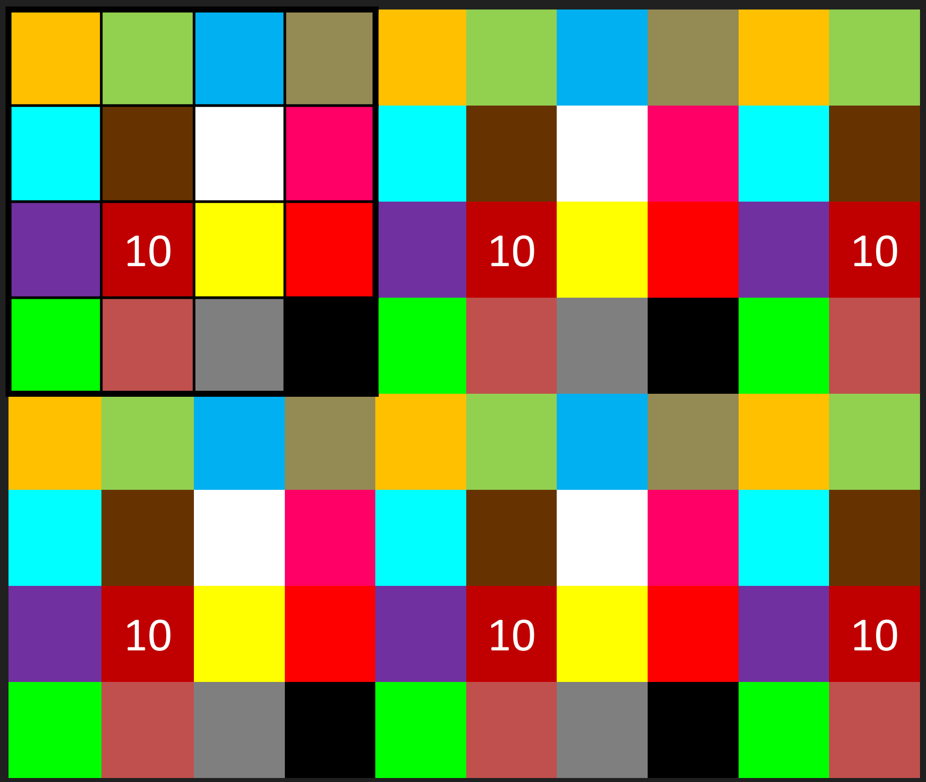

Each sub PPC of
the reference PPC
is a sub PPC (a portion) of
a primary PPC.

Example:

The 10th sub PPC of the
13th secondary PPC
is the 17th sub PPC of the
4th primary PPC.

Solving for the Expansion Coefficients: LS Solution

Assume:
1 – The LR basis is both complete and noiseless

$$\Rightarrow \underline{U}_4 = Y\underline{x}_4$$

2 - We got a corrupted version of the ref PPC (13th secondary PPC) $\Rightarrow$ All 16 sub PPCs are corrupted. The error is zero-mean white Gaussian noise.

Q: Given a noisy version of its 17th sub PPC, how can we estimate the 4th primary PPC?

$$\underline{U}_{ref,10} = \underline{U}_{4,17} = D_{17}\underline{U}_4 + \underline{e}$$

Ans: Solve for its expansion coefficients

$$\min_{\underline{x}_4} \left\| \underline{U}_{4,17} - D_{17}Y\underline{x}_4 \right\|^2 \equiv \min_{\underline{x}} \left\| \underline{b} - A\underline{x} \right\|^2$$

This is the ML estimate of the expansion coeffs of the 4th PPC.

Note: Y contains the (primary) LR images (unwrapped by column). We refer to it as the data matrix or the LR basis matrix.
- The matrix, A, is therefore a data submatrix (a submatrix of the LR basis matrix).

Solving for the Expansion Coefficients: Overdeterminedness of the Systems of Linear Equations

The system of linear equations:

$$A\underline{x} = \underline{b}$$

is overdetermined when

$$p = \frac{1}{J^2}\frac{M_1 M_2}{I^2} \geq K$$

⇒ Processing large LR images increases the overdeterminedness. In practice, this means we should super-resolve the largest possible subregions of the HR image (to within memory limitations).

⇒ For a given size of the LR images, we must have

$$J = I+1,$$

for the maximum possible overdeterminedness.

