When Spatially-Variant Filtering Meets Low-Rank Regularization: Exploiting Non-Local Similarity for Single Image Interpolation

Lantao Yu　　Michael Orchard

Department of Electrical and Computer Engineering
Rice University

RICE

ICIP 2019

Ubiquitous Need

iMac (Late 2013) : 1920×1080

Princess Returning Pearl (debut 1998) : 704×520

The Interpolation Problem

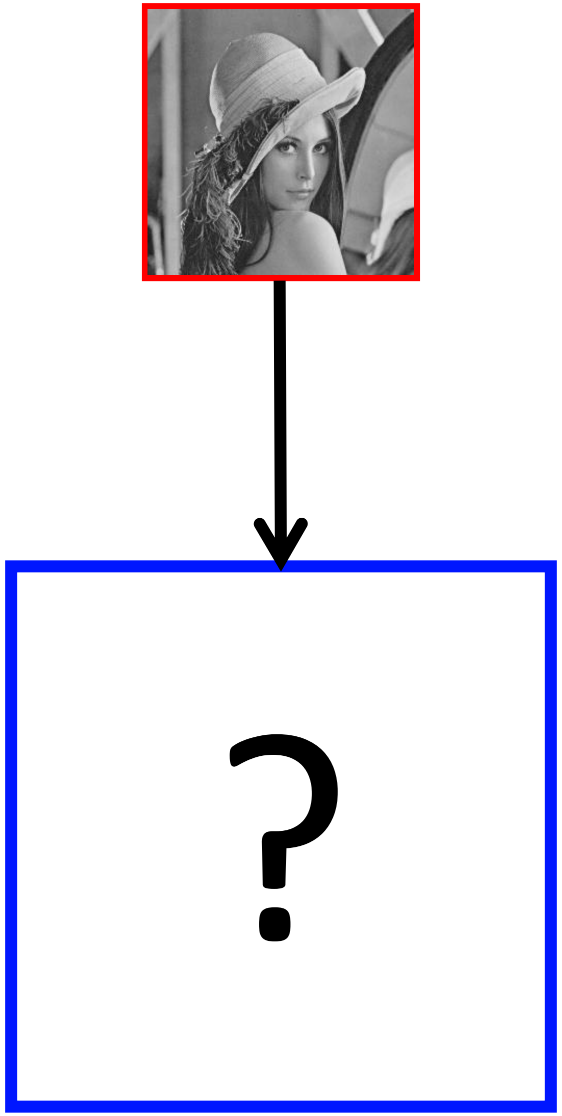

Low-Resolution (LR) image

High-Resolution (HR) image

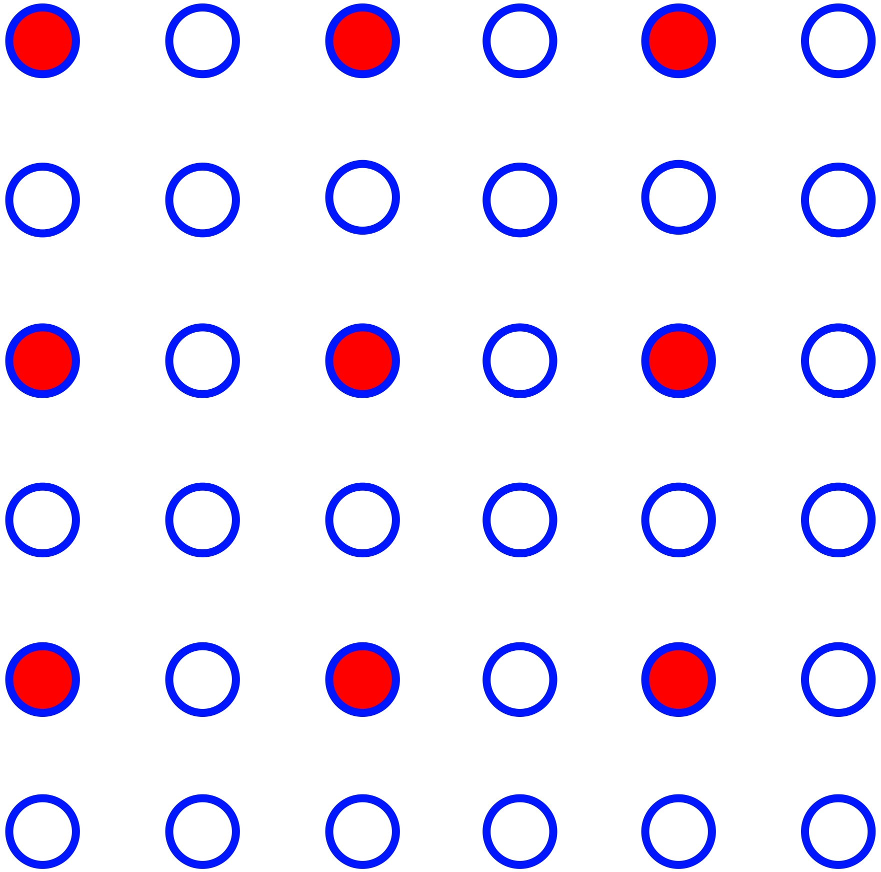

Measured Pixels: from LR image

Missing Pixels: to be estimated

Challenge

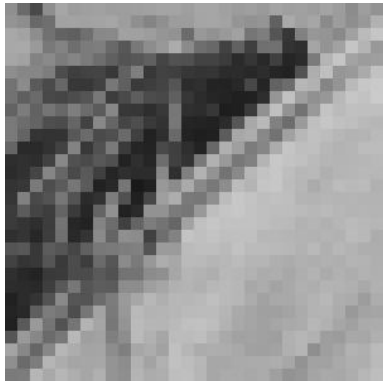

LR image

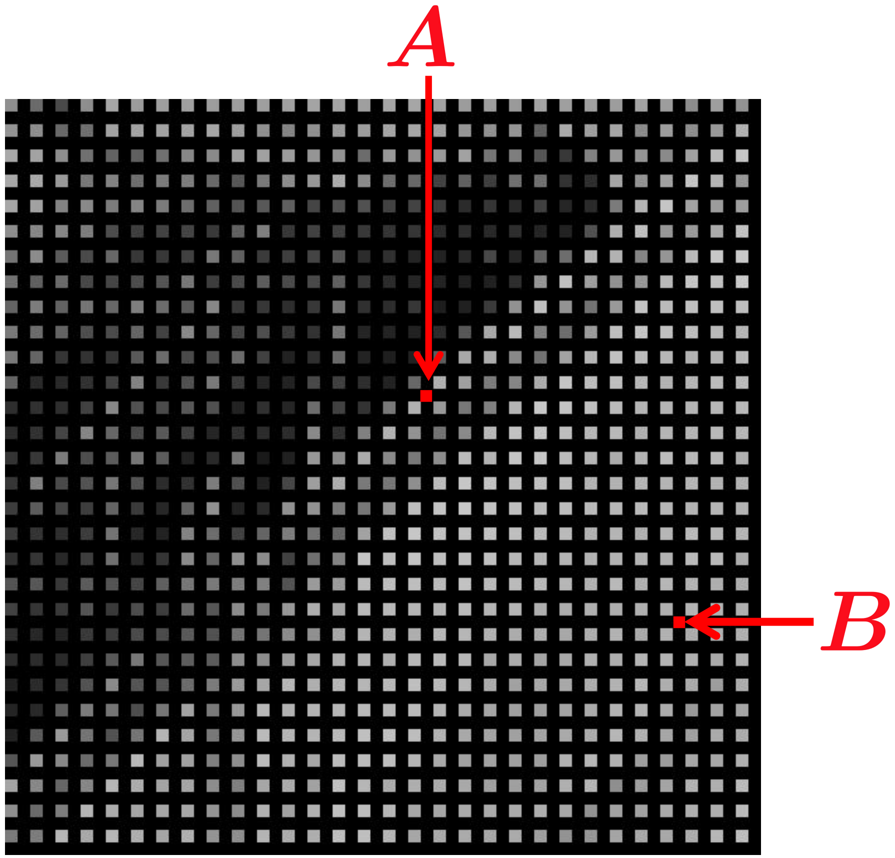

Upsampled-by-2
LR image

HR Image
(Ground-Truth)

Motivation

For pixels at different contexts, we should take different estimation strategies.

$$\hat{p}_{\mathrm{A}} = f_{\mathrm{A}}(\mathbf{N}_{\mathrm{A}})$$
$$\hat{p}_{\mathrm{B}} = f_{\mathrm{B}}(\mathbf{N}_{\mathrm{B}})$$

$\hat{p}_{\mathrm{A}}, \hat{p}_{\mathrm{B}}$: the estimation of pixel vlaues at A, B.

$\mathbf{N}_{\mathrm{A}}, \mathbf{N}_{\mathrm{B}}$: neighboring measured pixels centered around A, B.

$f_{\mathrm{A}}, f_{\mathrm{B}}$: the function of $\mathbf{N}_{\mathrm{A}}, \mathbf{N}_{\mathrm{B}}$.

Motivation

A simplistic choice of f:
linear combination of neighboring measured pixel values

[Buades et al. '05]

$$\hat{p} = \sum_i \sum_j f_{\mathbf{N}}(i,j)\mathbf{N}(i,j)$$

$\mathbf{N} \in \mathbb{R}^{W \times W}$: neighboring measured pixels surrounding the target pixel
$f_{\mathbf{N}} \in \mathbb{R}^{W \times W}$: weights of neighboring measured pixels

Image Interpolation from filtering perspective:
to learn spatially-varying filters' coefficients.

Motivation

Alternative choice of f :
linear combination of underlying structures (atoms)

[Arahon et al. '06]

$$\hat{p} = \sum_k \boldsymbol{\omega}_{\mathbf{N}}^k \boldsymbol{\sigma}_{\mathbf{N}}^k$$

$\boldsymbol{\sigma}_{\mathbf{N}}^k \in \mathbb{R}^{n^2 \times 1}$: k -th atom of size $n \times n$ stored in vectorized form.
$\boldsymbol{\omega}_{\mathbf{N}}^k \in \mathbb{R}^{1 \times n^2}$: $\hat{p}$'s corresponding weight (only one non-zero entry)

Image Interpolation from atomic perspective:
to learn atoms and their corresponding weights.

Motivation

think image interpolation
from two perspectives:

filtering	**atomic**
Simple	Complicated
Less Flexible	More Flexible
Potentially Worse Approximation	Potentially Better Approximation

Can we have an interpolation algorithm that is
both simple and promise to good approximation?

