Edge-Preserving Laplacian Pyramid

Stella X. Yu
Computer Science
Boston College

Acknowledgements:
Clare Boothe Luce Professorship
NSF CAREER IIS-0644204

Edge-Preserving Pyramid for Image Synopsis

1. remove spatial redundancy
2. retain perceptual saliency
3. refine over scale

image

smaller

Multiscale based on Signal or Perceptual Analysis

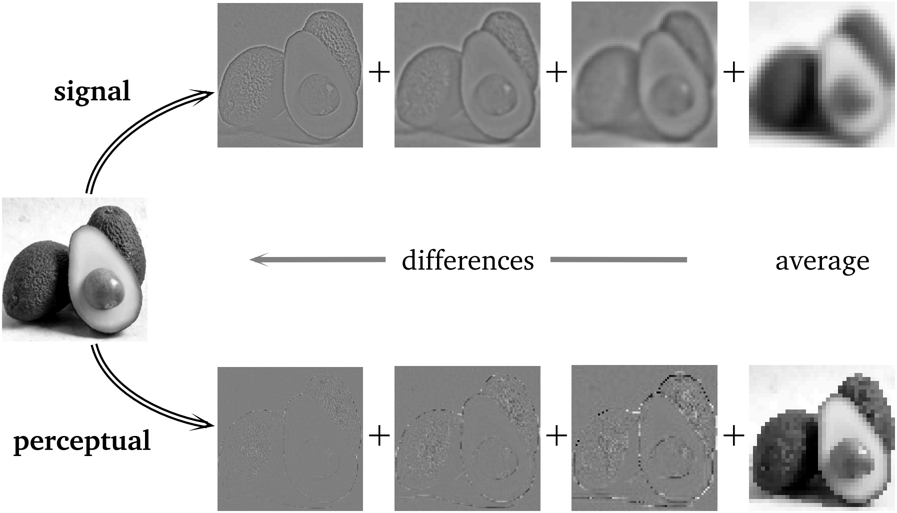

Local Average to Remove Spatial Redundancy

$$I(p) \approx \bar{I}(p) = \sum_{q=p'\text{s local neighbour}} W(p,q) \cdot I(q)$$

- $W(p,q)$ describes how q contributes to predicting p's intensity
- $\bar{I}$ becomes smoother than I and can be downsampled
- Pyramid: recursive application of averaging + downsampling

In the signal-based multiscale analysis:

- W is pre-chosen
- W is fixed over the entire image
- No consideration for either $\bar{I}$ or $I - \bar{I}$

Laplacian Pyramid Construction and Collapsing

Analysis: Given analysis weights $W_\triangledown$ and synthesis weights $W_\vartriangle$,

$$\begin{array}{lll} \text{average:} & A_{s+1} = \downarrow(A_s, W_\triangledown), & s = 1 \to n, \quad A_1 = I \\ \text{difference:} & D_s = A_s - \uparrow(A_{s+1}, W_\vartriangle), & s = n \to 1, \quad D_{n+1} = A_{n+1} \end{array}$$

Synthesis: Given difference pyramid D,

$$\begin{array}{ll} \text{average:} & A_s = D_s + \uparrow(A_{s+1}, W_\vartriangle), \quad s = n \to 1, \quad A_{n+1} = D_{n+1} \\ \text{reconstruction:} & I = A_1 \end{array}$$

Analysis and Synthesis Weights Are Independent

Before: $W_\triangledown$ and $W_\vartriangle$ are identical and spatially invariant.

$$\text{proximity: } W_\triangledown(p,q) = W_\vartriangle(p,q) = G(\|\overrightarrow{p} - \overrightarrow{q}\|;\sigma)$$

$$G(d;\sigma) = \exp\left(-\frac{d^2}{2\sigma^2}\right)$$

Fact: $W_\triangledown$ and $W_\vartriangle$ can be **independently** defined and in fact **arbitrary** without jeopardizing a perfect reconstruction.

After: $W_\triangledown \neq W_\vartriangle$, both vary according to edges at each pixel.

$$\begin{aligned} W_\triangledown(p,q;I) &= G(\|\overrightarrow{p} - \overrightarrow{q}\|;\sigma_\triangledown) \times K_g(p,q;I) \\ W_\vartriangle(p',q;I) &= G(\|\overrightarrow{p}' - \overrightarrow{q}\|;\sigma_\vartriangle) \times K_g(p',q;I) \end{aligned}$$