Solving for the Expansion Coefficients: Biasedness of LS Solution

- LR images are always noisy ⇒ LS is (asymptotically) biased (and thus inconsistent).

- So the bias does not decrease much with overdeterminedness (*larger* LR images), but it DOES decrease a lot with *more* LR images (noise, effectively, renders a complete basis incomplete).

- To benefit more of the overdeterminedness of the systems of linear equations, we need a consistent estimator.

Solving for the Expansion Coefficients: TLS Solution

- TLS takes into account the noise in the data submatrix, A.
- If A and $\underline{b}$ are both corrupted with zero mean i.i.d (Gaussian) noise, TLS is asymptotically unbiased (ML) estimator. (It's at least weakly consistent, if error is zero mean, same variance and uncorrelated)

$$\underline{\hat{x}}_{\text{TLS}} = \left(A^T A - \tilde{\sigma}_{K+1}^2 I \right)^{-1} A^T \underline{b}$$

where $\tilde{\sigma}_{K+1}$ is the last singular value of the augmented matrix $\left[A \mid \underline{b} \right]$.

- Since LR images are highly correlated $\Rightarrow \left[A \mid \underline{b} \right]$ has last few singular values that are close to each other $\Rightarrow$ TLS can be very *unstable* since $A^T A - \tilde{\sigma}_{K+1}^2 I$ can have vanishing last singular values (interlacing theorem).

Solving for the Expansion Coefficients: TRTLS Solution

- Tikhonov regularized TLS:

$$\underline{\hat{x}}_{\mathrm{TRTLS}} = \left(A^T A + \left(\lambda^2 - \tilde{\sigma}_{K+1}^2 \right) I \right)^{-1} A^T \underline{b}$$

- LS is simply a Tikhonov regularized TLS at

$$\lambda^2 = \tilde{\sigma}_{K+1}^2$$

- Tikhonov regularization roughly translates to the *a priori* assumption that the *expansion coefficients* are zero-mean, equal variance, *uncorrelated*, and Gaussian distributed.

Solving for the Expansion Coefficients: A Regularization Term as a Function of the PPC?

Why not write the regularization term as a function of the PPC?

$$\min_{\hat{A},\underline{x}} \left\| \left[A \;\vdots\; \underline{b} \right] - \left[\hat{A} \;\vdots\; \hat{A}\underline{x} \right] \right\|_F^2 + \lambda\, \Gamma\left(Y\underline{x} \right)$$

- non-convex.

- no analytical solution.

- roughness penalties are meaningless, in our case.

Solving for the Expansion Coefficients: The MSE of an Estimated PPC

- Let $Y = Y_0 + \mathcal{V}$ and $\underline{\hat{x}} = \underline{x} + \underline{w}$

- Assume the noise in the data matrix is zero mean, same variance and uncorrelated.

- For tractability, assume that $\underline{w}$ is independent of $\mathcal{V}$ then:

$$\mathrm{MSE}\left(\underline{\hat{U}}_n\right) = \mathrm{Tr}\left(Y_0 R_w Y_0^T\right) + d\sigma_v^2\left(\left\|\underline{x}\right\|^2 + \left\|\underline{\mu}_w\right\|^2 + \mathrm{Tr}\left(R_w\right) + 2\underline{\mu}_w^T\underline{x}\right) + \left\|Y_0\underline{\mu}_w\right\|^2$$

$$= d\sigma_v^2\left\|\underline{x}\right\|^2 \quad \text{if} \quad \underline{w} = \underline{0}.$$

- *Part of* the variance of error in an estimated PPC is independent of the bias-variance tradeoff associated with any estimator of its expansion coefficients.

  $\Rightarrow$ penalizing the TLS solution for the expansion coefficients using a penalty term that is a function of the PPC is NOT useful.

Solving for the Expansion Coefficients: PCA Pre-denoising (Why?)

- LS → biased.
- TLS → asymptotically unbiased but numerically unstable.
- TRTLS → inaccurate assumptions on the expansion coefficients.
- Cannot use a roughness penalty term as a function of the PPC.
- Noise augmentation, in the estimated PPC, results from multiplying the expansion coefficients with the data matrix.
- Noise augmentation is *partly* independent of the bias-variance tradeoff of any estimator of its expansion coefficients.