Non-Local Similarity

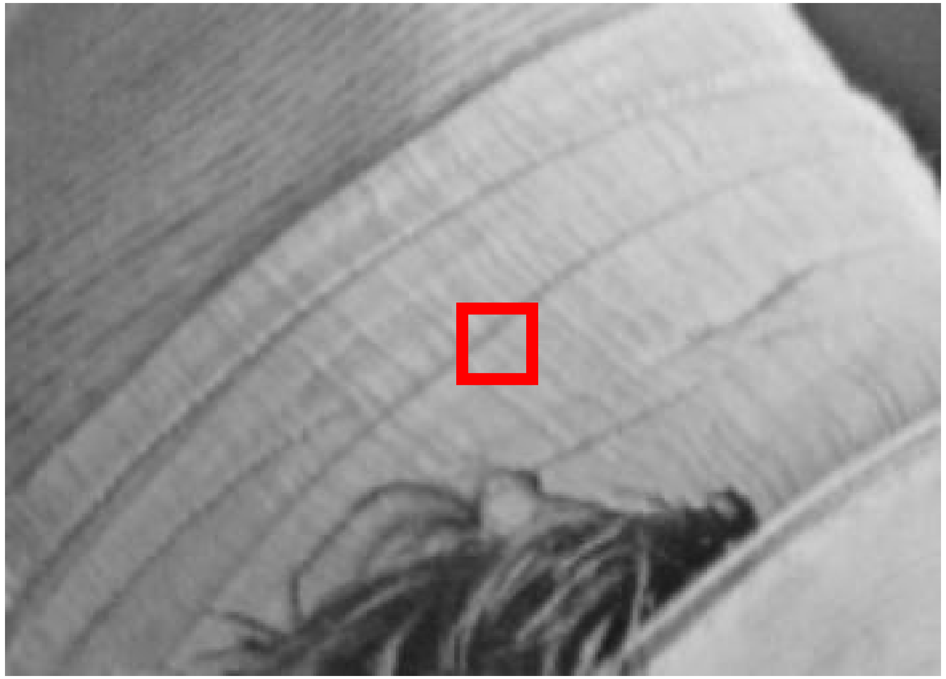

A Patch (the red block of pixels) and its Neighboring Similar Patches (the green blocks of pixels) in *Lena.*

Non-Local Similarity

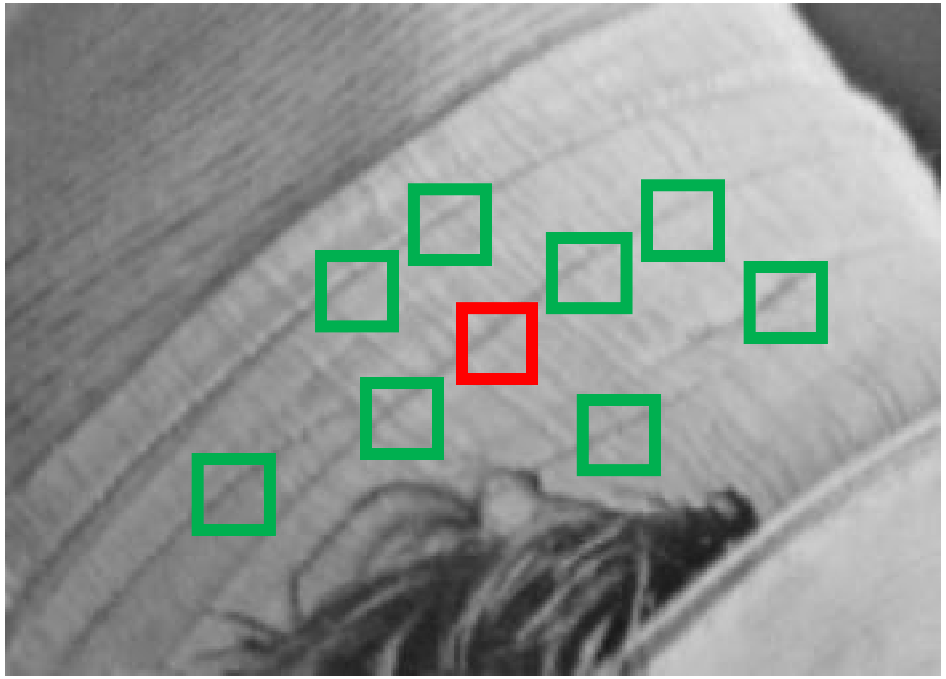

A Patch (the red block of pixels) and its Neighboring Similar Patches (the green blocks of pixels) in *Lena.*

Non-Local Similarity

$$\mathbf{P}_i \approx \sum_j \omega_j \mathbf{P}_i^j$$

An individual patch in a natural image

Similar patches searched in $\mathbf{P}_i$'s spatial neighborhood

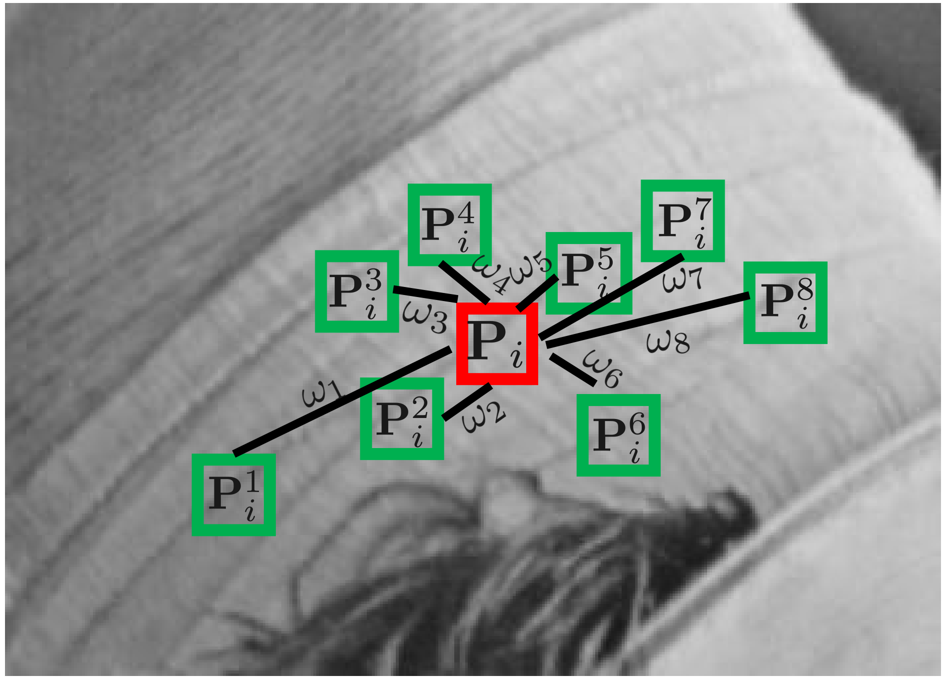

A Patch (the red block of pixels) and its Neighboring Similar Patches (the green blocks of pixels) in *Lena.*

[Buades et al. '05]
[Dong et al. '13]

