Super-Resolution Reconstruction of Images - An Overview

Michael Elad*

The Computer Science Department

The Technion, Israel

* Joint work with Arie Feuer – The EE department, Technion,
Yaacov Hel-Or – IDC, Israel,
Peyman Milanfar – The EE department, UCSC, &
Sina Farsiu – The EE department, UCSC.

Basic Super-Resolution Idea

Given: A set of low-quality images:

How?

Required: Fusion of these images into a higher resolution image

Comment: This is an actual super-resolution reconstruction result

Agenda

- Modeling the Super-Resolution Problem
  Defining the relation between the given and the desired images
- The Maximum-Likelihood Solution
  A simple solution based on the measurements
- Bayesian Super-Resolution Reconstruction
  Taking into account behavior of images
- Some Results and Variations
  Examples, Robustifying, Handling color
- Super-Resolution: A Summary
  The bottom line

Note: Our work thus-far has not addressed astronomical data, and this talk will be thus focusing on the fundamentals of Super-Resolution.

Chapter 1:
Modeling the Super-Resolution Problem

The Model

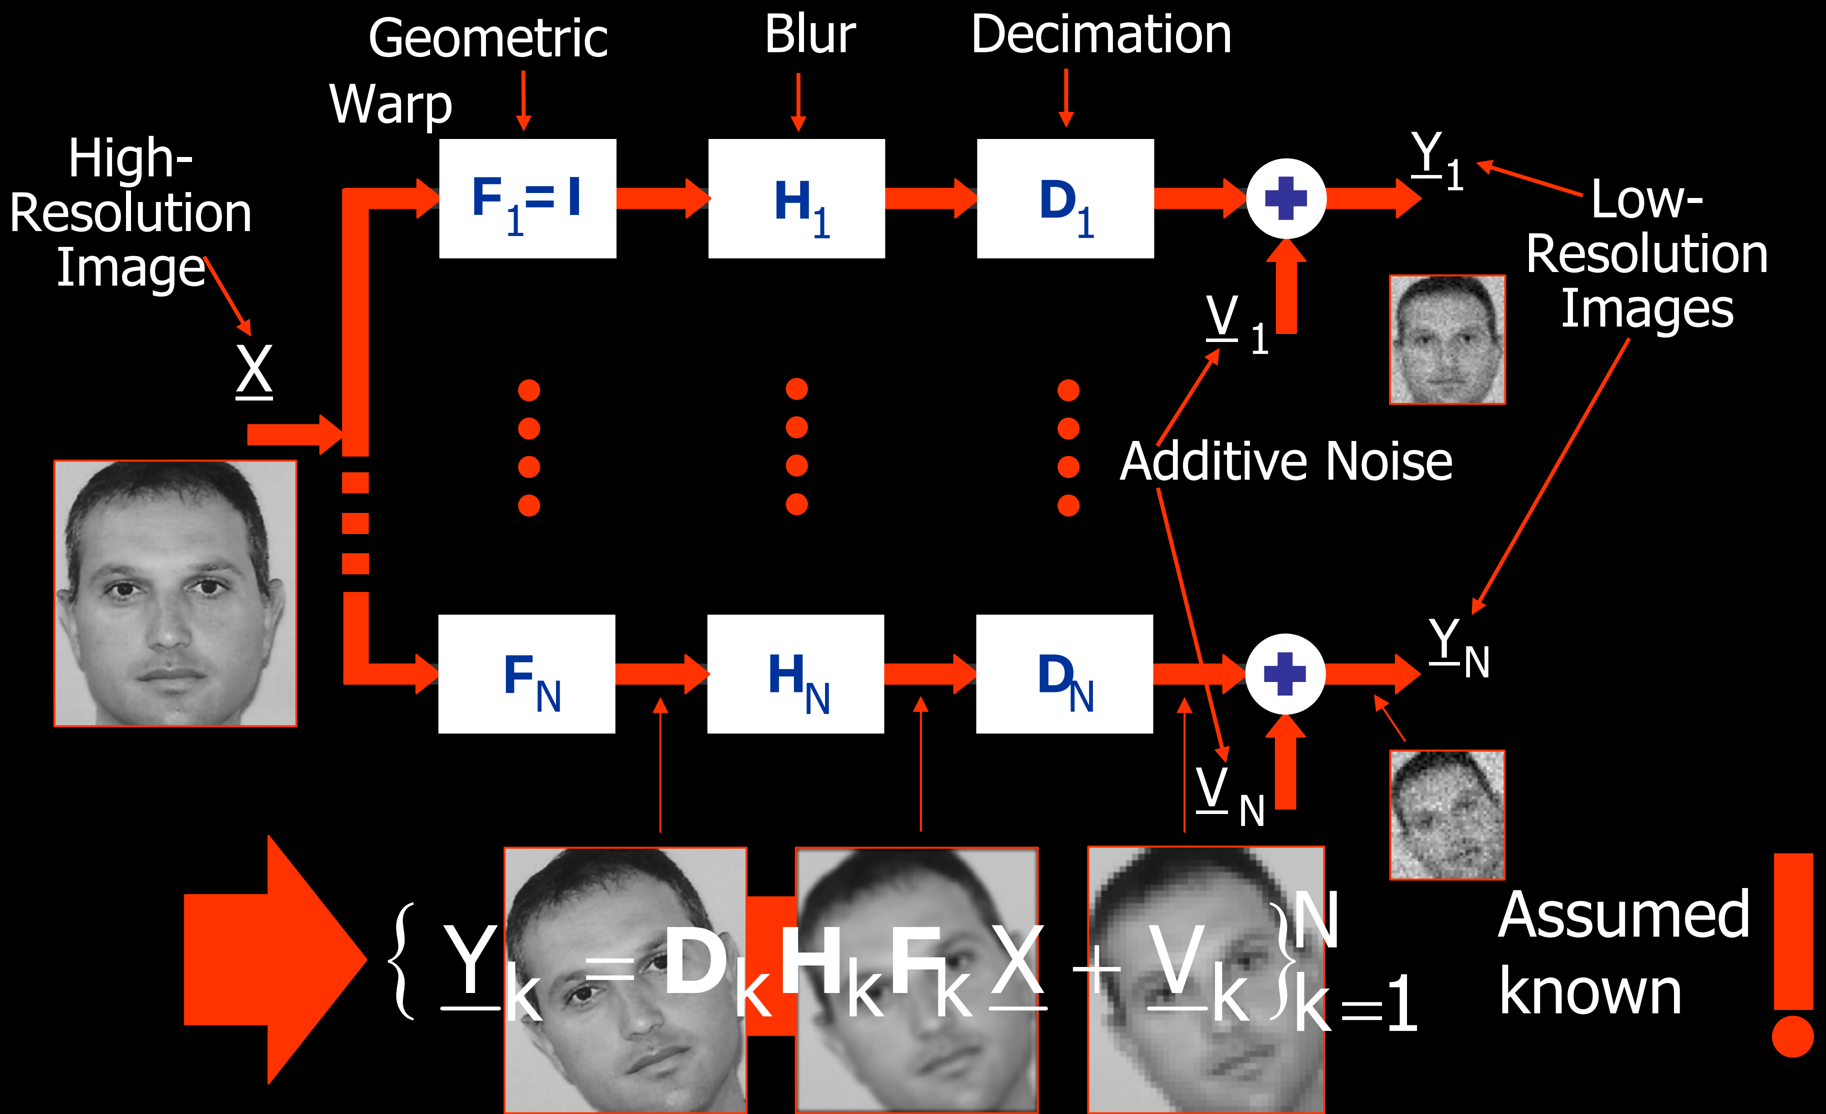

$$\left\{ \underline{Y}_k = \mathbf{D}_k \mathbf{H}_k \mathbf{F}_k \underline{X} + \underline{V}_k \right\}_{k=1}^{N}$$