- CONCLUSION? Pre-denoise the data!
- ✓ This reduces bias, and removes the advantage of TLS over LS.
- ✓ It reduces the noise augmentation (which exists even with zero-error in the estimation of the expansion coefficients).

Solving for the Expansion Coefficients: PCA Pre-denoising (How?)

- When the noise is uncorrelated with the same variance, PCA maximizes the SNR along the first few axes.

- By dropping the last few PCs, the (reconstruction) MSE corresponds mostly to noise.

- Estimate the covariance matrix of the *sub* LR images obtained by downsampling the primary and secondary sets of LR images, by ↓***J*x*J*** and ↓***I*x*I*,** respectively, thus obtaining

$$KJ^2 + K^S I^2$$

  samples to compute the sample covariance matrix.

- The sub LR images are then denoised via

$$\underline{\hat{y}}_k^{sub} = DD^T\left(\underline{y}_k^{sub} - \underline{\hat{\mu}}\right) + \underline{\hat{\mu}}$$

  where D is the PCA matrix containing the eigenvectors corresponding to the largest eigenvalues of the sample covariance matrix (corresponding to low order PCs).

ESTIMATION OF THE REFERENCE PPC

Estimation of the Reference Polyphase Component: Introduction

- In order to estimate the ref. PPC (ONLY ONE of the 25 secondary PPCs, e.g. the 13th), we need an *additional* (secondary) set of lower resolution images acquired by an *additional* (secondary) imaging sensor (of lower pixel density).

- Assume that the primary and secondary LR sets are two complete basis for representing the 16 primary PPCs and the 25 secondary PPCs, respectively.

- A sub PPC of the ref. PPC (the 13th secondary PPC) is equal to a sub PPC of a primary PPC.

  ⇒ pick any primary PPC, say, the 5th one.
  ⇒ the 15th (out of 16) sub PPC of the ref. PPC is the 24th (out of 25) sub PPC of the 5th primary PPC.

$$D_{24} \underbrace{Y \underline{x}_5}_{\underline{U}_5} = D_{15} \underbrace{Y^S \underline{x}_{13}}_{\underline{U}_{ref}}$$

Primary LR basis matrix

Secondary LR basis matrix

Estimation of the Reference Polyphase Component: Introduction (cont'd)

The previous equation is rewritten as:

$$A_1 \underline{x}_1 = A_2 \underline{x}_2$$

submatrix of the primary LR basis matrix.

the expansion coeff. of the 5th primary PPC.

the expansion coeff. of the ref. PPC.

submatrix of the secondary LR basis matrix

- Solve this (for $\underline{x}_2$) only once.
- Min the Euclidean dist. in the pixel domain.
- Min the Euclidean dist. in a decorrelated subspace.
- OR: Don't solve it!

Estimation of the Reference Polyphase Component: Min. Euclidean Dist. in the Pixel Domain

- Since we expect that the pair of vectors, $\underline{f} \in R(A_1)$ and $\underline{g} \in R(A_2)$ that best approximate the common sub PPC, have the shortest distance between them, we solve

$$\min_{\underline{x}_1, \underline{x}_2} \left\| A_1 \underline{x}_1 - A_2 \underline{x}_2 \right\|^2 \quad \text{subject to} \quad \left\| \underline{x}_1 \right\|^2 + \left\| \underline{x}_2 \right\|^2 = 1$$

which is equivalent to

$$\min_{\underline{x}} \left\| A \underline{x} \right\|^2 = \underline{x}^T A^T A \underline{x} \quad \text{subject to} \quad \left\| \underline{x} \right\|^2 = 1$$

where

$$A = \left[A_1 \;\vdots\; -A_2 \right]$$

$$\underline{x} = \left[\underline{x}_1^T \,,\; \underline{x}_2^T \right]^T$$

- Non-convex, but has the well known analytical solution: $\underline{x}$ = last right singular vector of A.

Estimation of the Reference Polyphase Component: Min. Euclidean Dist. in Decorrelated Subspace (Why?)