Exploit Non-Local Similarity

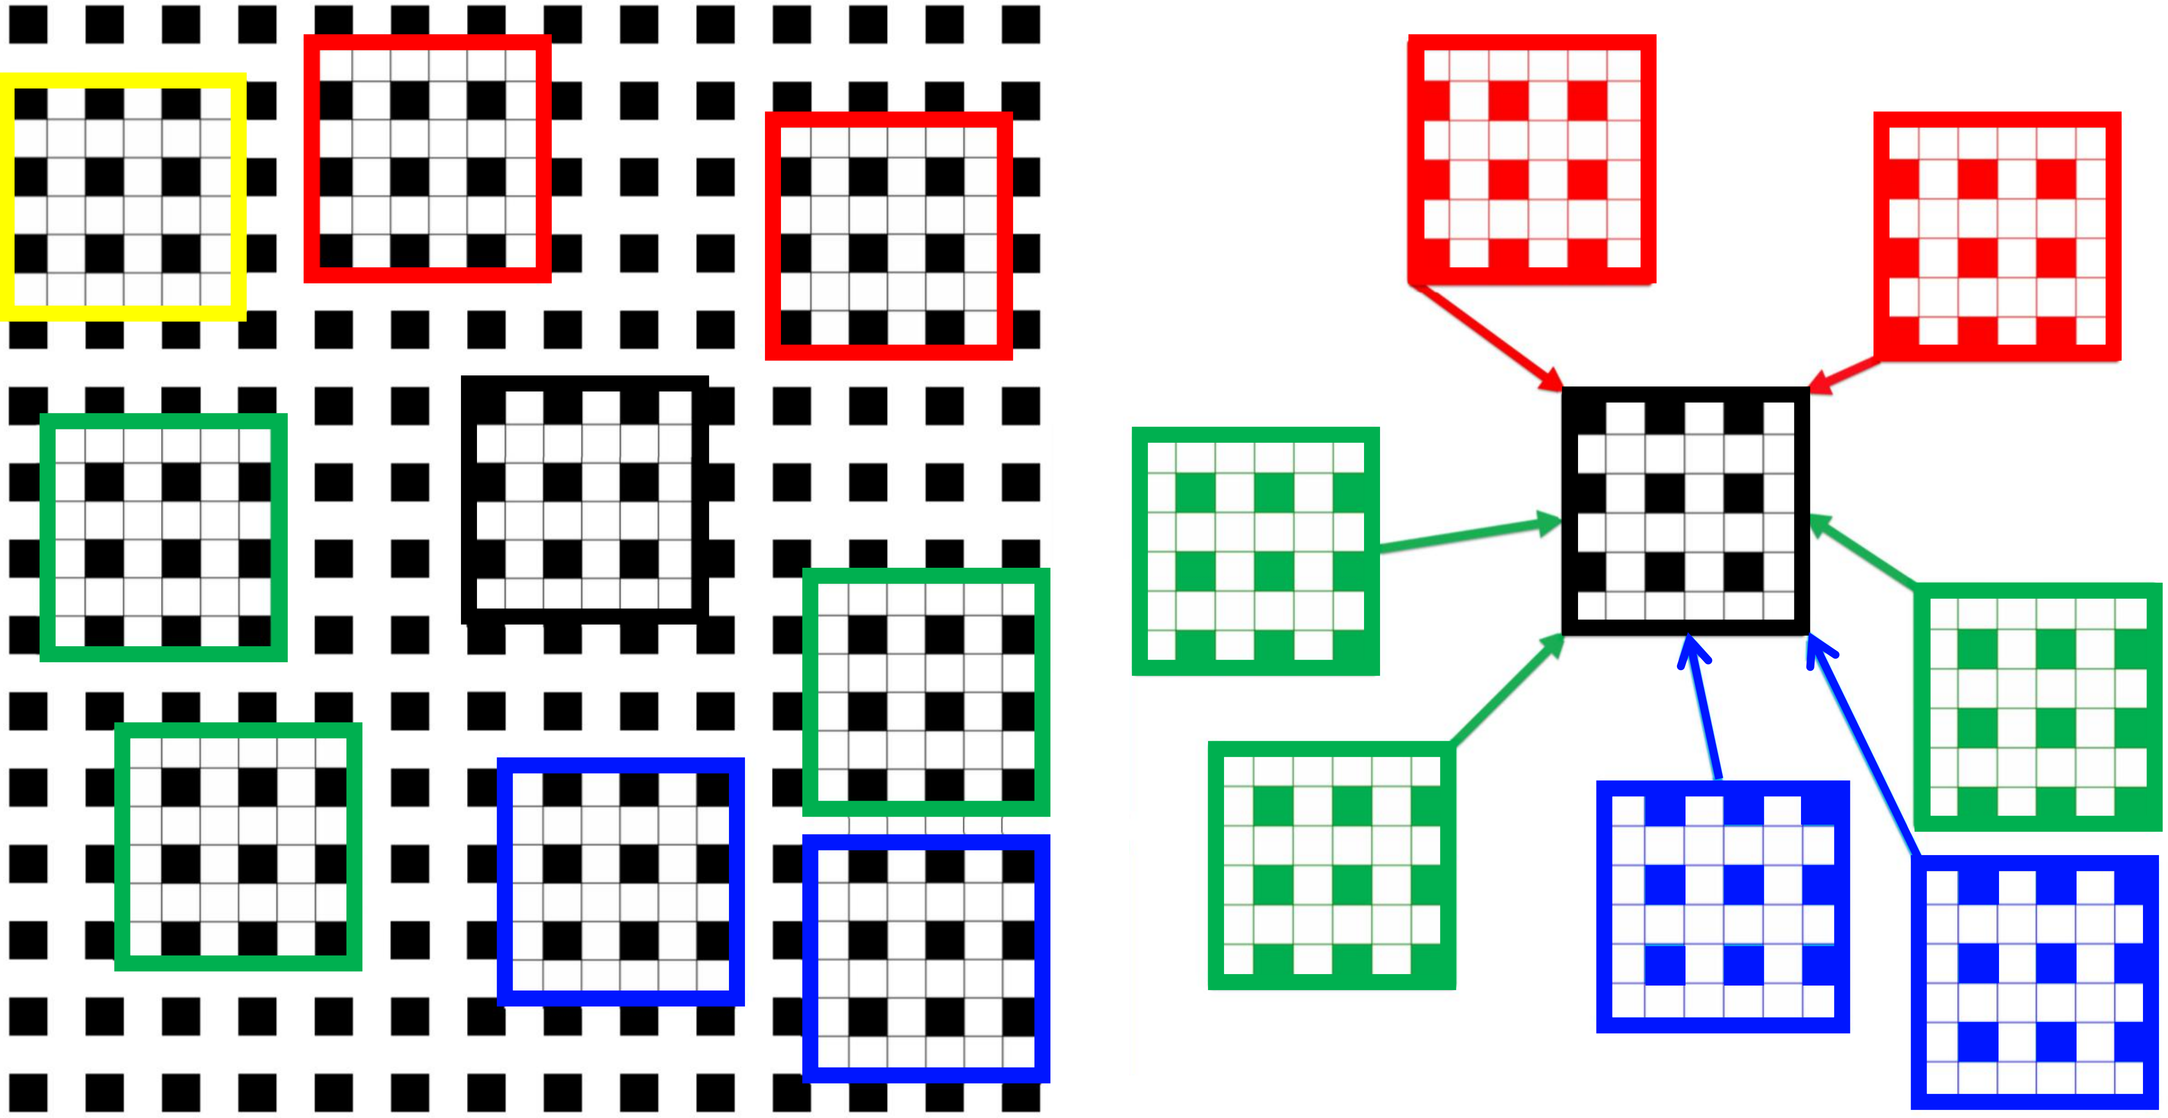

Exploit Non-Local Similarity

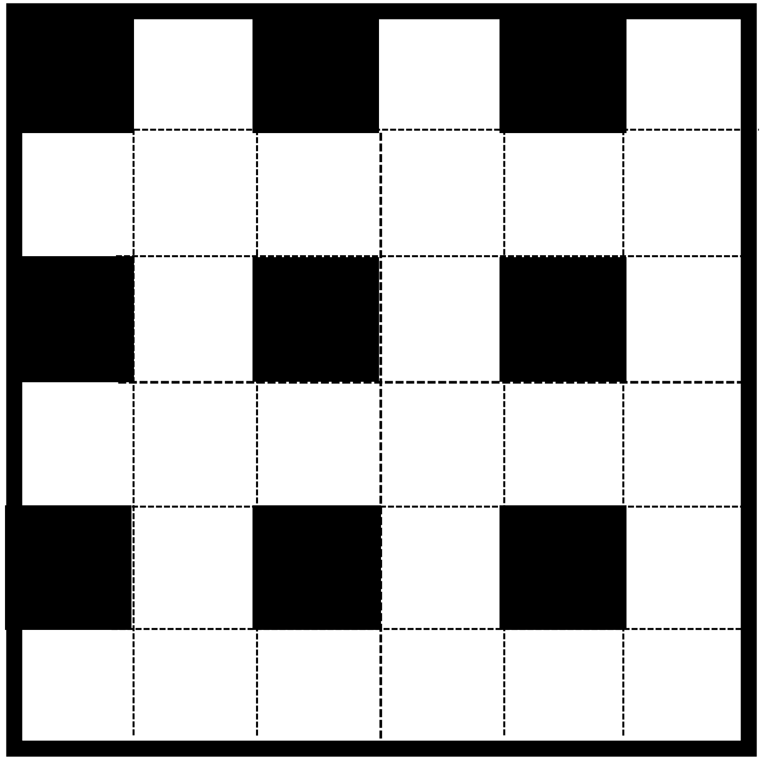

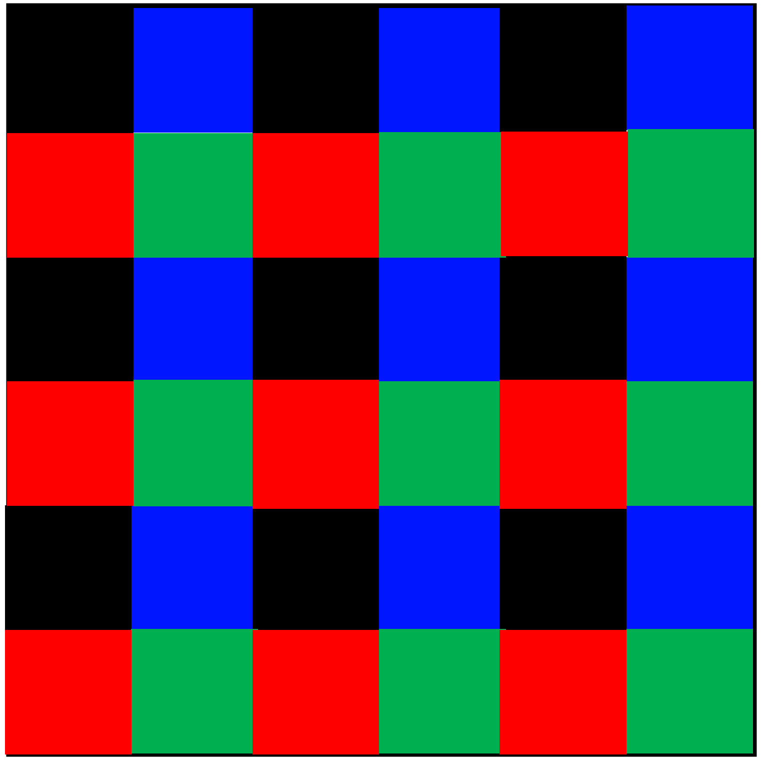

[Sun et al. '16]

Scheme of Each Iteration

- Decompose an image into overlapping patches
- Update each patch:
  - Identify the positions of similar patches (filter support)
  - Compute the weights of similar patches (non-zero filter coefficients)
- Average the contribution of overlapping patches to each missing pixel

