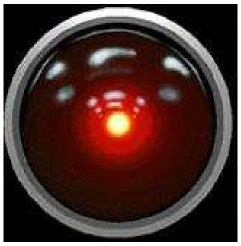

HCI/ComS 575X:
Computational Perception

Instructor: Alexander Stoytchev
http://www.cs.iastate.edu/~alex/classes/2007_Spring_575X/

Administrative Stuff

Deadlines Clarification

- On-campus Students
  - 11:59pm Central Time on the day it is due
- Off-campus students
  - 2 days after it is due, but …
  - If it is due on Monday – 4:00pm on Wednesday
  - If it is due on Wednesday - 11:59pm on Friday

Image Pyramids

January 29, 2007

HCI/ComS 575X: Computational Perception
Iowa State University, SPRING 2007

"The Laplacian Pyramid as a Compact Image Code"

By PETER J. BURT and **EDWARD H. ADELSON**
IEEE TRANSACTIONS ON COMMUNICATIONS, VOL. COM-3l, NO. 4, APRIL 1983

(Available on the Class Web page)
[Readings Section]

The Authors


Image Pyramids?

What is an Image Pyramid?

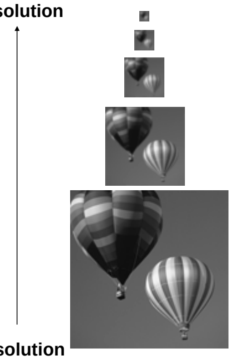

[www.wisdom.weizmann.ac.il/~mblank/CVfall04/handouts/lec4.ppt]

Fast Pattern Matching

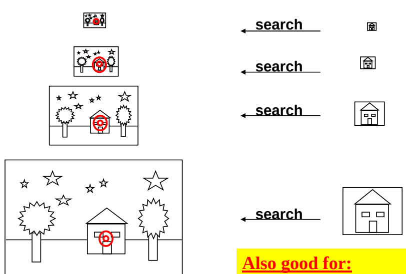

Also good for:
- motion analysis
- image compression
- other applications

[www.wisdom.weizmann.ac.il/~mblank/CVfall04/handouts/lec4.ppt]

The Theory

Dropping Pixels v.s. Smoothing and then dropping Pixels

Why does this look so bad?

1/2 1/4 (2x zoom) 1/8 (4x zoom)

[From Alexei Efros' lecture on Sampling and Pyramids]

Subsampling with Gaussian pre-filtering

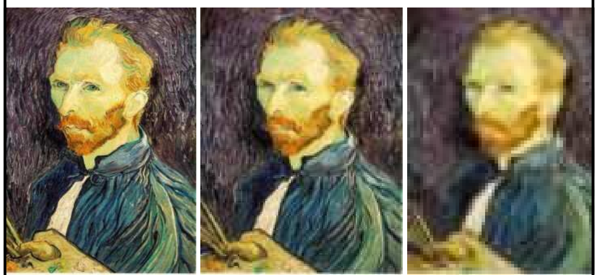

Gaussian 1/2 G 1/4 G 1/8

- Solution: filter the image, *then* subsample
  - Filter size should double for each ½ size reduction.

[From Alexei Efros' lecture on Sampling and Pyramids]

Sampling

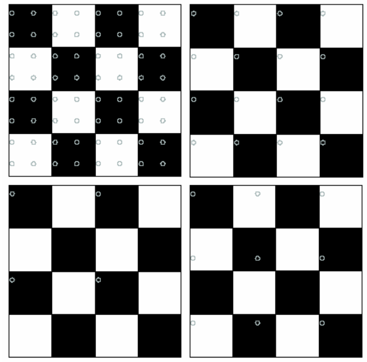

Good sampling:
•Sample often or,
•Sample wisely

Bad sampling:
•see aliasing in action!