The Model as One Equation

$$\left\{ \underline{Y}_k = \mathbf{D}_k \mathbf{H}_k \mathbf{F}_k \underline{X} + \underline{V}_k \right\}_{k=1}^{N}$$

$$\underline{Y} = \begin{bmatrix} \underline{Y}_1 \\ \underline{Y}_2 \\ \vdots \\ \underline{Y}_N \end{bmatrix} = \begin{bmatrix} \mathbf{D}_1 \mathbf{H}_1 \mathbf{F}_1 \\ \mathbf{D}_2 \mathbf{H}_2 \mathbf{F}_2 \\ \vdots \\ \mathbf{D}_N \mathbf{H}_N \mathbf{F}_N \end{bmatrix} \underline{X} + \begin{bmatrix} \underline{V}_1 \\ \underline{V}_2 \\ \vdots \\ \underline{V}_N \end{bmatrix} = \mathbf{H}\underline{X} + \underline{V}$$

A Thumb Rule

In the noiseless case we have

$$\underline{Y} = \begin{bmatrix} \underline{Y}_1 \\ \underline{Y}_2 \\ \vdots \\ \underline{Y}_N \end{bmatrix} = \begin{bmatrix} \mathbf{D}_1\mathbf{H}_1\mathbf{F}_1 \\ \mathbf{D}_2\mathbf{H}_2\mathbf{F}_2 \\ \vdots \\ \mathbf{D}_N\mathbf{H}_N\mathbf{F}_N \end{bmatrix} \underline{X} = \mathbf{H}\underline{X}$$

Clearly, this linear system of equations should have more **equations** than **unknowns** in order to make it possible to have a unique Least-Squares solution.

Example: Assume that we have N images of 100-by-100 pixels, and we would like to produce an image $\underline{X}$ of size 300-by-300. Then, we should require $N \geq 9$.