Exploit Non-Local Similarity

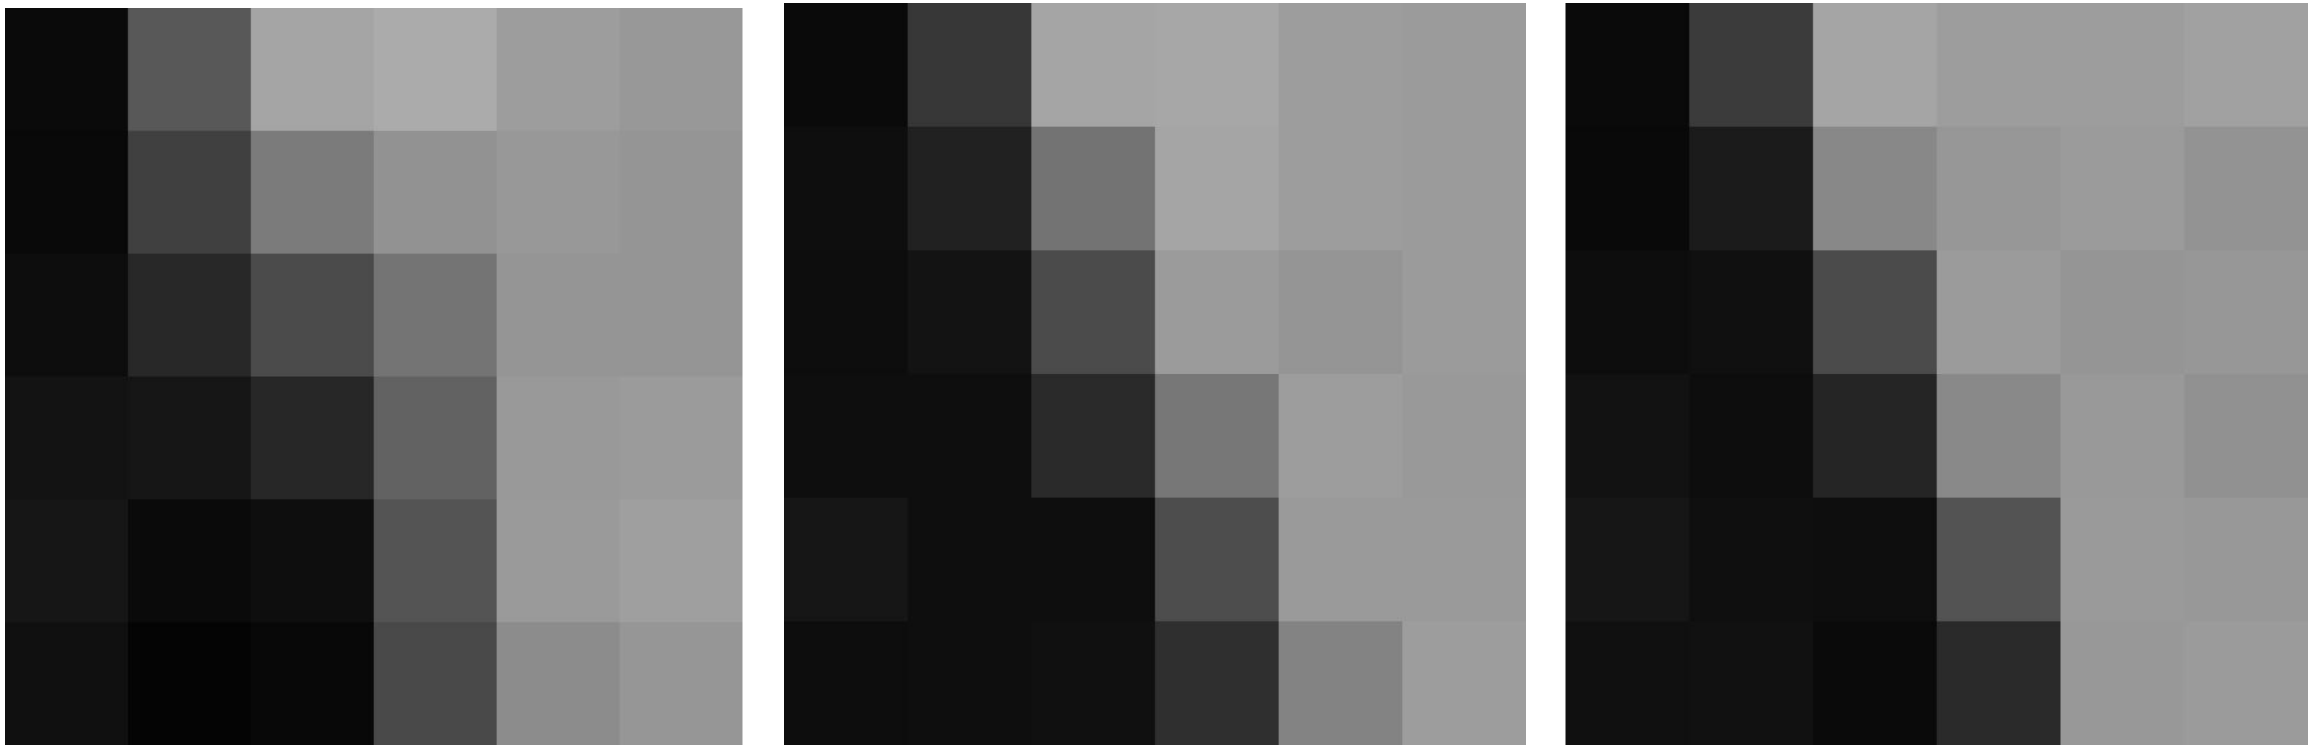

(a) initial estimate (b) after first iteration (c) ground-truth

Unique Features

- ☐ Robust initialization of the positions of similar patches
- ☐ Regularization of the weights of similar patches