[From Alexei Efros' lecture on Sampling and Pyramids]

Gaussian pre-filtering

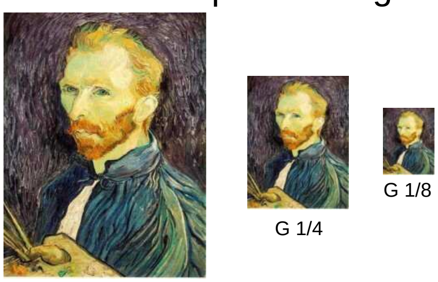

G 1/8

G 1/4

Gaussian 1/2

- Solution: filter the image, *then* subsample
  - Filter size should double for each ½ size reduction.

[From Alexei Efros' lecture on Sampling and Pyramids]

The Gaussian Pyramid

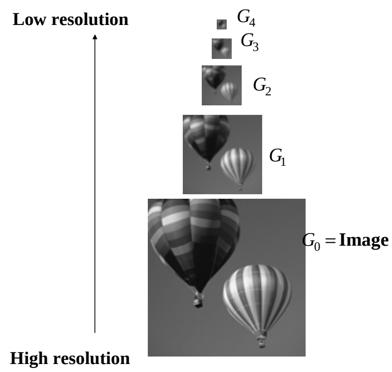

[www.wisdom.weizmann.ac.il/~mblank/CVfall04/handouts/lec4.ppt]

The Gaussian Pyramid

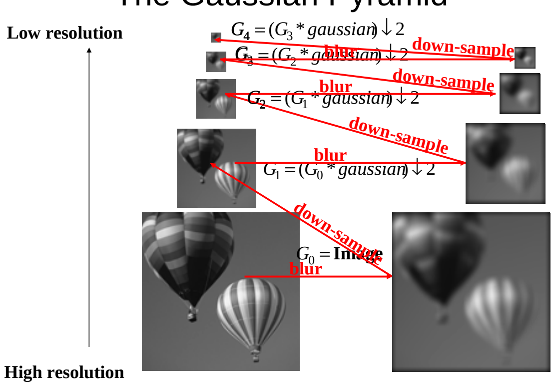

[www.wisdom.weizmann.ac.il/~mblank/CVfall04/handouts/lec4.ppt]

[www.wisdom.weizmann.ac.il/~mblank/CVfall04/handouts/lec4.ppt]

The Laplacian Pyramid

$L_i = G_i - \text{expand}(G_{i+1})$

$G_i = L_i + \text{expand}(G_{i+1})$

Gaussian Pyramid Laplacian Pyramid

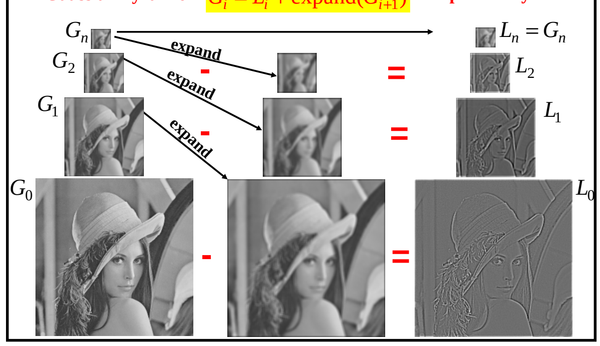

Laplacian ~ Difference of Gaussians

"The value at each node in the Laplacian Pyramid is the difference between the convolutions of two equivalent weighting functions h_l, h_{l+1} with the original image. Again, this is similar to convolving an appropriately scaled Laplacian weighting function with the image. The node value could have been obtained directly by applying this operator, although at considerably greater computational cost."

[Burt and Adelson 1983]

Laplacian ~ Difference of Gaussians

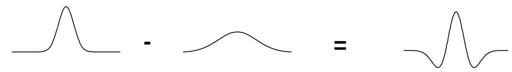

DOG = Difference Of Gaussians

More details on Gaussian and Laplacian pyramids can be found in the paper by Burt and Adelson

[www.wisdom.weizmann.ac.il/~mblank/CVfall04/handouts/lec4.ppt]