Chapter 2: The Maximum-Likelihood Solution

The Maximum-Likelihood Approach

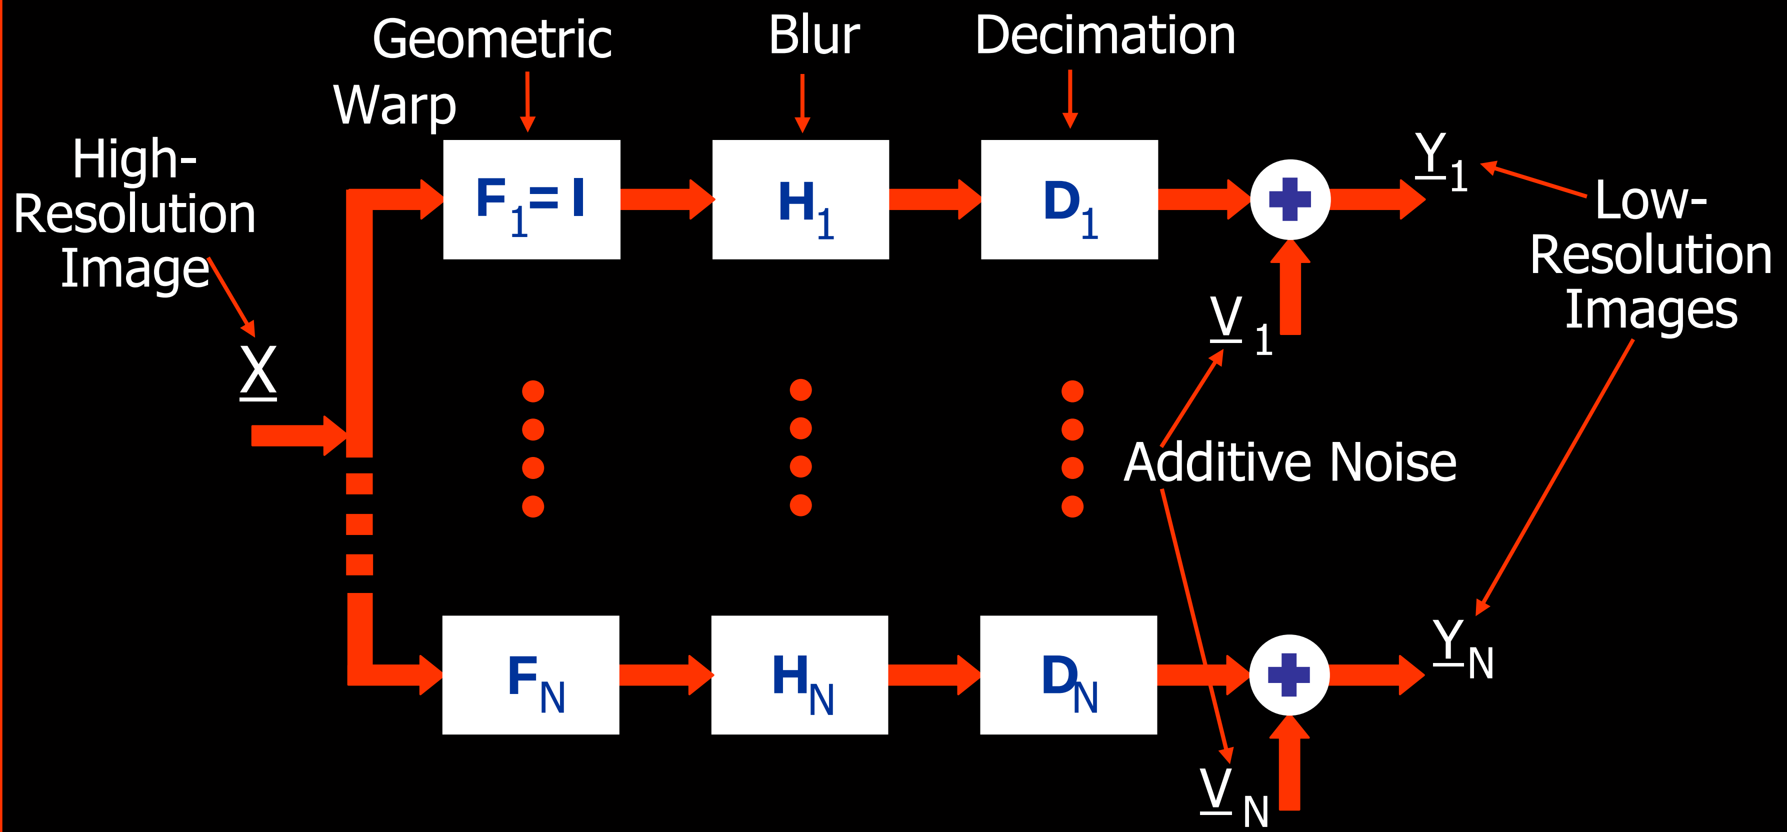

Which $\underline{X}$ would be such that when fed to the above system it yields a set $\underline{Y}_k$ closest to the measured images **?**

ML Reconstruction

Minimize:

$$\varepsilon_{ML}^2(\underline{X}) = \sum_{k=1}^{N} \left\| \underline{Y}_k - \mathbf{D}_k \mathbf{H}_k \mathbf{F}_k \underline{X} \right\|^2$$

$$= \left\| \underline{Y} - \mathbf{H}\underline{X} \right\|^2$$

Thus, require:

$$\frac{\partial \varepsilon_{ML}^2(\underline{X})}{\partial \underline{X}} = 0 \Rightarrow \mathbf{H}^T \mathbf{H} \hat{\underline{X}} = \mathbf{H}^T \underline{Y}$$

A Numerical Solution

$$\mathbf{H}^T\mathbf{H}\underline{\hat{X}} = \mathbf{H}^T\underline{Y}$$

- This is a (huge !!!) linear system of equations with #equations and unknowns = #of desired pixels (e.g. 10^6).

- This system of equations is solved iteratively using classic optimization techniques. Surprisingly, 10-15 simple iterations (CG or even SD) are sufficient in most cases.

- In case $\mathbf{H}^T\mathbf{H}$ is non-invertible (insufficient data), it means that no unique solution exists.