Robust Position Initialization

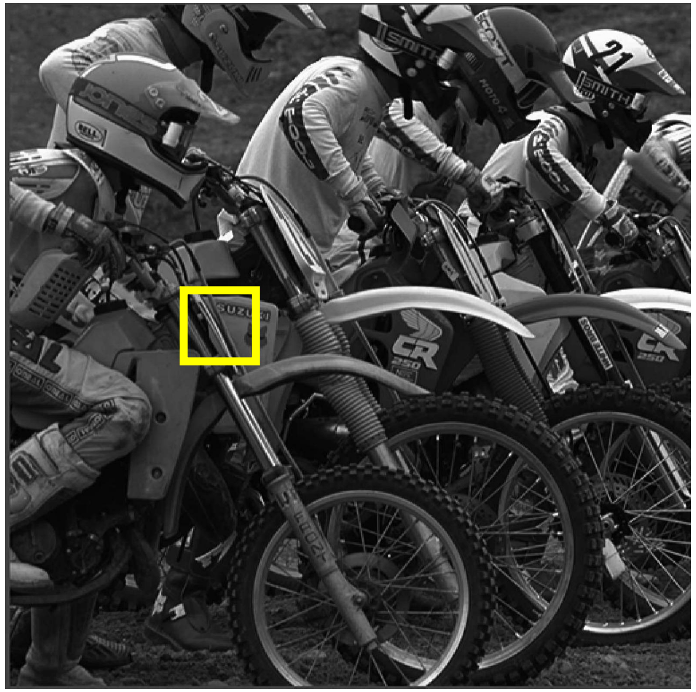

HR Image and a window of interest

Bicubic Initial Estimate

HR image

Robust Position Initialization

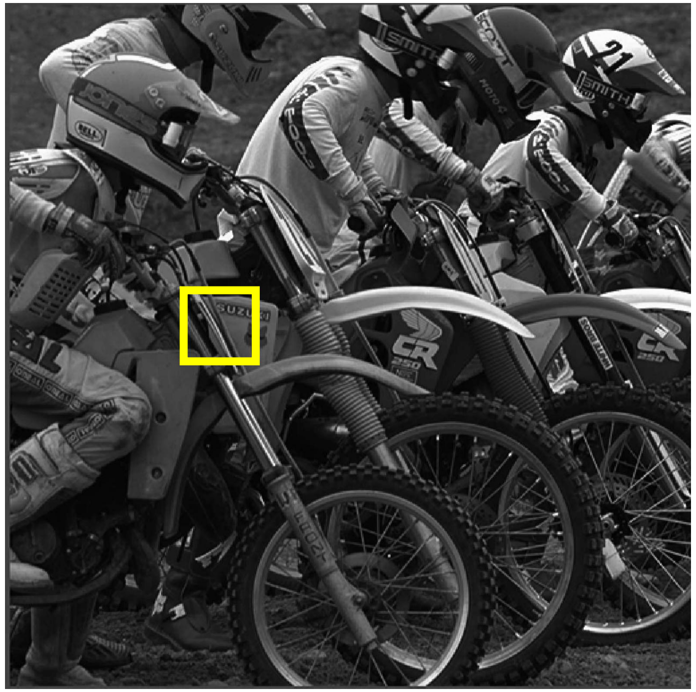

HR Image and a
window of interest

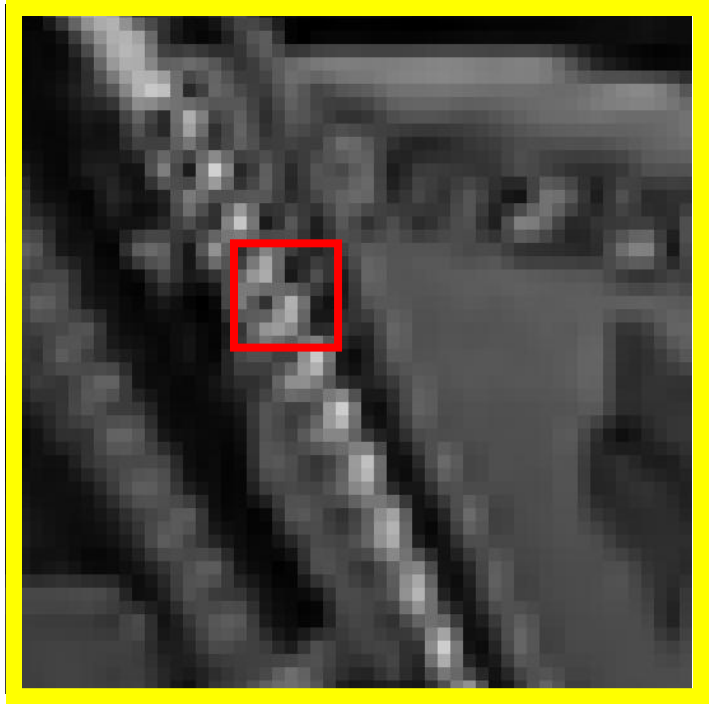

Bicubic Initial Estimate
target patch

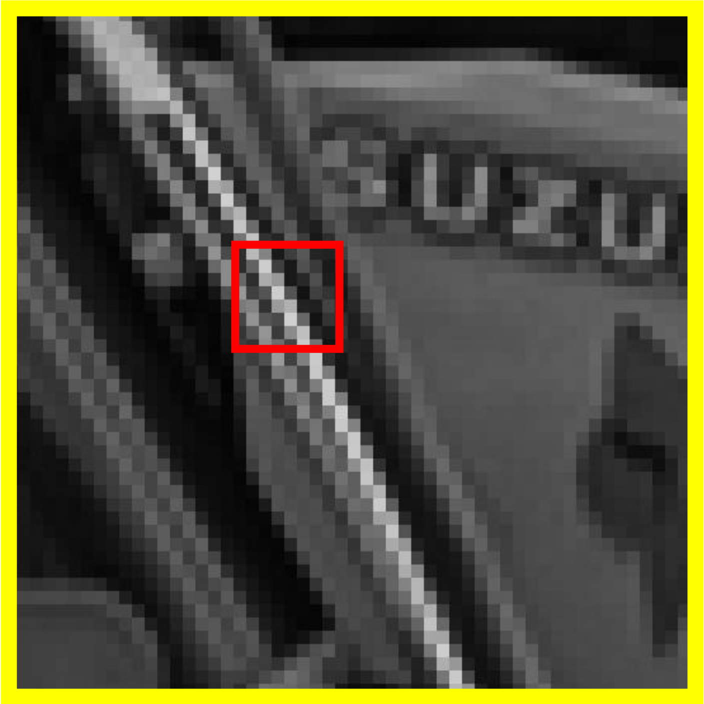

HR image
target patch

Robust Position Initialization

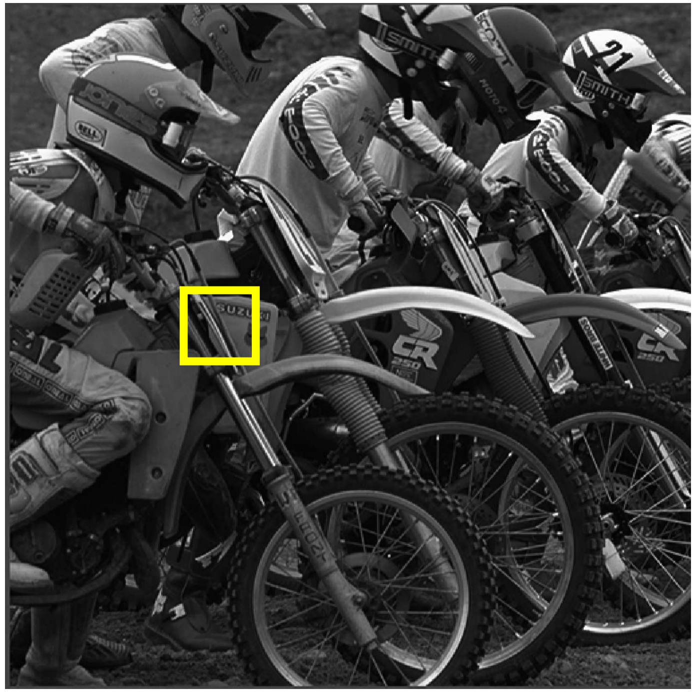

HR Image and a window of interest

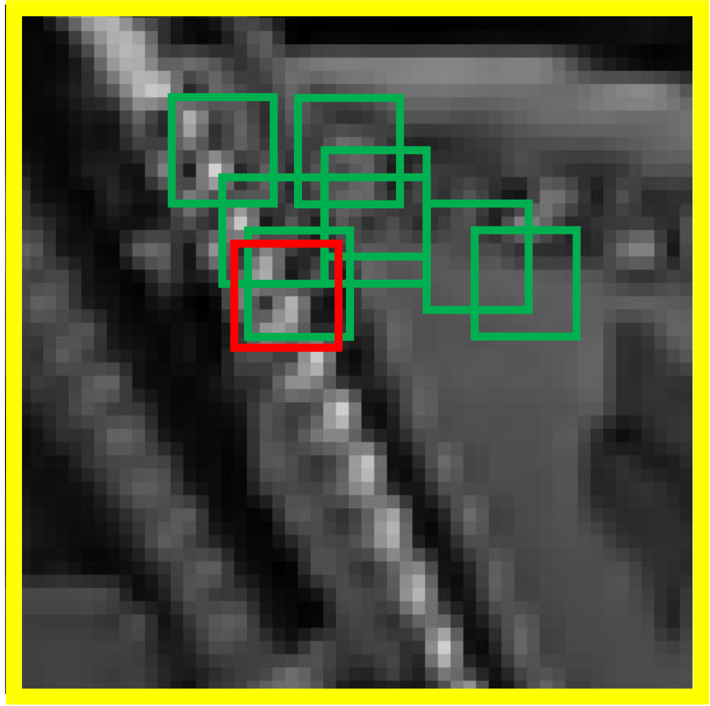

Bicubic Initial Estimate
target patch and
similar patches

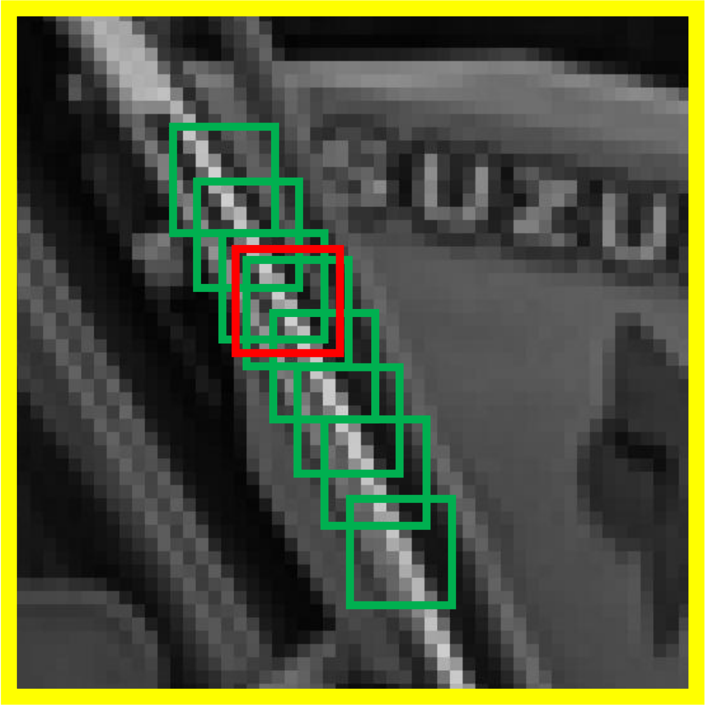

HR image
target patch and
similar patches

[Sun et al. '16]

Robust Position Initialization

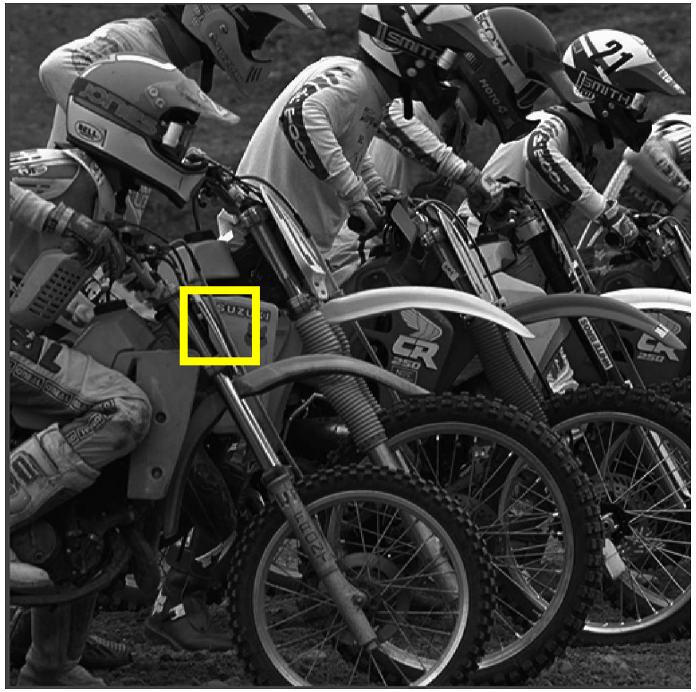

HR Image and a window of interest

Guide Image

HR image

[Yu and Orchard '19]

Robust Position Initialization

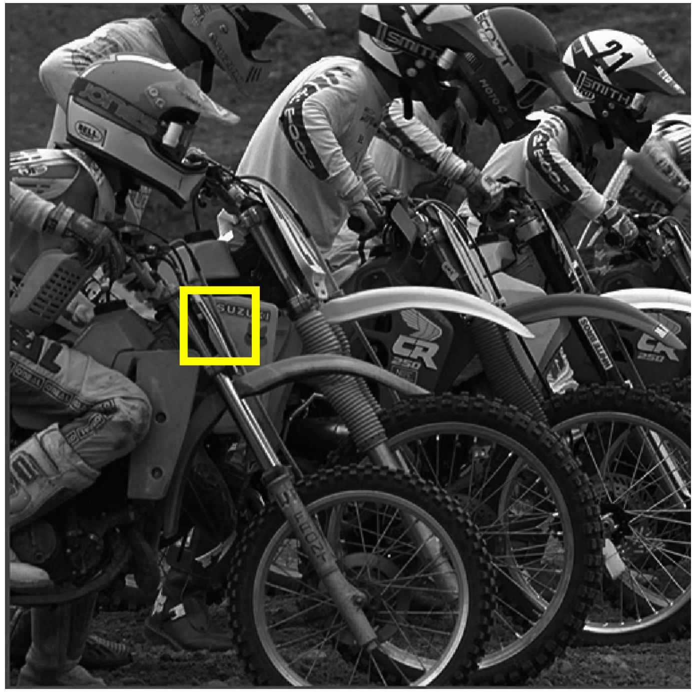

HR Image and a window of interest

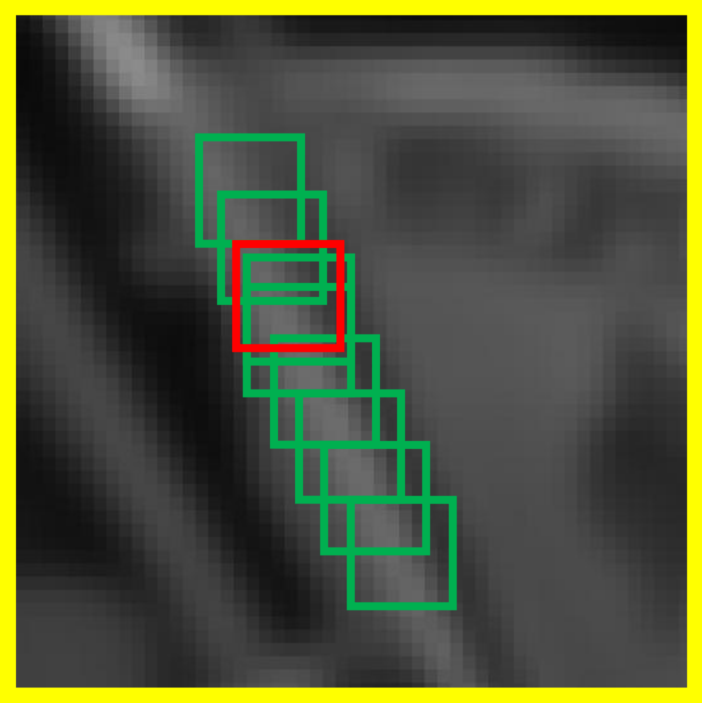

Guide Image
target patch and
similar patches

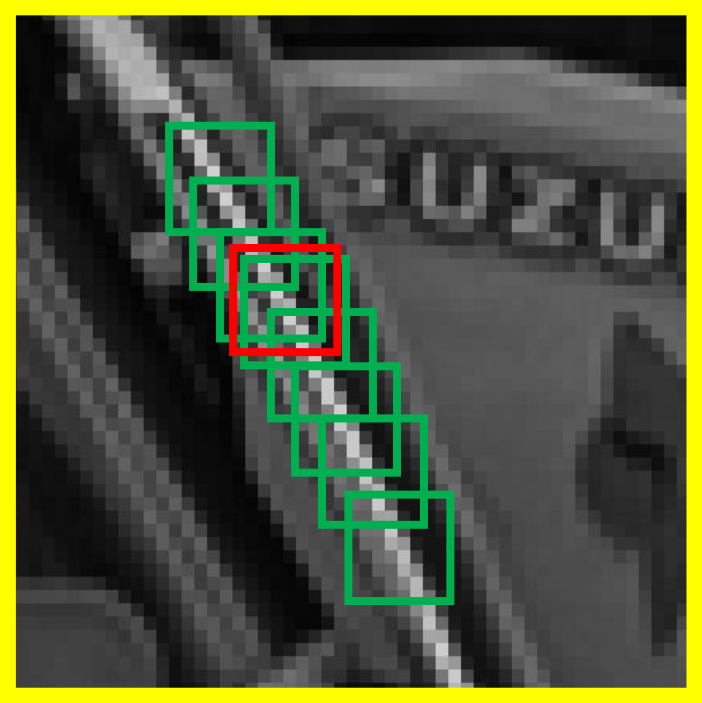

HR image
target patch and
similar patches

[Yu and Orchard '19]