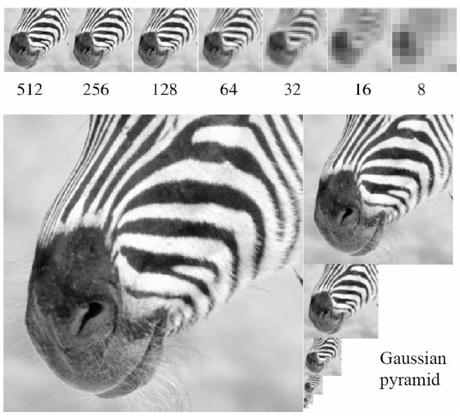

[From Alexei Efros' lecture on Sampling and Pyramids]

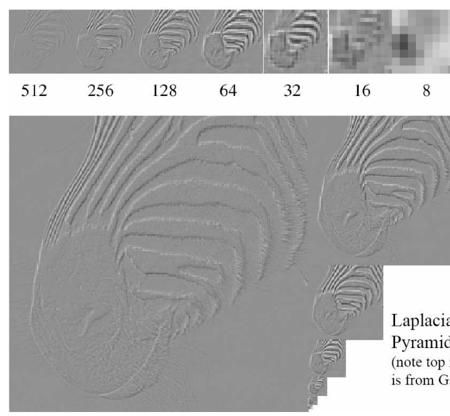

[From Alexei Efros' lecture on Sampling and Pyramids]

Constructing the Gaussian Pyramid

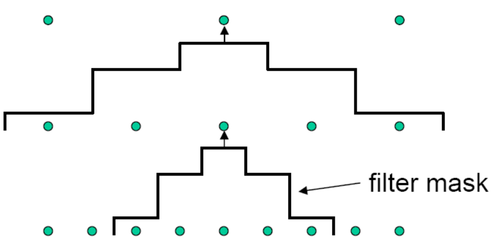

[From Alexei Efros' lecture on Sampling and Pyramids]

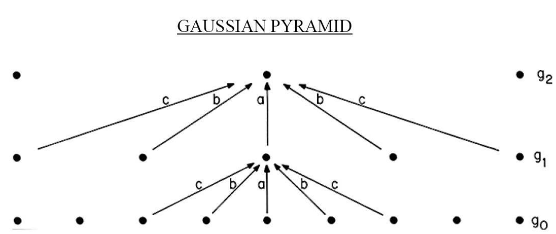

g_0 = IMAGE

g_L = REDUCE [g_{L-1}]

The Generating Kernel

- One Example:

[0.05, 0.25, 0.4, 0.25, 0.05]

Properties

- Separable: $w(m, n) = \hat{w}(m)\ \hat{w}(n).$
- Normalized: $\sum_{m=-2}^{2} \hat{w}(m) = 1$
- Symmetric: $\hat{w}(i) = \hat{w}(-i)$ for $i = 0, 1, 2.$

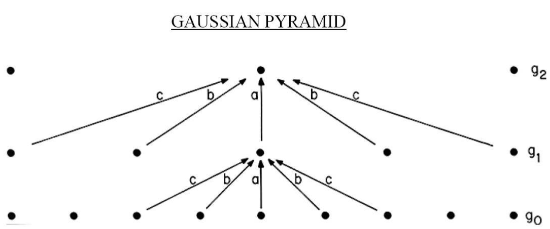

g_0 = IMAGE

g_L = REDUCE $[g_{L-1}]$

Symmetric Property

w=[c b a b c]

$\hat{w}(0) = a$

$\hat{w}(-1) = \hat{w}(1) = b$

$\hat{w}(-2) = \hat{w}(2) = c$

Symmetric + Normalized

- Symmetric: w=[c b a b c]
- Normalized: $\sum_{m=-2}^{2} \hat{w}(m) = 1$
- This must hold: a + 2b + 2c = 1

Equal Distribution:

all nodes at a given level must contribute the same total weight to nodes at the next higher level

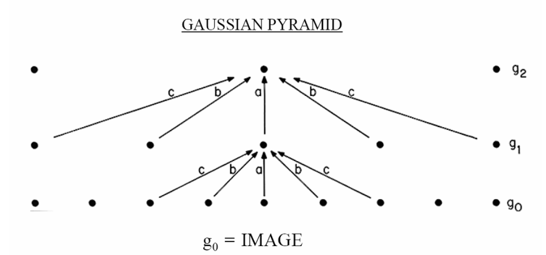

What is the contribution of this node?