Chapter 3:
Bayesian Super-Resolution Reconstruction

The Model – A Statistical View

$$\underline{Y} = \begin{bmatrix} \underline{Y}_1 \\ \underline{Y}_2 \\ \vdots \\ \underline{Y}_N \end{bmatrix} = \begin{bmatrix} \mathbf{D}_1\mathbf{H}_1\mathbf{F}_1 \\ \mathbf{D}_2\mathbf{H}_2\mathbf{F}_2 \\ \vdots \\ \mathbf{D}_N\mathbf{H}_N\mathbf{F}_N \end{bmatrix} \underline{X} + \begin{bmatrix} \underline{V}_1 \\ \underline{V}_2 \\ \vdots \\ \underline{V}_N \end{bmatrix} = \mathbf{H}\underline{X} + \underline{V}$$

We assume that the noise vector, $\underline{V}$, is Gaussian and white.

$$\mathrm{Prob}\{\underline{V}\} = \mathrm{Const} \cdot \exp\left\{ -\frac{\underline{V}^T \underline{V}}{2\sigma_v^2} \right\}$$

For a known $\underline{X}$, $\underline{Y}$ is also Gaussian with a "shifted mean"

$$\mathrm{Prob}\{\underline{Y} \mid \underline{X}\} = \mathrm{Const} \cdot \exp\left\{ -\frac{(\underline{Y} - \mathbf{H}\underline{X})^T (\underline{Y} - \mathbf{H}\underline{X})}{2\sigma_v^2} \right\}$$

Maximum-Likelihood ... Again

The ML estimator is given by

$$\underline{\hat{X}}_{ML} = \underset{\underline{X}}{\mathrm{ArgMax}}\,\mathrm{Prob}\{\underline{Y} \mid \underline{X}\}$$

which means: Find the image $\underline{X}$ such that the measurements are the most likely to have happened.

In our case this leads to what we have seen before

$$\underline{\hat{X}}_{ML} = \underset{\underline{X}}{\mathrm{ArgMax}}\,\mathrm{Prob}\{\underline{Y} \mid \underline{X}\} = \underset{\underline{X}}{\mathrm{ArgMin}}\left\| \mathbf{H}\underline{X} - \underline{Y} \right\|^2$$

ML Often Sucks !!! For Example ...

For the image denoising problem we get

$$\hat{\underline{X}}_{ML} = \underset{\underline{X}}{\text{ArgMin}} \left\| \underline{X} - \underline{Y} \right\|^2 \Rightarrow \hat{\underline{X}} = \underline{Y}$$

We got that the best ML estimate for a noisy image is ... the noisy image itself.

The ML estimator is quite useless, when we have insufficient information. A better approach is needed. The solution is **the Bayesian approach**.

Using The Posterior

Instead of maximizing the Likelihood function

$$\mathrm{Prob}\left\{\underline{Y} \mid \underline{X}\right\}$$

maximize the Posterior probability function

$$\mathrm{Prob}\left\{\underline{X} \mid \underline{Y}\right\}$$

This is the Maximum-Aposteriori Probability (MAP) estimator:
Find the most probable X, given the measurements

A major conceptual change – $\underline{X}$ is assumed to be random

Why Called Bayesian?

Bayes formula states that

$$\text{Prob}\{\underline{X}|\underline{Y}\} = \frac{\text{Prob}\{\underline{Y}|\underline{X}\}\text{Prob}\{\underline{X}\}}{\text{Prob}\{\underline{Y}\}}$$

and thus MAP estimate leads to

$$\hat{\underline{X}}_{MAP} = \underset{\underline{X}}{\text{ArgMax}}\,\text{Prob}\{\underline{X}|\underline{Y}\} = \underset{\underline{X}}{\text{ArgMax}}\,\boxed{\text{Prob}\{\underline{Y}|\underline{X}\}}\,\boxed{\text{Prob}\{\underline{X}\}}$$

This part is already known

What shall it be?

Image Priors?

$$\text{Prob}\{\underline{X}\} = ?$$

- This is the probability law of images. How can we describe it in a relatively simple expression?
- Much of the progress made in image processing in the past 20 years (PDE's in image processing, wavelets, MRF, advanced transforms, and more) can be attributed to the answers given to this question.

MAP Reconstruction

If we assume the **Gibbs** distribution with some energy function $A(\underline{X})$ for the prior, we have

$$\text{Prob}\{\underline{X}\} = \text{Const} \cdot \exp\{-A\{\underline{X}\}\}$$

$$\hat{\underline{X}}_{MAP} = \underset{\underline{X}}{\text{ArgMax}}\ \text{Prob}\{\underline{Y} | \underline{X}\}\ \text{Prob}\{\underline{X}\}$$

$$= \underset{\underline{X}}{\text{ArgMin}} \left\| \mathbf{H}\underline{Y} - \underline{X} \right\|^2 + \boxed{\lambda A\{\underline{X}\}}$$

This additional term is also known as regularization

Choice of Regularization

$$\varepsilon^2_{MAP}(\underline{X}) = \sum_{k=1}^{N} \left\| \underline{Y}_k - \mathbf{D}_k \mathbf{H}_k \mathbf{F}_k \underline{X} \right\|^2 + \lambda A\{\underline{X}\}$$

Possible Prior functions - Examples:

1. $A\{\underline{X}\} = \left\| \mathbf{S}\underline{X} \right\|^2$ - simple smoothness (Wiener filtering),
2. $A\{\underline{X}\} = \underline{X}^T \mathbf{S}^T \mathbf{W}(\underline{X}_0) \mathbf{S}\underline{X}$ - spatially adaptive smoothing,
3. $A\{\underline{X}\} = \rho\{\mathbf{S}\underline{X}\}$ - M-estimator (robust functions),

4. The bilateral prior – the one used in our recent work:

$$A\{\underline{X}\} = \sum_{n=-P}^{P} \sum_{m=-P}^{P} a_{mn} \cdot \rho\left(\underline{X} - \mathbf{S}_h^n \mathbf{S}_v^m \underline{X}\right)$$