- The solution of the previous minimization problem can be numerically unstable, although we can denoise the solution.
- As a dissimilarity measure, the Euclidean distance is very sensitive to error when the variables (pixels) are highly correlated.
- Removing dependencies among pixels in $\underline{f} \in R(A_1)$ and $\underline{g} \in R(A_2)$ before deciding which $\underline{f}$ and $\underline{g}$ are with minimal dissimilarity, gives a less biased decision.
- PCA gives us a basis, in terms of which, the expansion coefficients (PCs) of (centered) $\underline{f}$ and $\underline{g}$ are uncorrelated.
- The highest variance PCs have the *greatest weight* in the choice of the pair of vectors with minimal dissimilarity. (these are assumed to represent *significant features*).

Estimation of the Reference Polyphase Component: Min. Euclidean Dist. in Decorrelated Subspace (How?)

- Find the shortest distance between decorrelated vectors:

$$\min_{\underline{x}_1, \underline{x}_2} \left\| D\left(A_1 \underline{x}_1 - A_2 \underline{x}_2\right) \right\|^2 \quad \text{subject to} \quad \left\|\underline{x}_1\right\|^2 + \left\|\underline{x}_2\right\|^2 = 1$$

which is equivalent to

$$\min_{\underline{x}} \left\| D^T A \underline{x} \right\|^2 = \underline{x}^T A^T D D^T A \underline{x} \quad \text{subject to} \quad \left\|\underline{x}\right\|^2 = 1$$

where

$$A = \left[A_1 \;\vdots\; -A_2 \right]$$

$$\underline{x} = \left[\underline{x}_1^T \, , \; \underline{x}_2^T \right]^T$$

and D is the same PCA matrix used for pre-denoising.

- Non-convex, but has the well known analytical solution: $\underline{x}$ = last right singular vector of $D^T A$.

An Alternative to Estimating the Reference PPC

- Solving $A_1\underline{x}_1=A_2\underline{x}_2$ is a case where the blind are leading the blind.
- Problems of this type are thus more sensitive to errors (compared to $A\underline{x}=\underline{b}$).
- Avoid solving it by picking one of the secondary LR images as our reference PPC.

<u>Intuition:</u>

Suppose we have secondary LR images corresponding to perfect global motions.

⇒ Each secondary LR image is a secondary PPC.
⇒ For more complex cases, secondary LR images are mixtures of secondary PPCs.
⇒ Pick the sharpest (non-outlier) secondary LR image.
⇒ Which sec. PPC, the chosen sec. LR image approximates best?
⇒ Solve for the I^2 primary PPCs, J^2 times. Pick the smoothest HR image.

SINCE THE PROPOSAL

Since the Proposal

PROPOSAL PROMISES

- Denoise (preprocess) data matrices using PCA (or kernel PCA?).
- Get rid of outlier images (trimming? MCD?)
- Better decorrelation, for a better estimate of the reference PPC?

WHAT HAS BEEN DONE

- Lower bias AND *noise augmentation* by PCA predenoising the data.
- LR images are highly correlated → outliers are the furthest from the mean.
- Use a decorrelating PCA matrix derived from all sub LR images (use the denoising matrix).
- Choose the best secondary LR frame, instead of estimating the reference PPC.
- Straightforward: Color images/post-processing.

Some Failed Attempts for Better Estimation of the Reference PPCs

Shrinkage.

$$\min_{\underline{x}_1, \underline{x}_2} \| D^T (A_1 \underline{x}_1 - A_2 \underline{x}_2) \|^2 + \lambda \| Y_2 \underline{x}_2 - \underline{y}_{best}^S \|^2$$

ICA, NMF.

Min angle, L1, Chebychev (all using CCP)

$$\min_{\underline{x}_1, \underline{x}_2, \underline{x}_3} \| D^T (A_1 \underline{x}_1 - A_2 \underline{x}_2) \|^2 + \| D^T (A_3 \underline{x}_3 - A_2 \underline{x}_2) \|^2$$

subject to $\| \underline{x}_2 \|^2 = 1.$

FUTURE WORK

Future Work
Pre-processing: Variants of PCA (1)

- As a *denoiser*, theoretical PCA performs best if the errors are uncorrelated with the same variance (regardless of distribution). Otherwise, solve:

$$\max_{\underline{e}_\ell} \frac{\underline{e}_\ell^T C_y \underline{e}_\ell}{\underline{e}_\ell^T R_v \underline{e}_\ell} \quad \text{subject to} \quad \underline{e}_q^T R_v \underline{e}_\ell = 0 \quad \text{for} \quad 1 \leq q < \ell, \quad \ell \neq 1$$

  This non-convex problem can be reduced to solving a system of non-linear equations.