Edge Geometry Kernel K_g for Downsampling

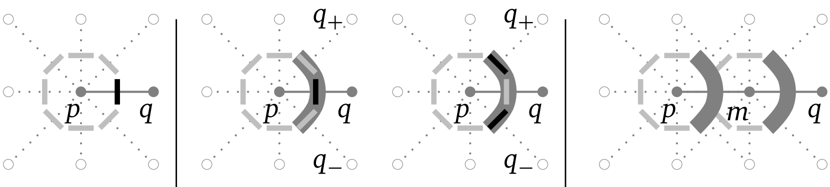

$C(p,q) \approx 0$ | $K_g(p,q) \approx 0$ via q_+ and q_- | $C(p,q)$ via m

$$L(p,q) = \begin{cases} \min(E(p), E(q)), & P(p) \neq P(q) \\ 0, & P(p) = P(q) \end{cases}$$

$$C(p,q) = \begin{cases} G(L(p,q);\, \sigma_g), & q \in N(p,1) \\ \min(K_g(p,m), K_g(m,q)), \quad \vec{m} = \frac{\vec{p}+\vec{q}}{2}, & q \in N(p,2) \end{cases}$$

$$K_g(p,q) = \min(C(p,q), \max_{o \in \{q_+, q_-\}} C(p,o)), |\measuredangle q_\pm p q| = 45^\circ, q, q_\pm \in N(p,r)$$

Edge Geometry Kernel K_g for Upsampling

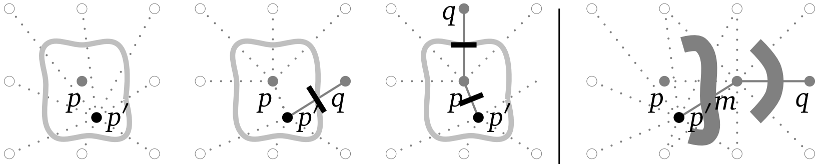

$C(p',q), q \in N(p,1)$ $\qquad$ $C(p',q), q \in N(p,2)$

$$C(p',q) = \begin{cases} G(L(p',q);\sigma_g), & \|\vec{p}\,' - \vec{q}\| \leq \|\vec{p} - \vec{q}\|, q \in N(p,1) \\ G(\max(L(p',p), L(p,q)); \sigma_g), & \|\vec{p}\,' - \vec{q}\| > \|\vec{p} - \vec{q}\|, q \in N(p,1) \\ \min(K_g(p',m), K_g(m,q)), & \vec{m} = \frac{\vec{p} + \vec{q}}{2}, q \in N(p,2) \end{cases}$$

$$K_g(p',q) = \min(C(p',q), \max_{o \in \{q_+, q_-\}} C(p',o)), |\measuredangle q_{\pm} p q| = 45^\circ, q, q_{\pm} \in N(p,r)$$