4. Other options: Total Variation, Beltrami flow, example-based, sparse representations, …

Chapter 4: Some Results and Variations

The Super-Resolution Process

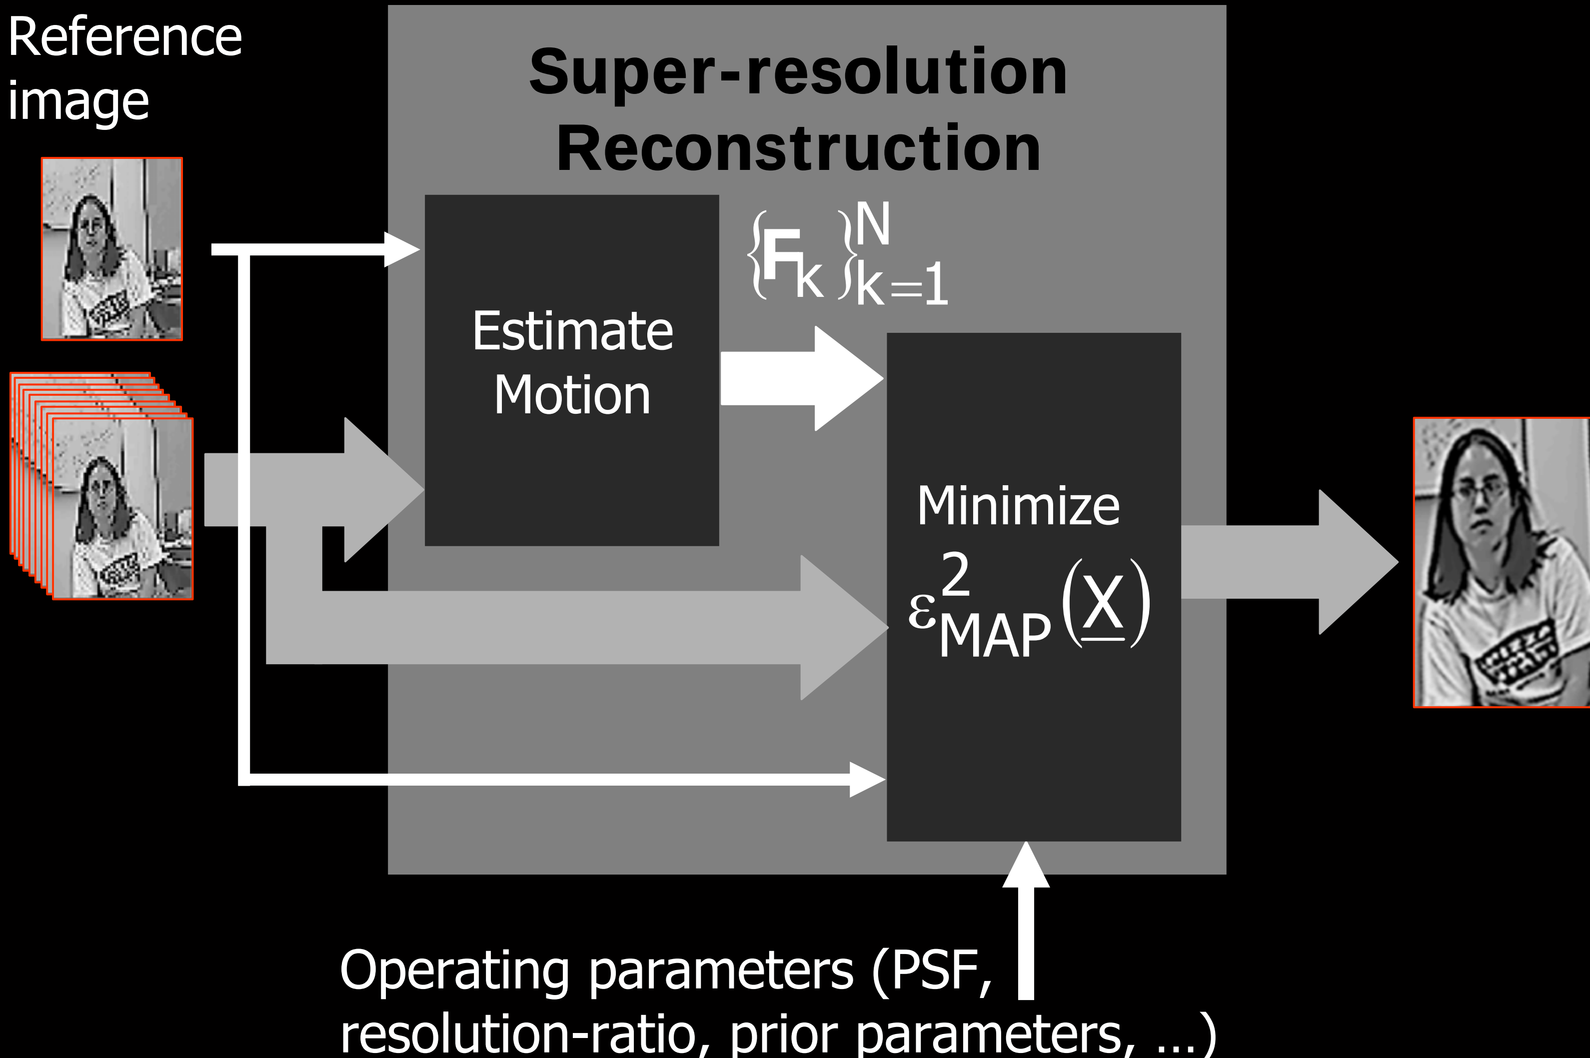

Example 0 – Sanity Check

Synthetic case:

9 images, no blur, 1:3 ratio

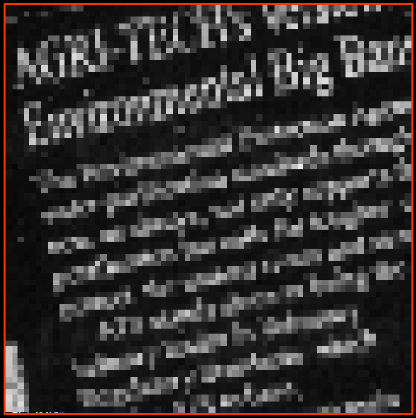

One of the low-resolution images

The higher resolution original

The reconstructed result

Example 1 – SR for Scanners

16 scanned images, ratio 1:2

Taken from one of the given images

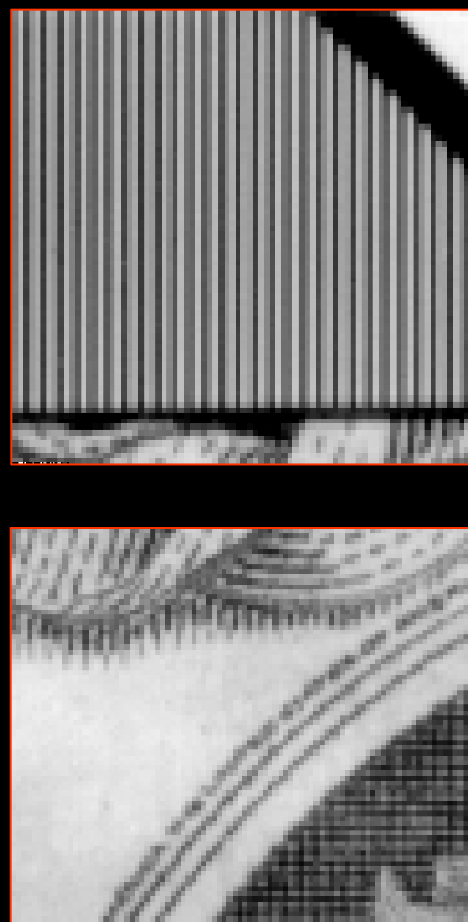

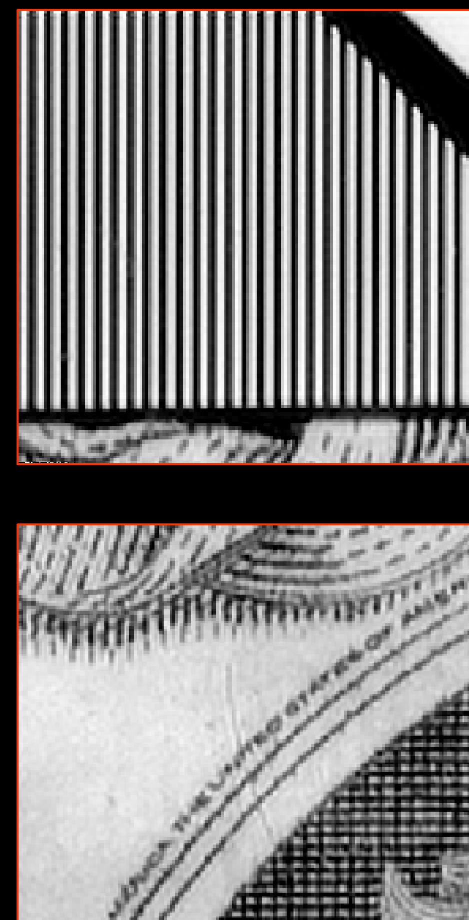

Taken from the reconstructed result

Example 2 – SR for IR Imaging

8 images*, ratio 1:4

* This data is courtesy of the US Air Force

Example 3 – Surveillance

40 images
ratio 1:4

Robust SR

$$\varepsilon^2_{MAP}(\underline{X}) = \sum_{k=1}^{N} \left\| \underline{Y}_k - \mathbf{D}_k \mathbf{H}_k \mathbf{F}_k \underline{X} \right\|^2 + \lambda A\{\underline{X}\}$$

Cases of measurements outlier:

- Some of the images are irrelevant,
- Error in motion estimation,
- Error in the blur function, or
- General model mismatch.

$$\Rightarrow \quad \varepsilon^2_{MAP}(\underline{X}) = \sum_{k=1}^{N} \left\| \underline{Y}_k - \mathbf{D}_k \mathbf{H}_k \mathbf{F}_k \underline{X} \right\|_1 + \lambda A\{\underline{X}\}$$