- Empirical PCA (using the sample covariance matrix) performs well, if the noise is also Gaussian (sensor readout noise is white Gaussian).

Future Work
Pre-processing: Variants of PCA (2)

Other types of noise

- Shot noise (poisson) is due to fluctuation of photon counts, but it becomes more Gaussian-like distributed with more photons (due to larger pixels in LR images).
- For impulsive noise, robust PCA performs better.
- Small scale problem.
  - # of columns of data submatrices = # of LR images.
  - # of rows is M_1M_2/I^2J^2 for data submatrices (this is an important factor if we choose to PCA pre-denoise the data or estimate the reference PPC).

Future Work
Post-processing

- The literature on denoising/deblurring is vast.
- Limit the search to edge-preserving methods.
- Adaptive TV denoising for textured HR images.
- Joint deblurring and TV denoising.
- TV denoising for other types of noise (e.g. impulsive or poisson noise).

THANK YOU

BACKUP SLIDES

What is a Polyphase Component?

A PPC is a shifted and downsampled version of the HR image.

Shown: 4 out of 16 PPCs corresponding to ↓4x4.

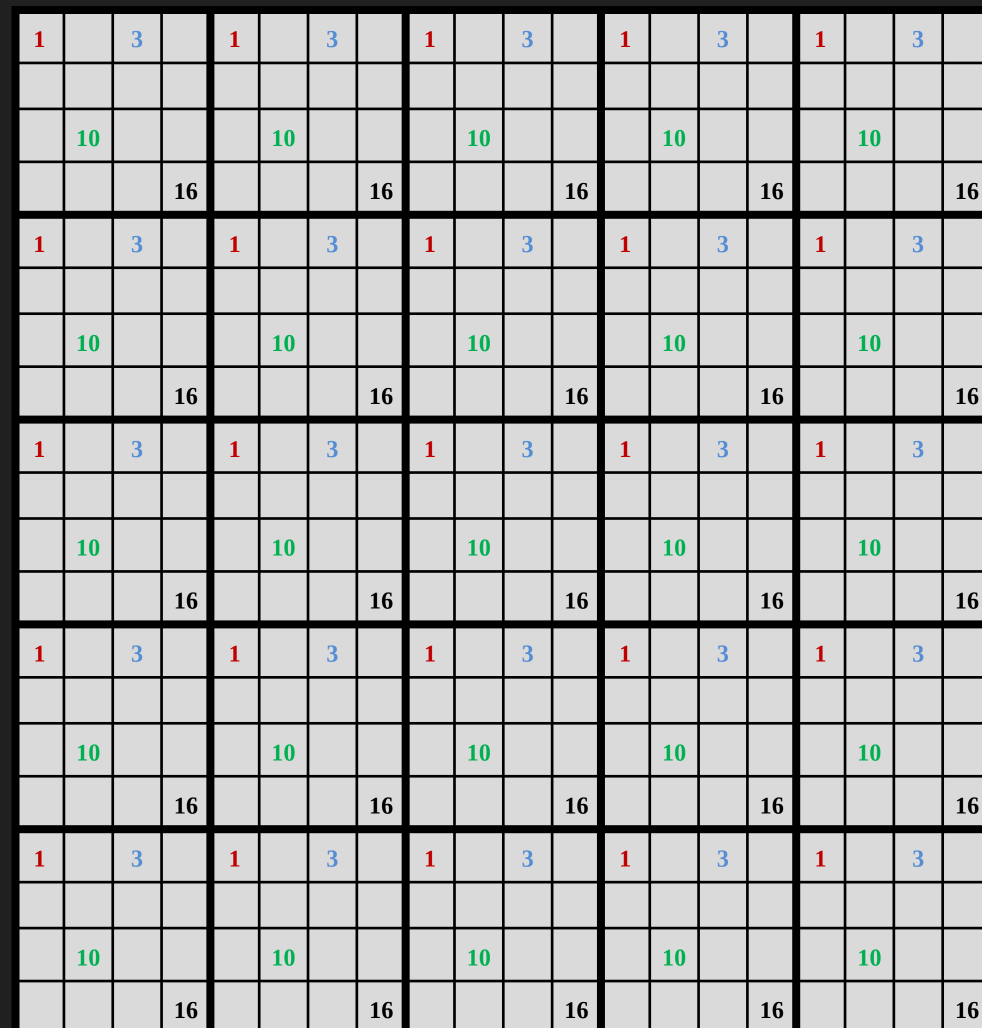

An illustration of the integration effect of the primary LR CCD array corresponding to ↓4x4

The light shaded areas represent the active portions of the *primary* LR pixels. The small blue squares represent the active portions of the pixels of the HR CCD array.