Comparison of Interpolation Methods

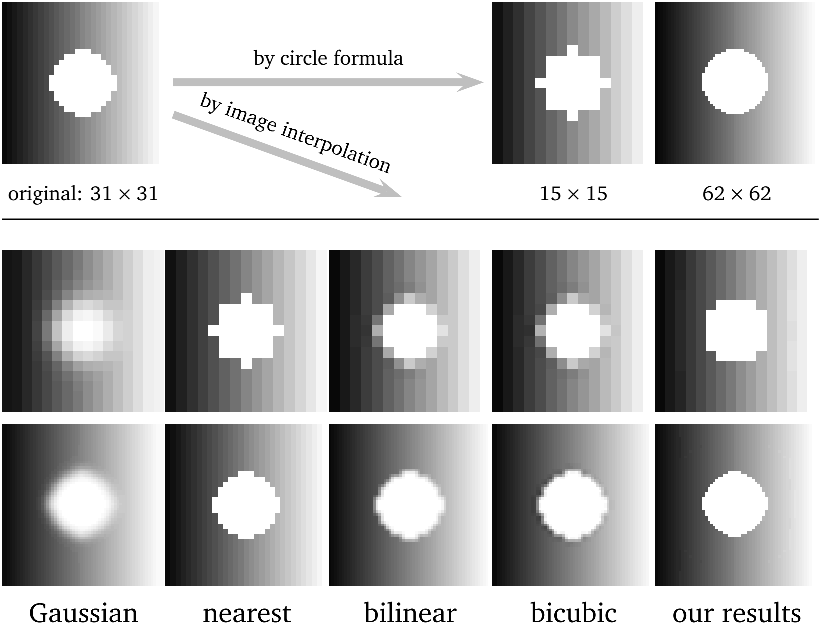

Comparison of the Average as an Image Synopsis

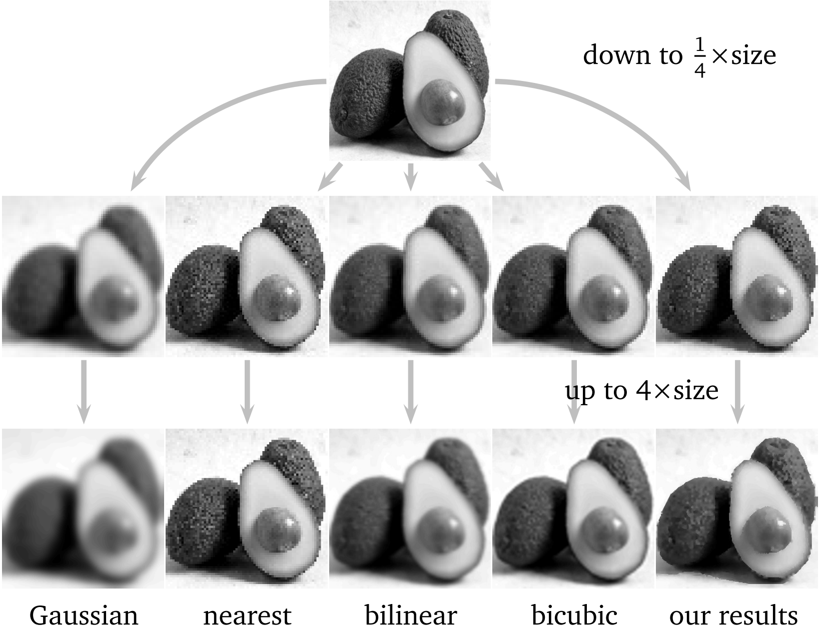

Comparison of the Difference as an Image Code

standard test images

Reduced Entropy in the Difference Images

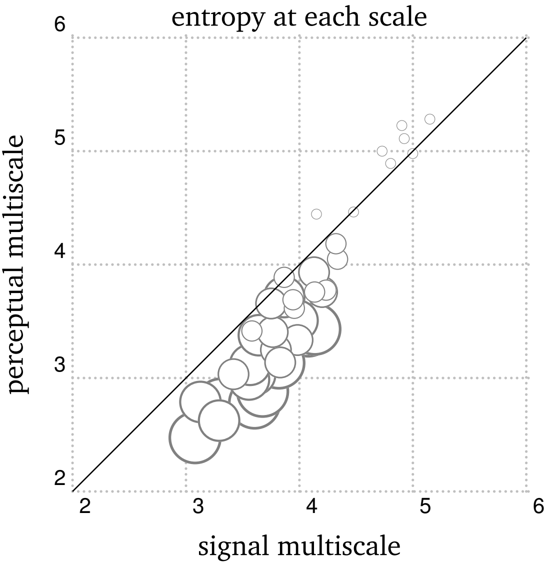

4-Time Additional Savings in Lossless Compression

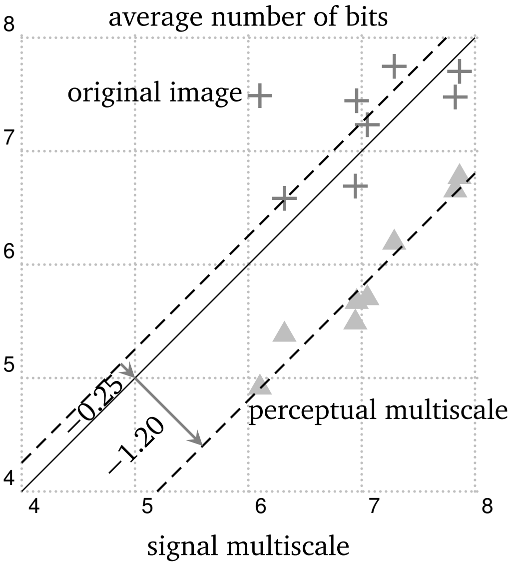

More Savings than Any Interpolation Methods

perceptual over	Gaussian	nearest	bilinear	bicubic
bits per pixel	-1.20	-0.16	-0.37	-0.38
confidence	± 0.14	± 0.14	± 0.05	± 0.05
p-value	1.5×10^{-7}	3.3×10^{-2}	4.1×10^{-7}	3.0×10^{-7}

As an image code, nearest neighbour is most efficient.
It is better than the widely known Laplacian Pyramid.

Faithful Image Synopsis and Effective Image Code

- Signal-based multiscale methods always face a trade-off:

  as an image synopsis, bicubic interpolation is most faithful;
  as an image code, nearest neighbour is most efficient.

- Edge-preserving pyramid outperforms on either account:

  The averages retain boundaries and shading at lower spatial and tonal resolutions;
  The differences refine edge locations and intensity details with a remarkably sparse code.

Distinction with Other Edge-Preserving Methods

- anisotropic diffusion: gradients → curvilinearity
- bilateral filtering: intensity similarity → boundary separation
- nonlocal mean filtering: staircasing effects → none
- wavelets: expand basis, e.g. ridgelets → locally adaptive basis