Robust Position Initialization

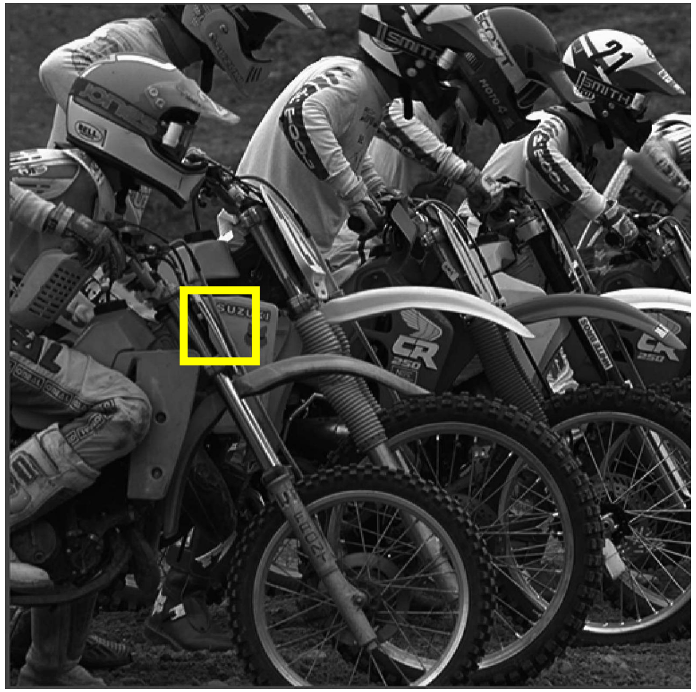

HR Image and a window of interest

Input Image in Last Iter
target patch and
similar patches

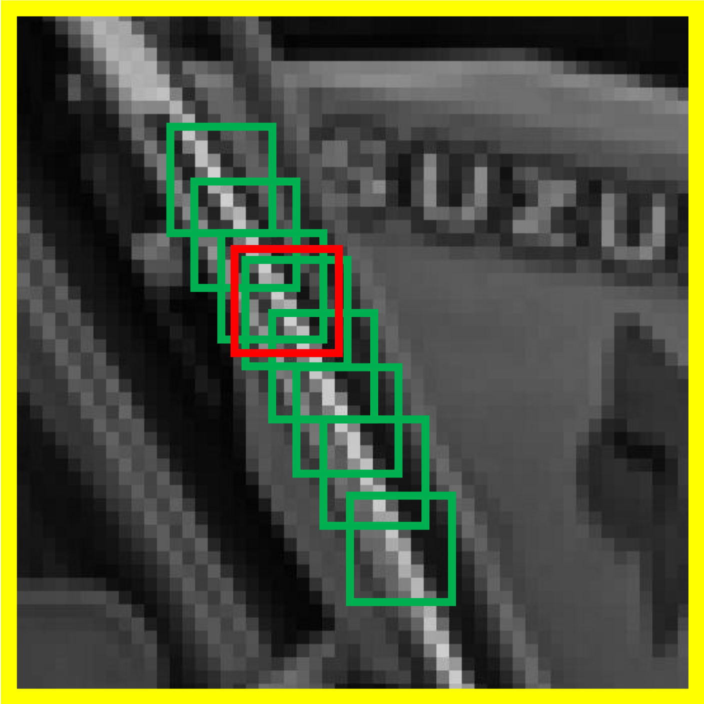

HR image
target patch and
similar patches

[Yu and Orchard '19]

Regularized Weights

$$\omega = \underset{\omega}{\operatorname{argmin}} \|\mathbf{Q}_i \omega - \mathbf{p}_i\|_2^2 + \omega^\top \mathbf{D} \omega$$