The weighting kernel (CCD PSF) is represented by a 4x4 Gaussian kernel with variance 1.

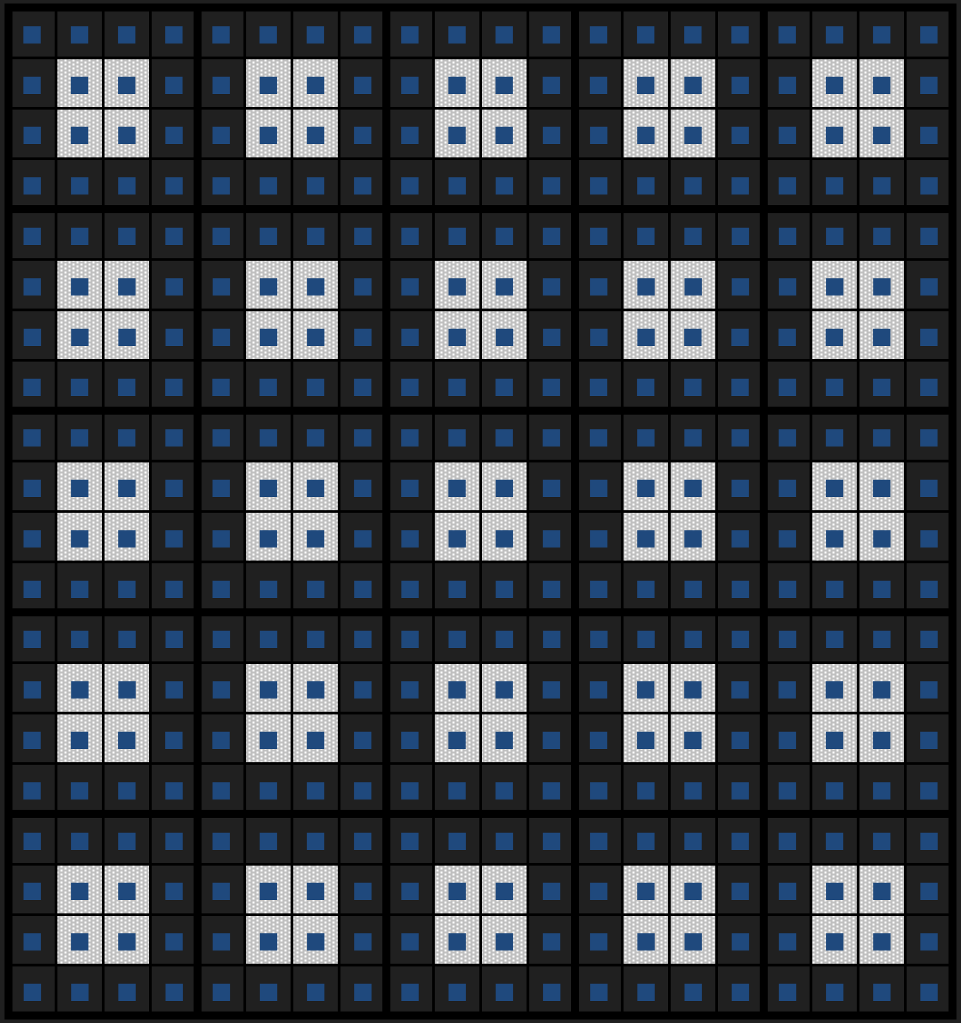

An illustration of the integration effect of the secondary LR CCD array corresponding to ↓5x5

The light shaded areas represent the active portions of the *secondary* LR pixels. The small blue squares represent the active portions of the pixels of the HR CCD array.