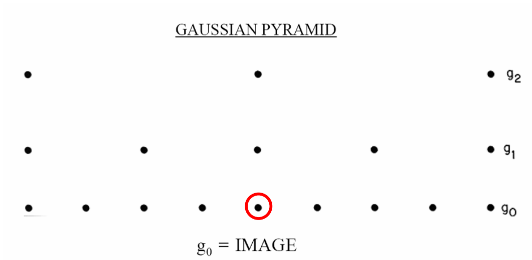

What is the contribution of this node?

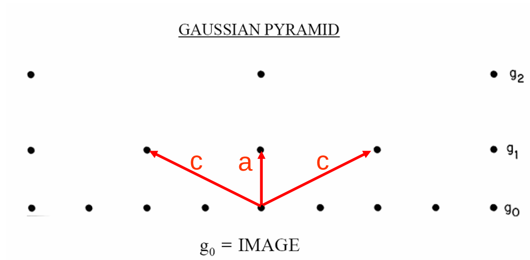

How About this one?

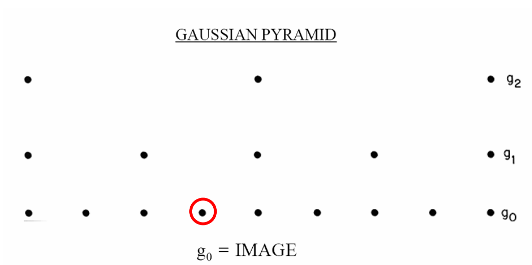

How About this one?

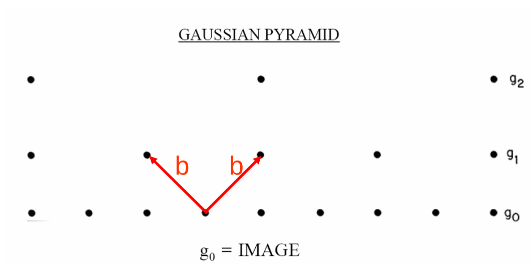

Equal Distribution

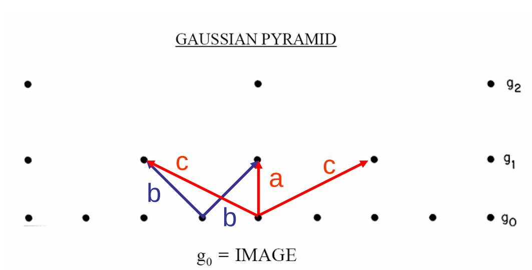

This must hold: $a + 2c = 2b$

2 Constraints, 3 Variables

$a + 2b + 2c = 1$

$a + 2c = 2b$

We can derive that $4b = 1$

Therefore, $b = 1/4$

The constraints are satisfied when:

$$\hat{w}(0) = a$$

$$\hat{w}(-1) = \hat{w}(1) = 1/4$$

$$\hat{w}(-2) = \hat{w}(2) = 1/4 - a/2.$$

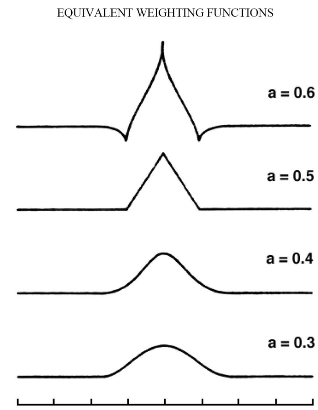

Optimal value for a for best compression results

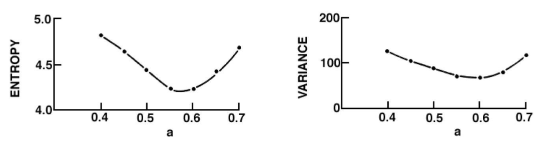

EXPAND and REDUCE

We now define a function EXPAND as the reverse of REDUCE. Its effect is to expand an $(M+1)$-by-$(N+1)$ array into a $(2M+1)$-by-$(2N+1)$ array by interpolating new node values between the given values. Thus, EXPAND applied to array g_l of the Gaussian pyramid would yield an array $g_{l,1}$ which is the same size as g_{l-1}.