Data Fidelity Term

Penalty Term

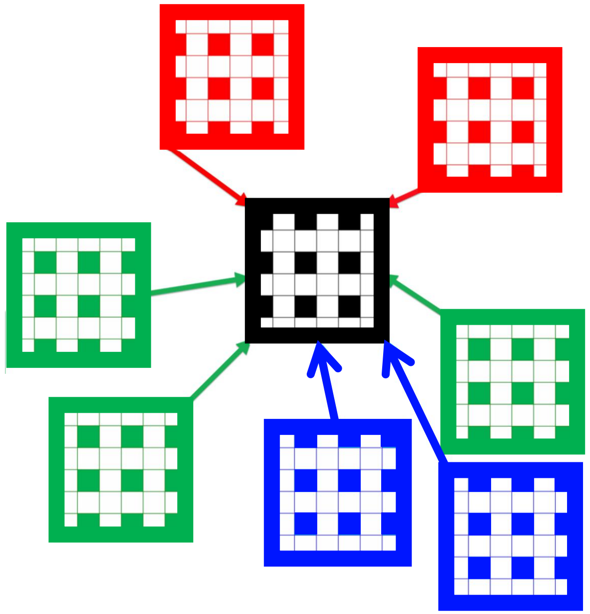

Typical Filter Coefficients

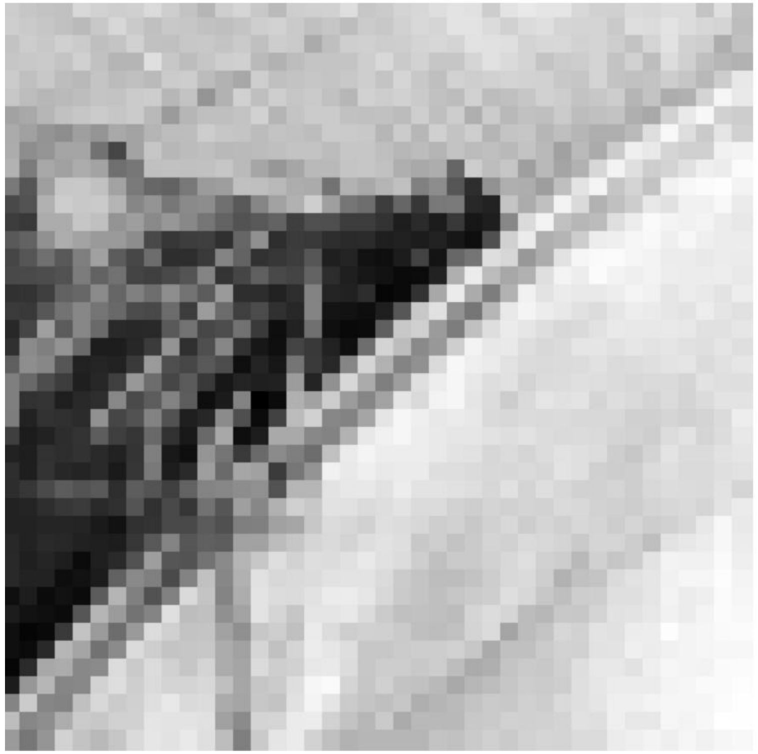

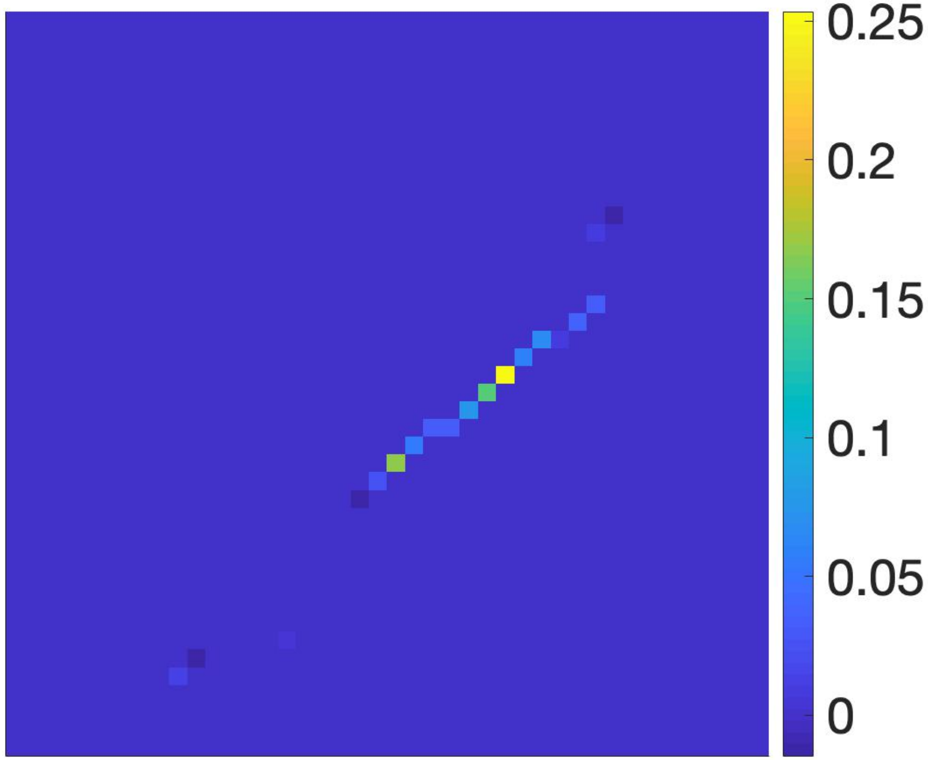

the neighboring measured pixels of the target missing pixel A

the weights of the neighboring measured pixels

Typical Filter Coefficients

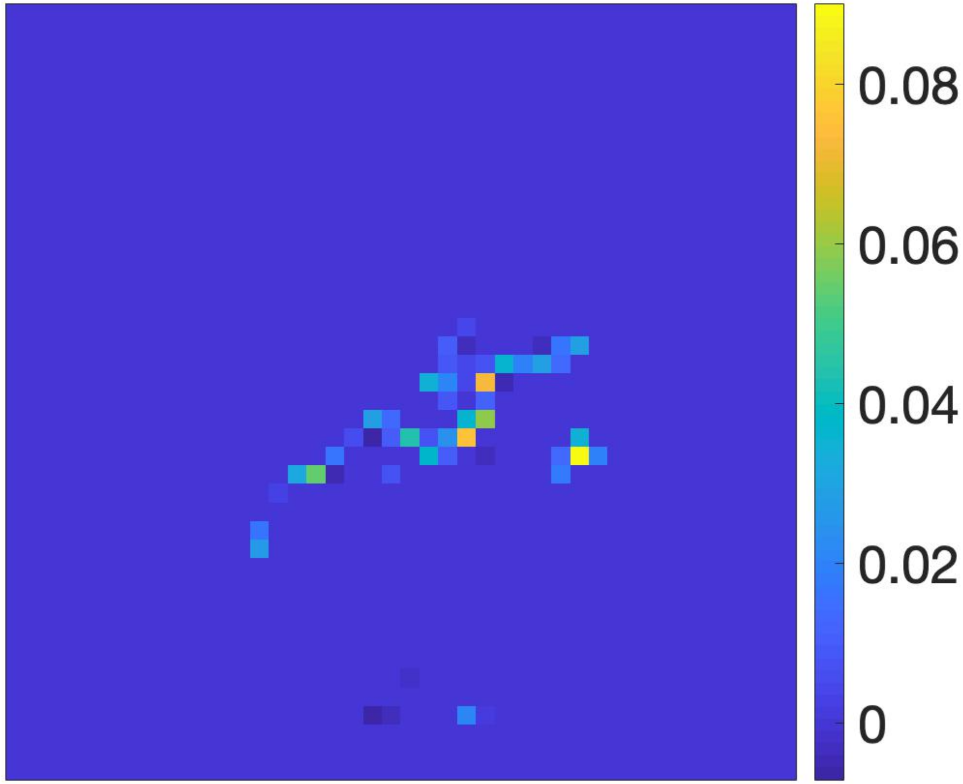

the neighboring measured pixels of the target missing pixel B

the weights of the neighboring measured pixels

Rethink Image Filtering

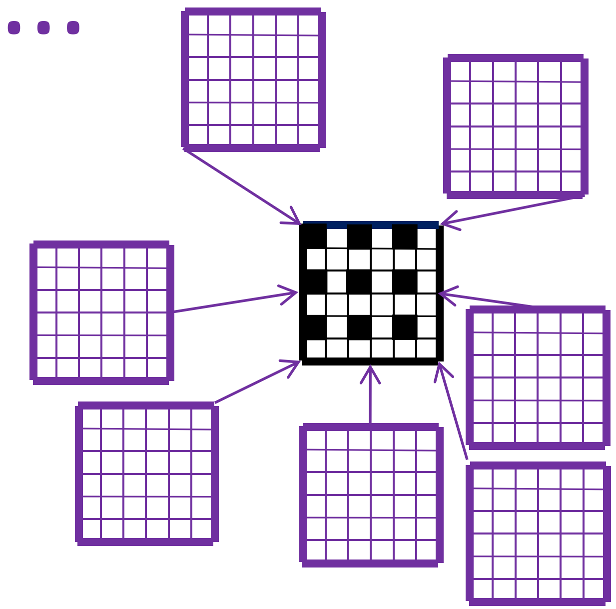

image filtering in final stage

$\mathbf{X}$:matrix of grouped patches

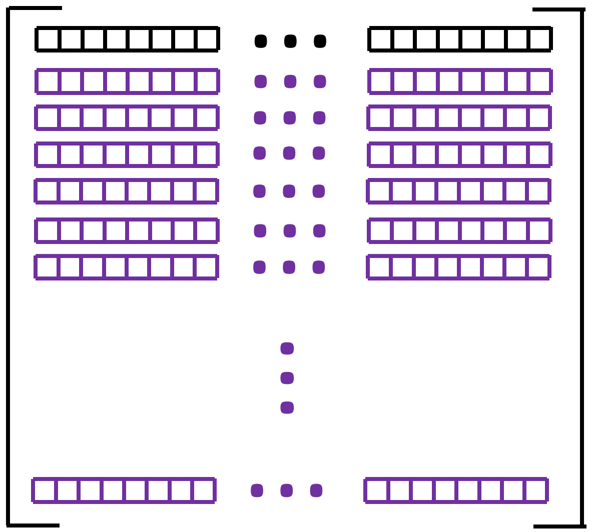

1. filtering = minus-one rank regularization
2. only the first row of $\mathbf{X}$ will be updated

Low-rank Regularization

$$\mathbf{L} = \underset{\mathbf{L}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{X} - \mathbf{L}\|_{\mathrm{F}}^2 + \lambda \cdot \operatorname{rank}(\mathbf{L})$$

We solve a more tractable surrogate:

Weighted Nuclear Norm Minimization (non-descending weights)

$$\mathbf{L} = \underset{\mathbf{L}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{X} - \mathbf{L}\|_{\mathrm{F}}^2 + \sum_i \omega_i |\sigma_i(\mathbf{L})|_1$$ [Gu et al. '17] [Dong et al. '14]

$$\mathbf{L}^\star = \mathbf{U} \mathcal{S}_w(\mathbf{\Sigma}) \mathbf{V}^\top$$ **(global minimizer)**

where: $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top = \mathbf{X}$ and $\mathcal{S}_\omega(\mathbf{\Sigma})_{ii} = \max(\mathbf{\Sigma}_{ii} - \omega_i, 0)$

Algorithm

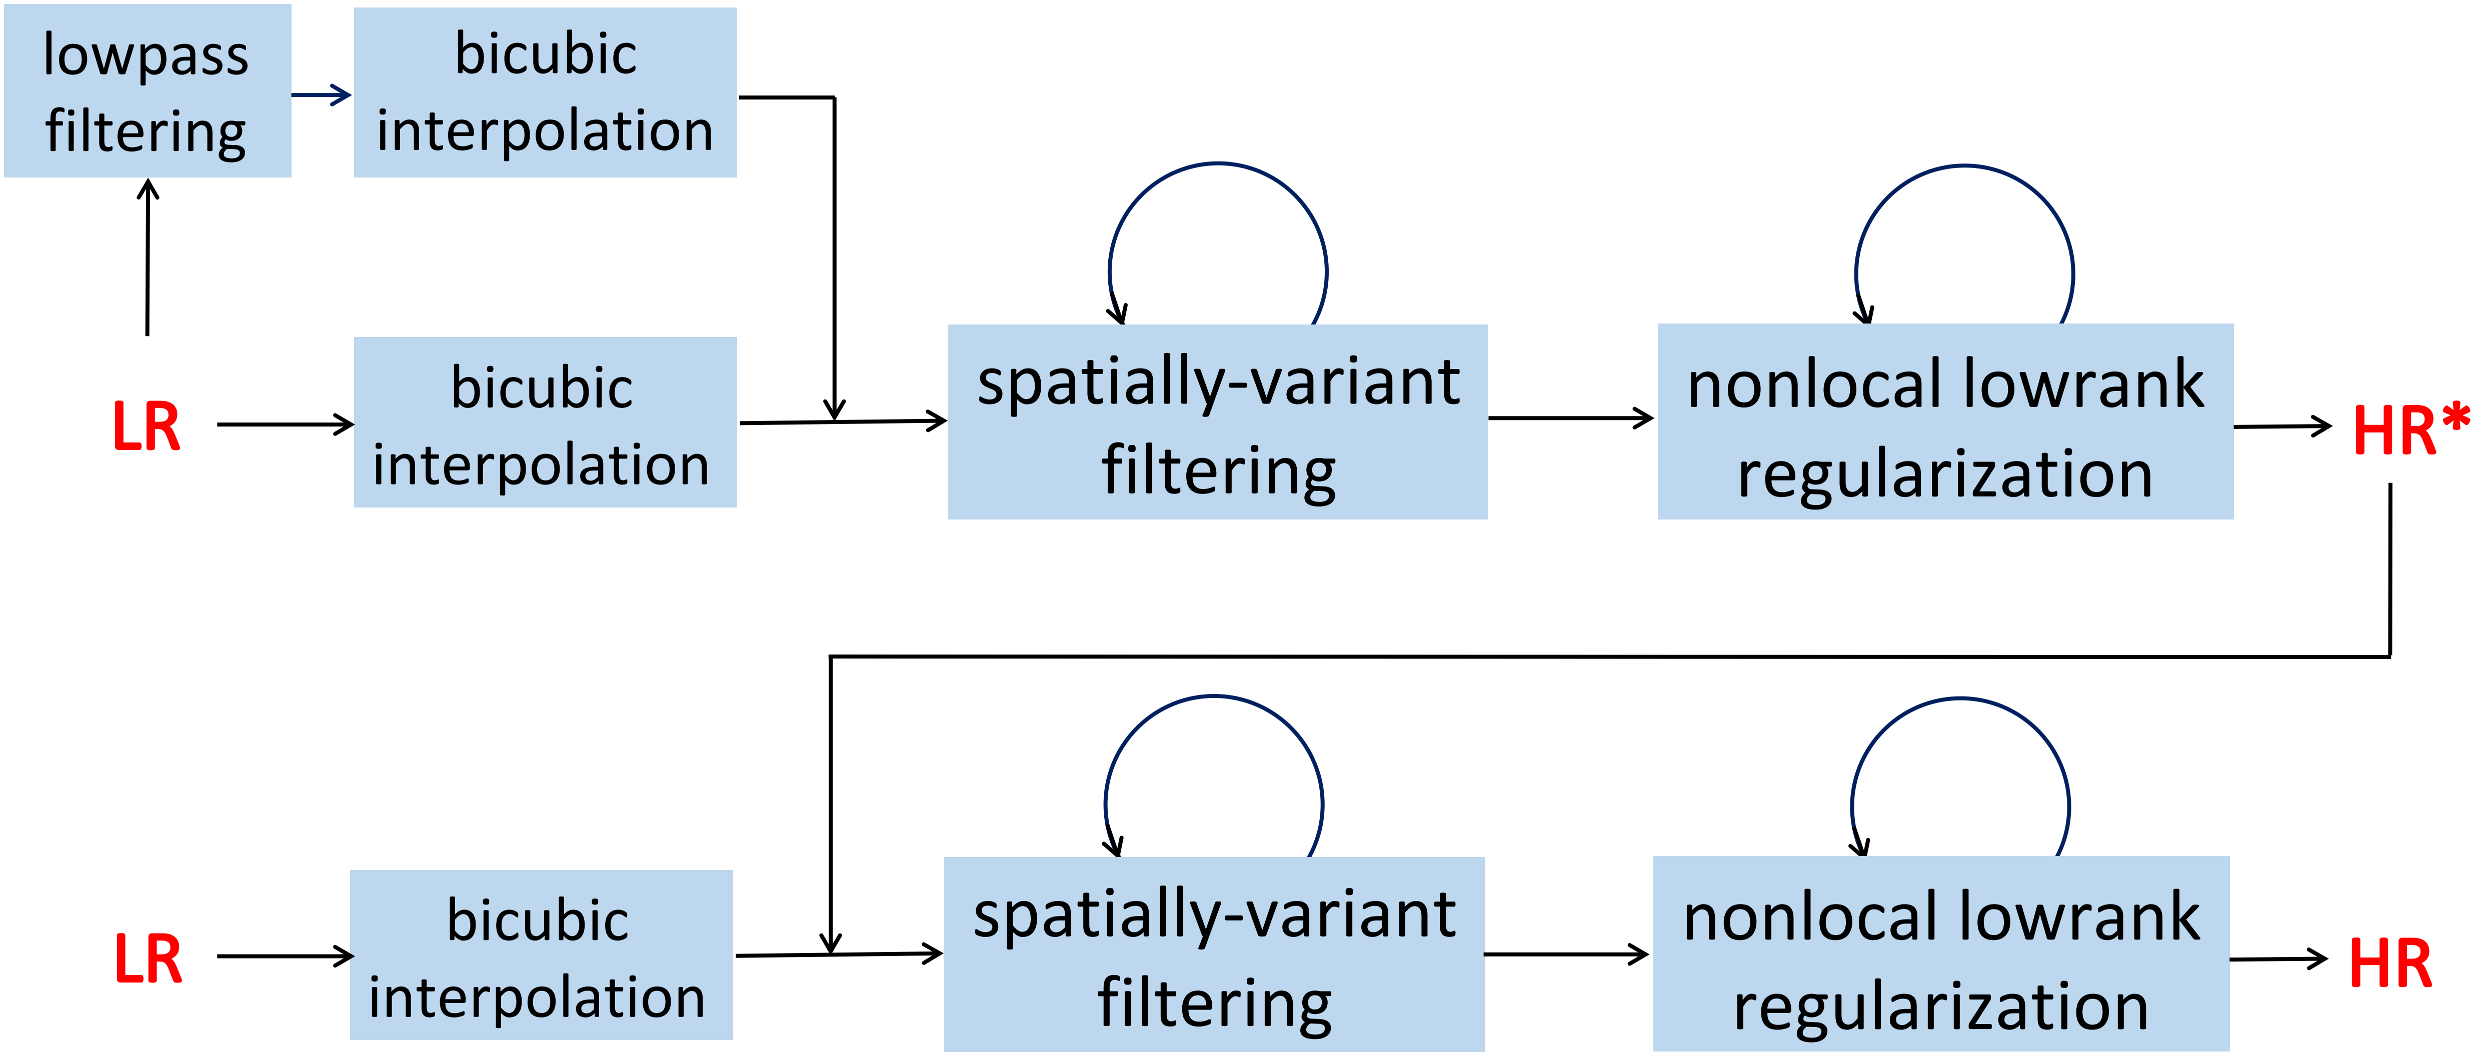
lowpass filtering
bicubic interpolation
LR
bicubic interpolation
spatially-variant filtering
nonlocal lowrank regularization
HR*
LR
bicubic interpolation
spatially-variant filtering
nonlocal lowrank regularization
HR

Testset

Images are from USC-SIPI Database, Berkeley Segmentation Dataset, Kodak and IMAX.
The name of the images in the first row(from left to right): Elk, Birds, Butterfly, Flower.
The name of the images in the first row(from left to right): Leaves, Male, Lena, House.