Example 4 – Robust SR

20 images, ratio 1:4

L_2 norm based

L_1 norm based

Example 5 – Robust SR

20 images, ratio 1:4

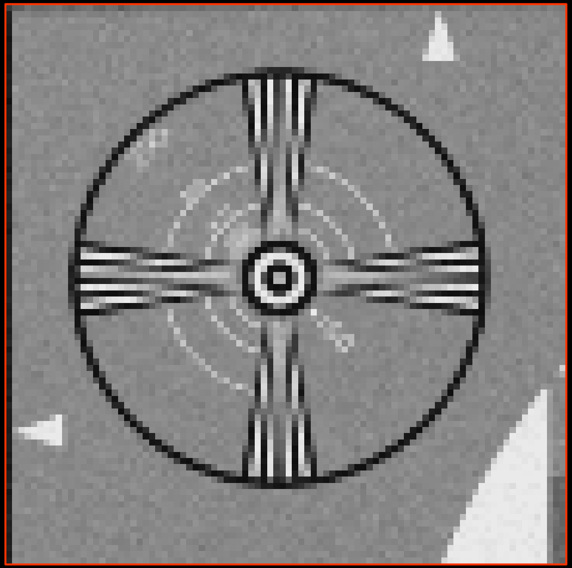

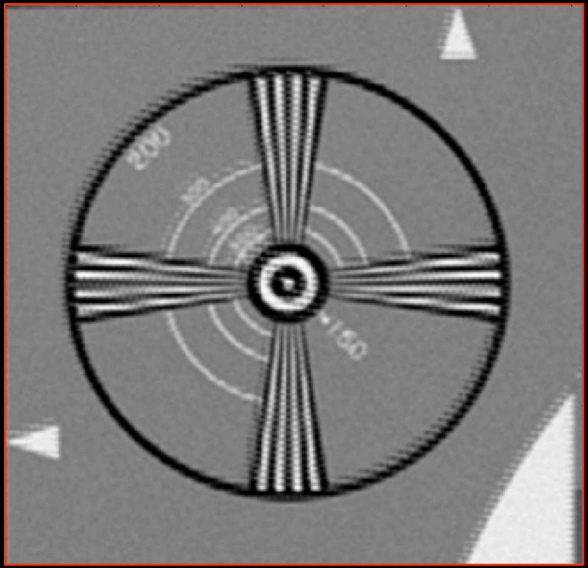

L_2 norm based

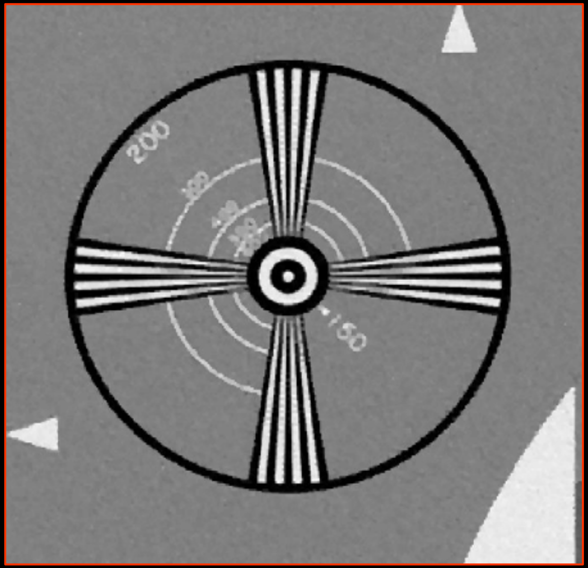

L_1 norm based

Handling Color in SR

$$\varepsilon^2_{MAP}(\underline{X}) = \sum_{k=1}^{N} \| \underline{Y}_k - \mathbf{D}_k\mathbf{H}_k\mathbf{F}_k\underline{X}\|^2 + \lambda A\{\underline{X}\}$$

- Handling color: the classic approach is to convert the measurements to YCbCr, apply the SR on the Y and use trivial interpolation on the Cb and Cr.
- Better treatment can be obtained if the statistical dependencies between the color layers are taken into account (i.e. forming a prior for color images).
- In case of mosaiced measurements, demosaicing followed by SR is sub-optimal. An algorithm that directly fuse the mosaic information to the SR is better.