REDUCE

$$g_k = \text{REDUCE}(g_{k-1}) \qquad (1)$$

which means, for levels $0 < l < N$ and nodes i, j, $0 \leq i < C_l$, $0 \leq j < R_l$,

$$g_l(i, j) = \sum_{m=-2}^{2} \sum_{n=-2}^{2} w(m, n) g_{l-1}(2i + m, 2j + n).$$

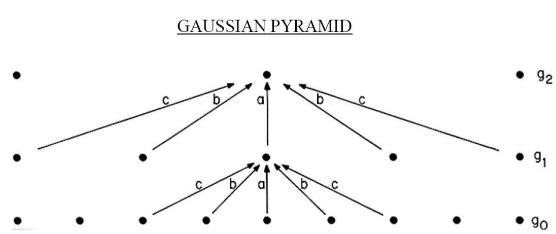

g_0 = IMAGE

g_L = REDUCE $[g_{L-1}]$

EXPAND

By EXPAND we mean, for levels $0 < l \le N$ and $0 \le n$ and nodes i, j, $0 \le i < C_{l-n}$, $0 \le j < R_{l-n}$,

$$g_{l,n}(ij) = 4 \sum_{m=-2}^{2} \sum_{n=-2}^{2} w(m,n) \cdot g_{l,n-1}\left(\frac{i-m}{2}, \frac{j-n}{2}\right). \qquad (2)$$

Only terms for which $(i-m)/2$ and $(j-n)/2$ are integers are included in this sum.

The Gaussian Pyramid

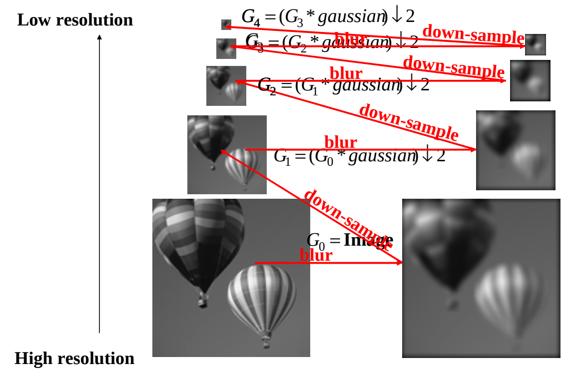

[www.wisdom.weizmann.ac.il/~mblank/CVfall04/handouts/lec4.ppt]

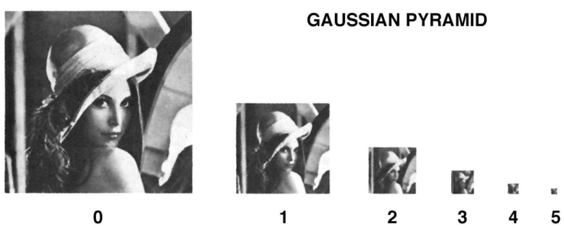

[www.wisdom.weizmann.ac.il/~mblank/CVfall04/handouts/lec4.ppt]