Quantitative Comparison

Table 1: Comparison of Average PSNR (in decibels) of interpolated images in the task of interpolating an image by a factor of 2.

Image	NARM [Dong et al. '13]	ANSM [Romano et al. '14]	NLPC [Sun et al. '16]	Ours
Elk	31.95	32.51	32.31	**32.80**
Birds	35.03	34.67	35.00	**35.15**
Butterfly	28.23	27.90	27.86	**28.83**
Flower	34.41	34.13	34.22	**34.67**
Leaves	29.38	28.84	29.23	**30.47**
Male	32.41	32.40	32.50	**32.72**
Lena	35.09	34.87	35.08	**35.25**
House	33.49	34.46	33.93	**34.85**
AVERAGE	32.50	32.47	32.52	**33.09**

Ablation Study

Table 2: Comparison of Average PSNR (in decibels) of interpolated images in the task of interpolating an image by a factor of 2.

Image	Single Pass Filtering without NLR	Two Pass Filtering without NLR	Entire Process
Elk	32.72	32.77	32.80
Birds	35.12	35.10	35.15
Butterfly	28.49	28.66	28.83
Flower	34.61	34.61	34.67
Leaves	29.95	30.19	30.47
Male	32.70	32.70	32.72
Lena	35.24	35.22	35.25
House	34.56	34.70	34.85
AVERAGE	32.93	32.99	33.09

Visual Comparison (X2)

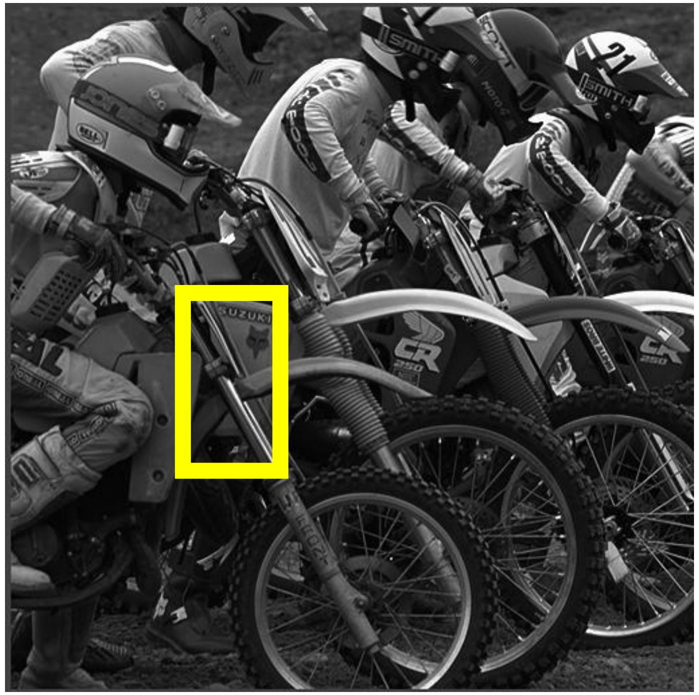

NARM

ANSM

NLPC

Ground-Truth

Proposed

NARM: [Dong et al. '13]
ANSM: [Romano et al. '14]
NLPC: [Sun et al. '16]

Visual Comparison (X2)

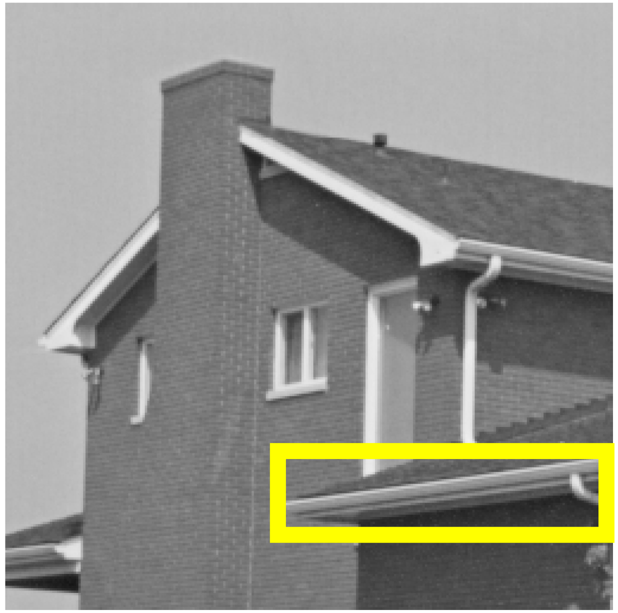

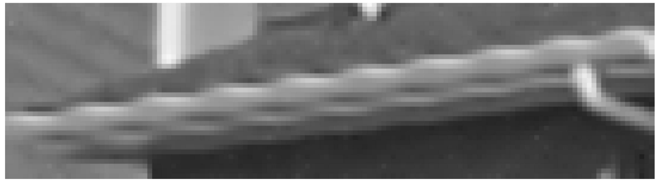

NARM

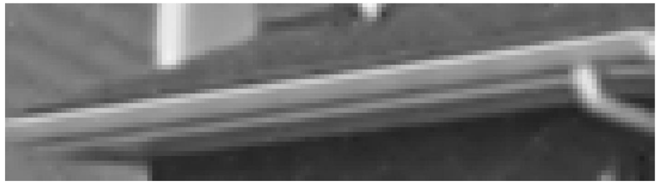

ANSM

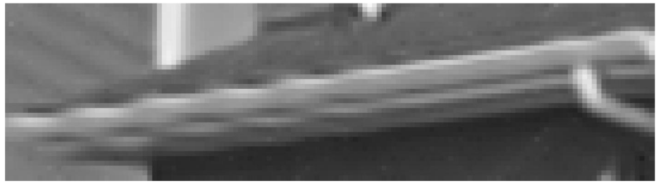

NLPC

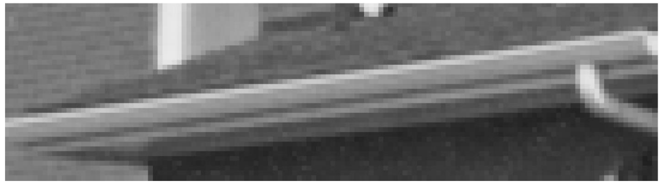

Ground-Truth

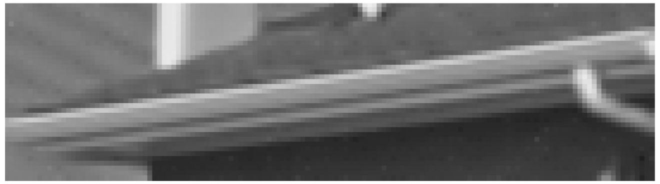

Proposed

NARM: [Dong et al. '13]
ANSM: [Romano et al. '14]
NLPC: [Sun et al. '16]

Visual Comparison (X2)

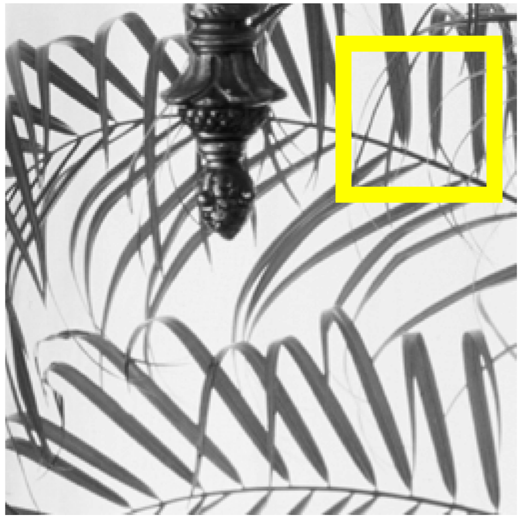

NARM

ANSM

NLPC

Ground-Truth

Proposed

NARM: [Dong et al. '13]
ANSM: [Romano et al. '14]
NLPC: [Sun et al. '16]

Reference

[1] Antoni Buades,Bartomeu Coll, and Jean-Michel Morel,“A non-local algorithm for image denoising,” CVPR 2005.

[2] M Aharon, M Elad, and A Bruckstein, “K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation,” TSP 2006.

[3] Weisheng Dong, Lei Zhang, Rastislav Lukac, and Guangming Shi, “Sparse representation based image interpolation with nonlocal autoregressive modeling,” TIP 2013.

[4] Weisheng Dong, Guangming Shi, Xin Li, Yi Ma, and Feng Huang, “Compressive sensing via nonlocal low-rank regularization,” TIP 2014.

[5] Yaniv Romano, M Protter, and M Elad, “Single image interpolation via adaptive non-local sparsity-based modeling,” TIP 2014.

[6] Dong Sun, Qingwei Gao, and Yixiang Lu, “Image interpolation via collaging its non-local patches,” Digital Signal Processing 2016.

[7] Shuhang Gu, Qi Xie, Deyu Meng, Wangmeng Zuo, Xianchu Feng, and Lei Zhang, “Weighted nuclear norm minimization and its applications to low level vision,” IJCV 2017.

[8] Lantao Yu, and Michael T. Orchard, “Single image interpolation exploiting semi-local similarity,” ICASSP 2019.

Conclusion

- ☐ Combine spatially-variant filtering and lowrank approximation to exploit non-local similarity
- ☐ State-of-the-art PSNR
- ☐ Simple, Parallerizable Algorithm