Example 6 – SR for Full Color

20 images, ratio 1:4

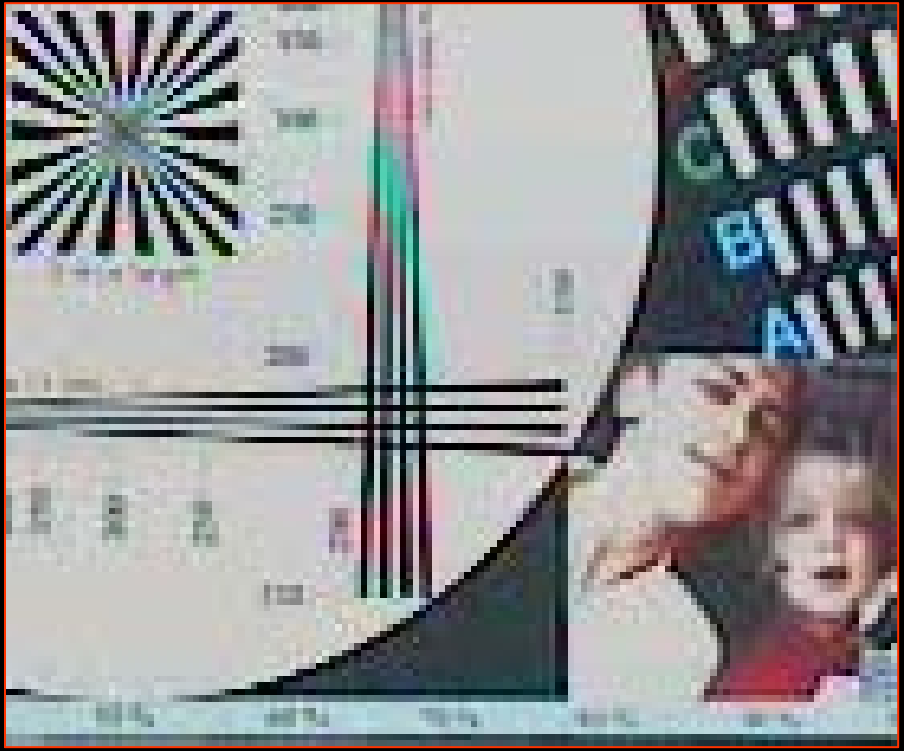

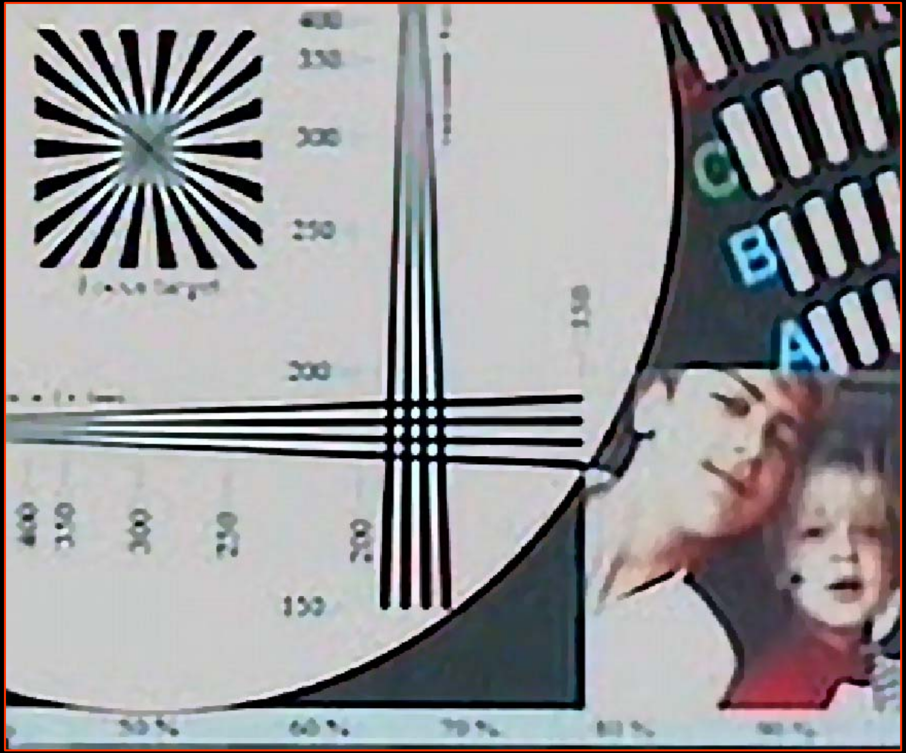

Example 7 – SR+Demoaicing

20 images, ratio 1:4

Mosaiced input

Mosaicing and then SR

Combined treatment

Chapter 5: Super-Resolution: A Summary

To Conclude

- SR reconstruction is possible, but … not always! (needs aliasing, accurate motion, enough frames, …).
- Accurate motion estimation remains the main bottle-neck for Super-Resolution success.
- Our recent work on robustifying the SR process, better treatment of color, and more, gives a significant step forward in the SR abilities and results.
- **The dream:** A robust SR process that operates on a set of low-quality frames, fuses them reliably, and gives an output image with quality never below the input frames, and with no strange artifacts.

➡ Unfortunately, **WE ARE NOT THERE YET**.

Our Work in this Field

1. M. Elad and A. Feuer, “Restoration of Single Super-Resolution Image From Several Blurred, Noisy and Down-Sampled Measured Images”, the IEEE Trans. on Image Processing, Vol. 6, no. 12, pp. 1646-58, December 1997.
2. M. Elad and A. Feuer, “Super-Resolution Restoration of Continuous Image Sequence - Adaptive Filtering Approach”, the IEEE Trans. on Image Processing, Vol. 8. no. 3, pp. 387-395, March 1999.
3. M. Elad and A. Feuer, “Super-Resolution reconstruction of Continuous Image Sequence”, the IEEE Trans. On Pattern Analysis and Machine Intelligence (PAMI), Vol. 21, no. 9, pp. 817-834, September 1999.
4. M. Elad and Y. Hel-Or, “A Fast Super-Resolution Reconstruction Algorithm for Pure Translational Motion and Common Space Invariant Blur”, the IEEE Trans. on Image Processing, Vol.10, No. 8, pp.1187-93, August 2001.
5. S. Farsiu, D. Robinson, M. Elad, and P. Milanfar, “Fast and Robust Multi-Frame Super-resolution”, IEEE Trans. On Image Processing, Vol. 13, No. 10, pp. 1327-1344, October 2004.
6. S. Farsiu, D. Robinson, M. Elad, and P. Milanfar, "Advanced and Challenges in Super-Resolution", the International Journal of Imaging Systems and Technology, Vol. 14, No. 2, pp. 47-57, Special Issue on high-resolution image reconstruction, August 2004.
7. S. Farsiu, M. Elad, and P. Milanfar, “Multi-Frame Demosaicing and Super-Resolution of Color Images”, IEEE Trans. on Image Processing, vol. 15, no. 1, pp. 141-159, Jan. 2006.
8. S. Farsiu, M. Elad, and P. Milanfar, "Video-to-Video Dynamic Superresolution for Grayscale and Color Sequences," EURASIP Journal of Applied Signal Processing, Special Issue on Superresolution Imaging , Volume 2006, Article ID 61859, Pages 1–15.

All, including these slides) are found in
http://www.cs.technion.ac.il/~elad

For our Matlab toolbox on Super-Resolution, see
http://www.soe.ucsc.edu/~milanfar/SR-Software.htm