The Laplacian Pyramid

$L_i = G_i - \text{expand}(G_{i+1})$

$G_i = L_i + \text{expand}(G_{i+1})$

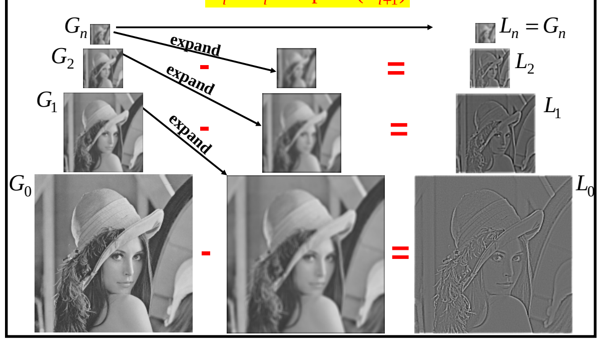

Histogram and Entropy of the Original Image

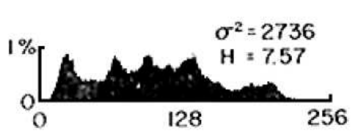

$$H = -\sum_{i=0}^{255} f(i) \log_2 f(i)$$

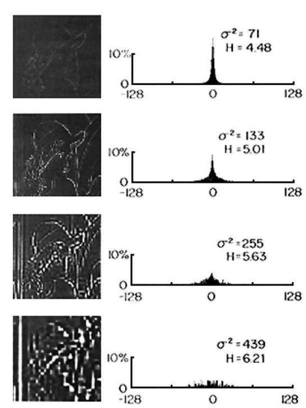

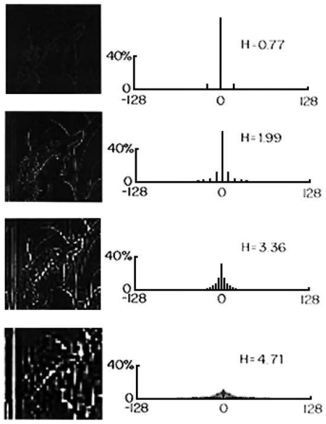

MATLAB DEMO

Sample Uses of Pyramids

Blending Regions

Horror Photo

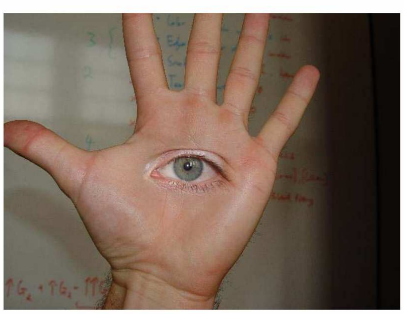

© prof. dmartin

Image Morphing

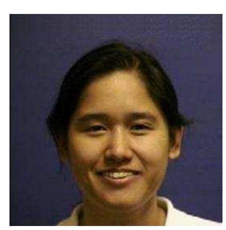

[Female Image]

[Male Image]

[Done by me for a homework long time ago]

Image Morphing

[No Blending]

[With Blending]

[Done by me for a homework long time ago]

Pyramid Blending

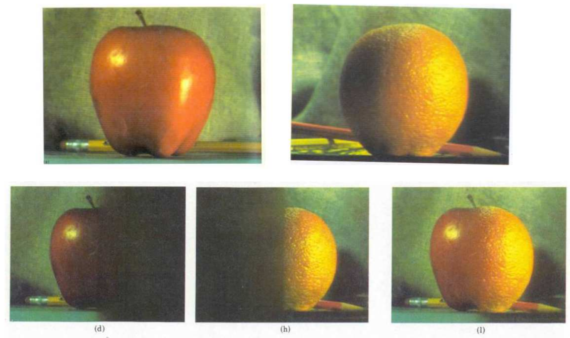

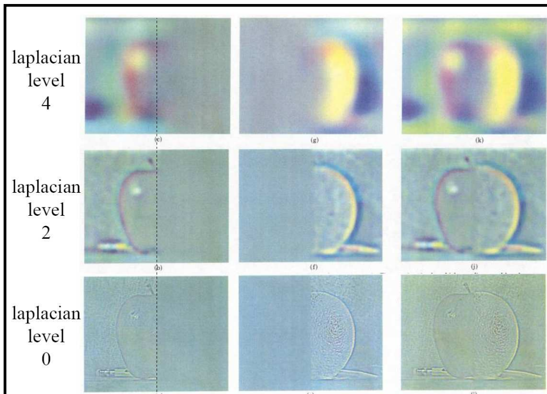

Laplacian Pyramid: Blending

General Approach:

1. Build Laplacian pyramids *LA* and *LB* from images *A* and *B*
2. Build a Gaussian pyramid *GR* from selected region *R*
3. Form a combined pyramid *LS* from *LA* and *LB* using nodes of *GR* as weights:
   - *LS(i,j) = GR(I,j,)*LA(I,j) + (1-GR(I,j))*LB(I,j)*
4. Collapse the *LS* pyramid to get the final blended image