The weighting kernel (CCD PSF) is represented by a 5x5 Gaussian kernel with variance 1.

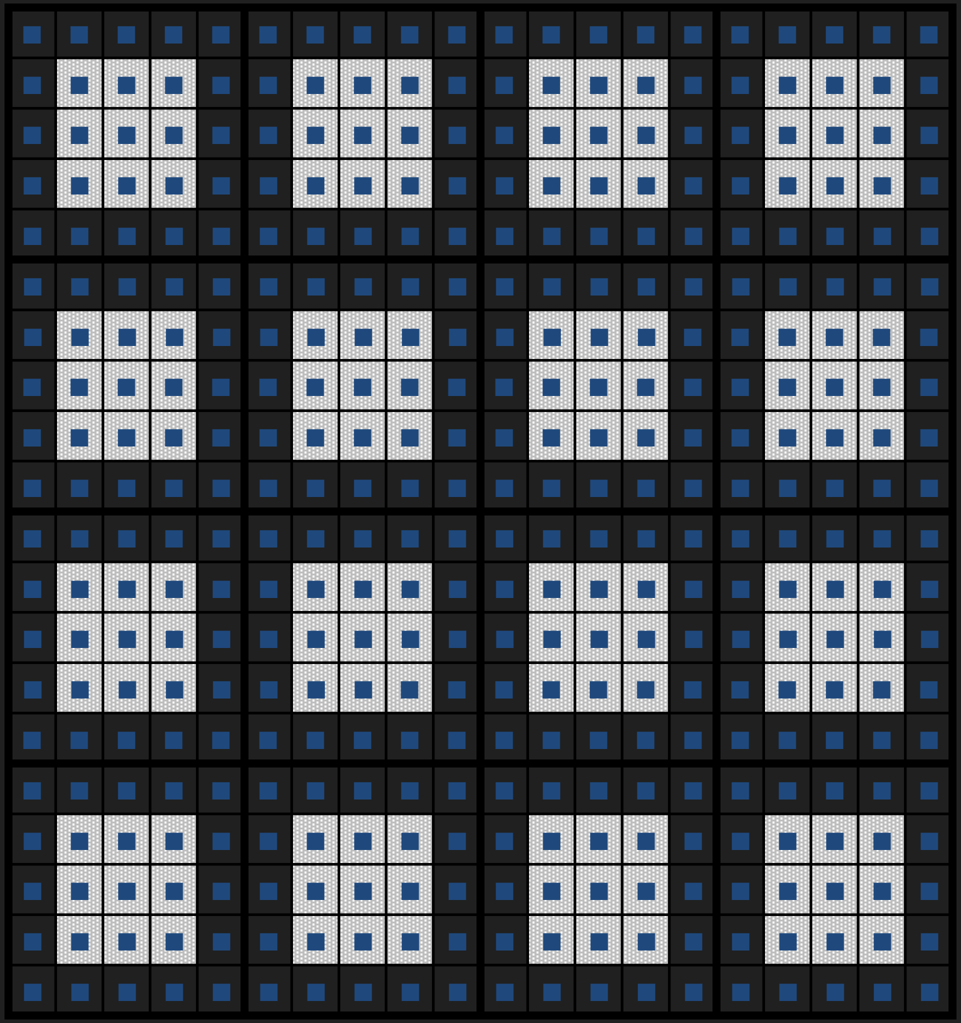

Future Work (Cont'd)

- Iteratively estimate the reference PPC based on a best secondary LR image (but how?).
- Build a prototype camera.
- Concentrate on satellite imaging of the Earth.
- Investigate whether single-frame SR might benefit from the property of sampling diversity, where the single (distortion-free) LR frame plays the role of the reference PPC. In particular, it might be easier to train a basis to reconstruct low resolution signals (PPCs) and as such, the sampling diversity idea could be extended to single frame SR and *without* the additional requirement of a secondary sensor.
- dynamic SR? we could use each secondary frame as a reference PPC, thus obtaining a sequence of SR images that are, in essence, HR versions of the secondary LR images. This, however, would probably require a temporal resolution high enough for a valid assumption of the rigidity of the scene within reasonably short time windows.

Results From the Proposal

SR image. Comp. time = 0.25 s.

Results From the Proposal

SR image.

Results From the Proposal

SR image. Comp. time = 0.65 s.