Optical flow result

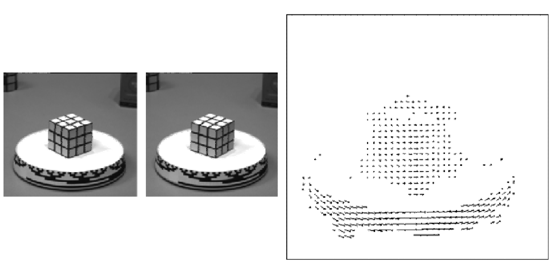

[www.cfar.umd.edu/~fer/cmsc426/lectures/imagemotion.ppt]

[www.cfar.umd.edu/~fer/cmsc426/lectures/imagemotion.ppt]

Coarse-to-fine optical flow estimation

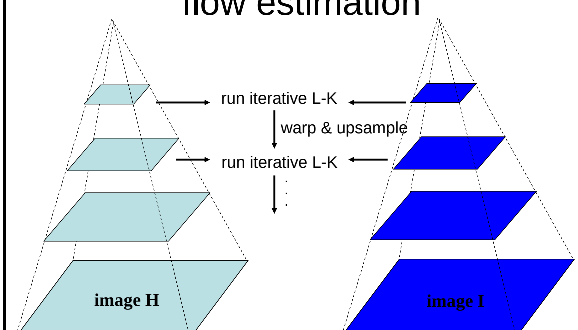

Some Fun Facts About the Paper

The paper was originally rejected

One of the reviews said something like:

– It is a catchy idea but it will not catch up.

[…that's what one of my professors told me]

Who was Lenna?

Answer: Playboy's Miss November 1972

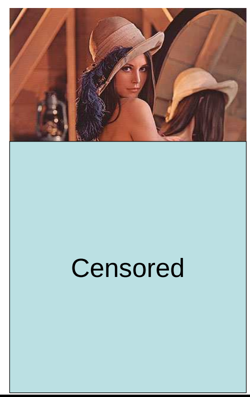

Yes, it's true!

Lenna attended the 50th Anniversary IS&T conference in Boston held in May 1997.

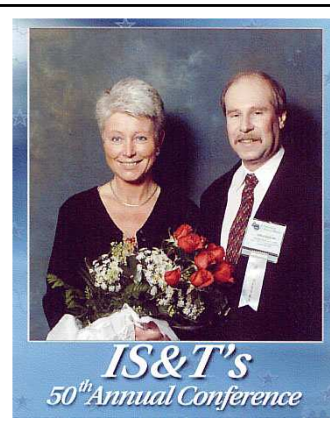

Lenna at the conference

Playmate Meets Geeks Who Made Her a Net Star

by Janelle Brown (Wired Magazine)
6:11pm 20.May.97.PDT

Having graced the desktops of millions of engineers, researchers, and digital imaging specialists for 25 years, *Playboy*'s Miss November 1972 - dubbed the "First Lady of the Internet" - is coming to meet her fans.

Lena Sjooblom became Net royalty when her centerfold was scanned in by programmers at the University of Southern California to use as a test image for digital compression and transmission over Arpanet, the precursor to the Internet. Years later, the "Lena" image (a closeup of her face and bare shoulder) is still the industry standard for tests. This week, Sjooblom is making her first public appearance at the 50th Annual Conference of the Society for Imaging Science in Technology, as part of an overview of the history of digital imaging.

"They must be so tired of me ... looking at the same picture for all these years!" comments Sjooblom, who was, until last November, unaware of her fame, and who has still never seen the Net. This week she is busy signing autographs, posing for pictures, and giving a presentation about herself at the conference.

Playboy helped track down the Swedish native in Stockholm, where she helps handicapped people work on (non-networked) computers. Although *Playboy* is notorious for cracking down on illegal uses of its images, it has decided to overlook the widespread distribution of this particular centerfold.

Says Eileen Kent, VP of new media at *Playboy:* "We decided we should exploit this, because it is a phenomenon."

Additional Reading

http://www.lenna.org/

So, What is the Moral of this Story?

So, What is the Moral of this Story?

- Believe in yourself and don't give up!
- Resubmit the paper!

So, What is the Moral of this Story?

- Believe in yourself and don't give up!
- Resubmit the paper!

- If that still does not work then use the image of a playboy playmate in your paper ☺

THE END