Matting & Compositing

Many slides from Freeman&Durand's Computational Photography course at MIT.
Some are from A.Efros at CMU.
Some from Z.Yin from PSU!
I even made a bunch of new ones…

Motivation: compositing

Combining multiple images. Typically, paste a foreground object onto a new background

- **Movie special effect**
- **Combining graphics & film**
- **Photo retouching**
  - Change background
  - Fake depth of field
  - Page layout: extract objects, magazine covers

Motivation

Slide from Alyosha Efros

Motivation

From Cinefex

Plate 94 *A composite image created for the film* Titanic.

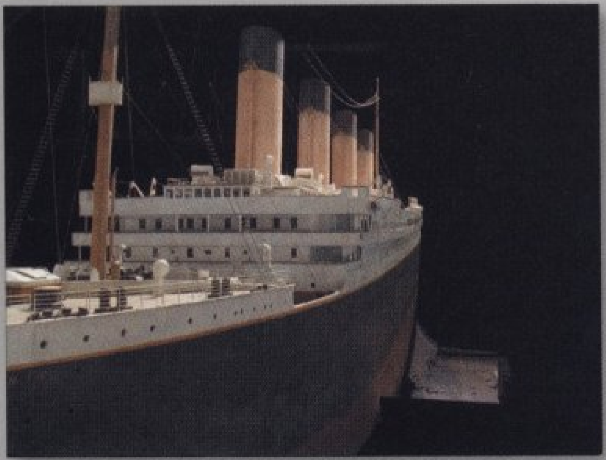

Plate 95 *An element that features a miniature of the ship.*

Plate 96 *An intermediate element that contains computer-generated water and an animated sky.*

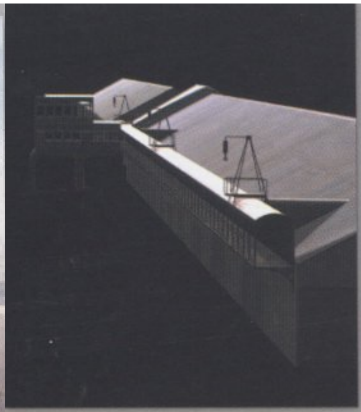

Plate 97 *A computer-generated dock element.*

Plate 98 *An element used to control the atmosphere on the dock.*

Plate 99 *An element featuring people that were on the ship.*

Plate 100 *An element featuring a group of people on the dock.*

Page layout, magazine covers

Photo editing

- **Edit the background independently from foreground**

Photo editing

- **Edit the background independently from foreground**

Technical Issues

- **Compositing**
  - How exactly do we handle transparency?
- **Smart selection**
  - Facilitate the selection of an object
- **Matte extraction**
  - Resolve sub-pixel accuracy, estimate transparency
- **Smart pasting**
  - Don't be smart with copy, be smart with paste
  - Example: pyramid splining (Burt and Adelson)
  - Example: gradient domain (Poisson blending)
- **Extension to video**
  - Where life is always harder

Key Idea: adding an Alpha channel

- α: 1 means opaque, 0 means transparent
- 32-bit images: R, G, B, α

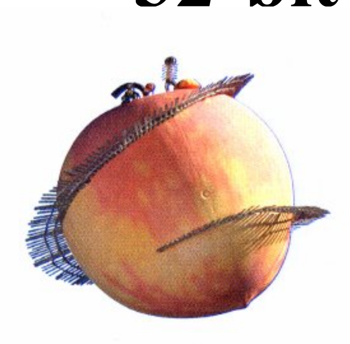

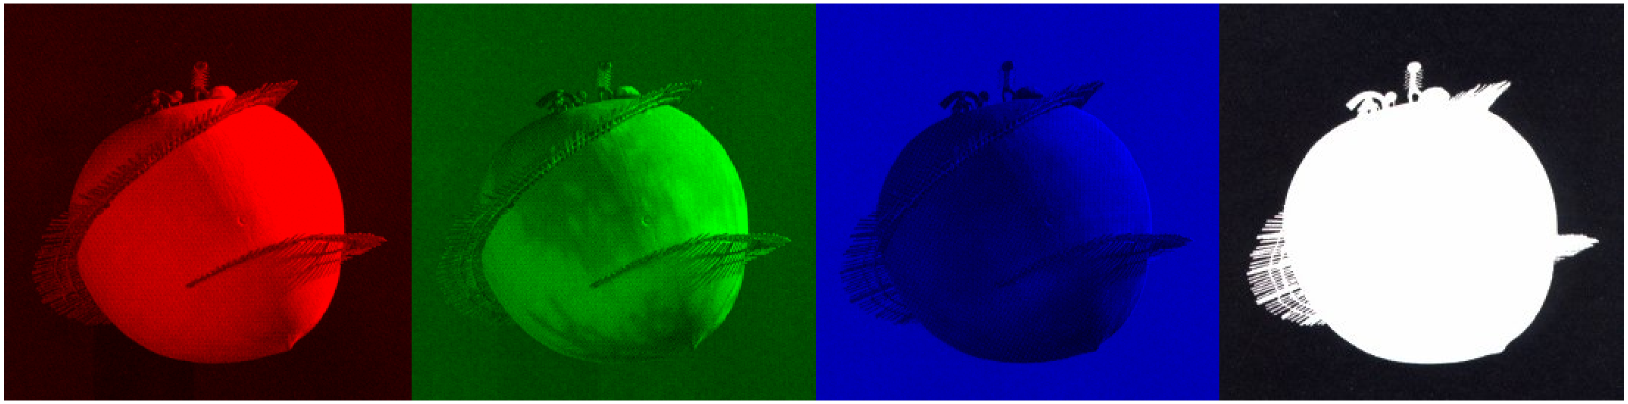

From the Art & Science of Digital Compositing

Photoshop layer masks

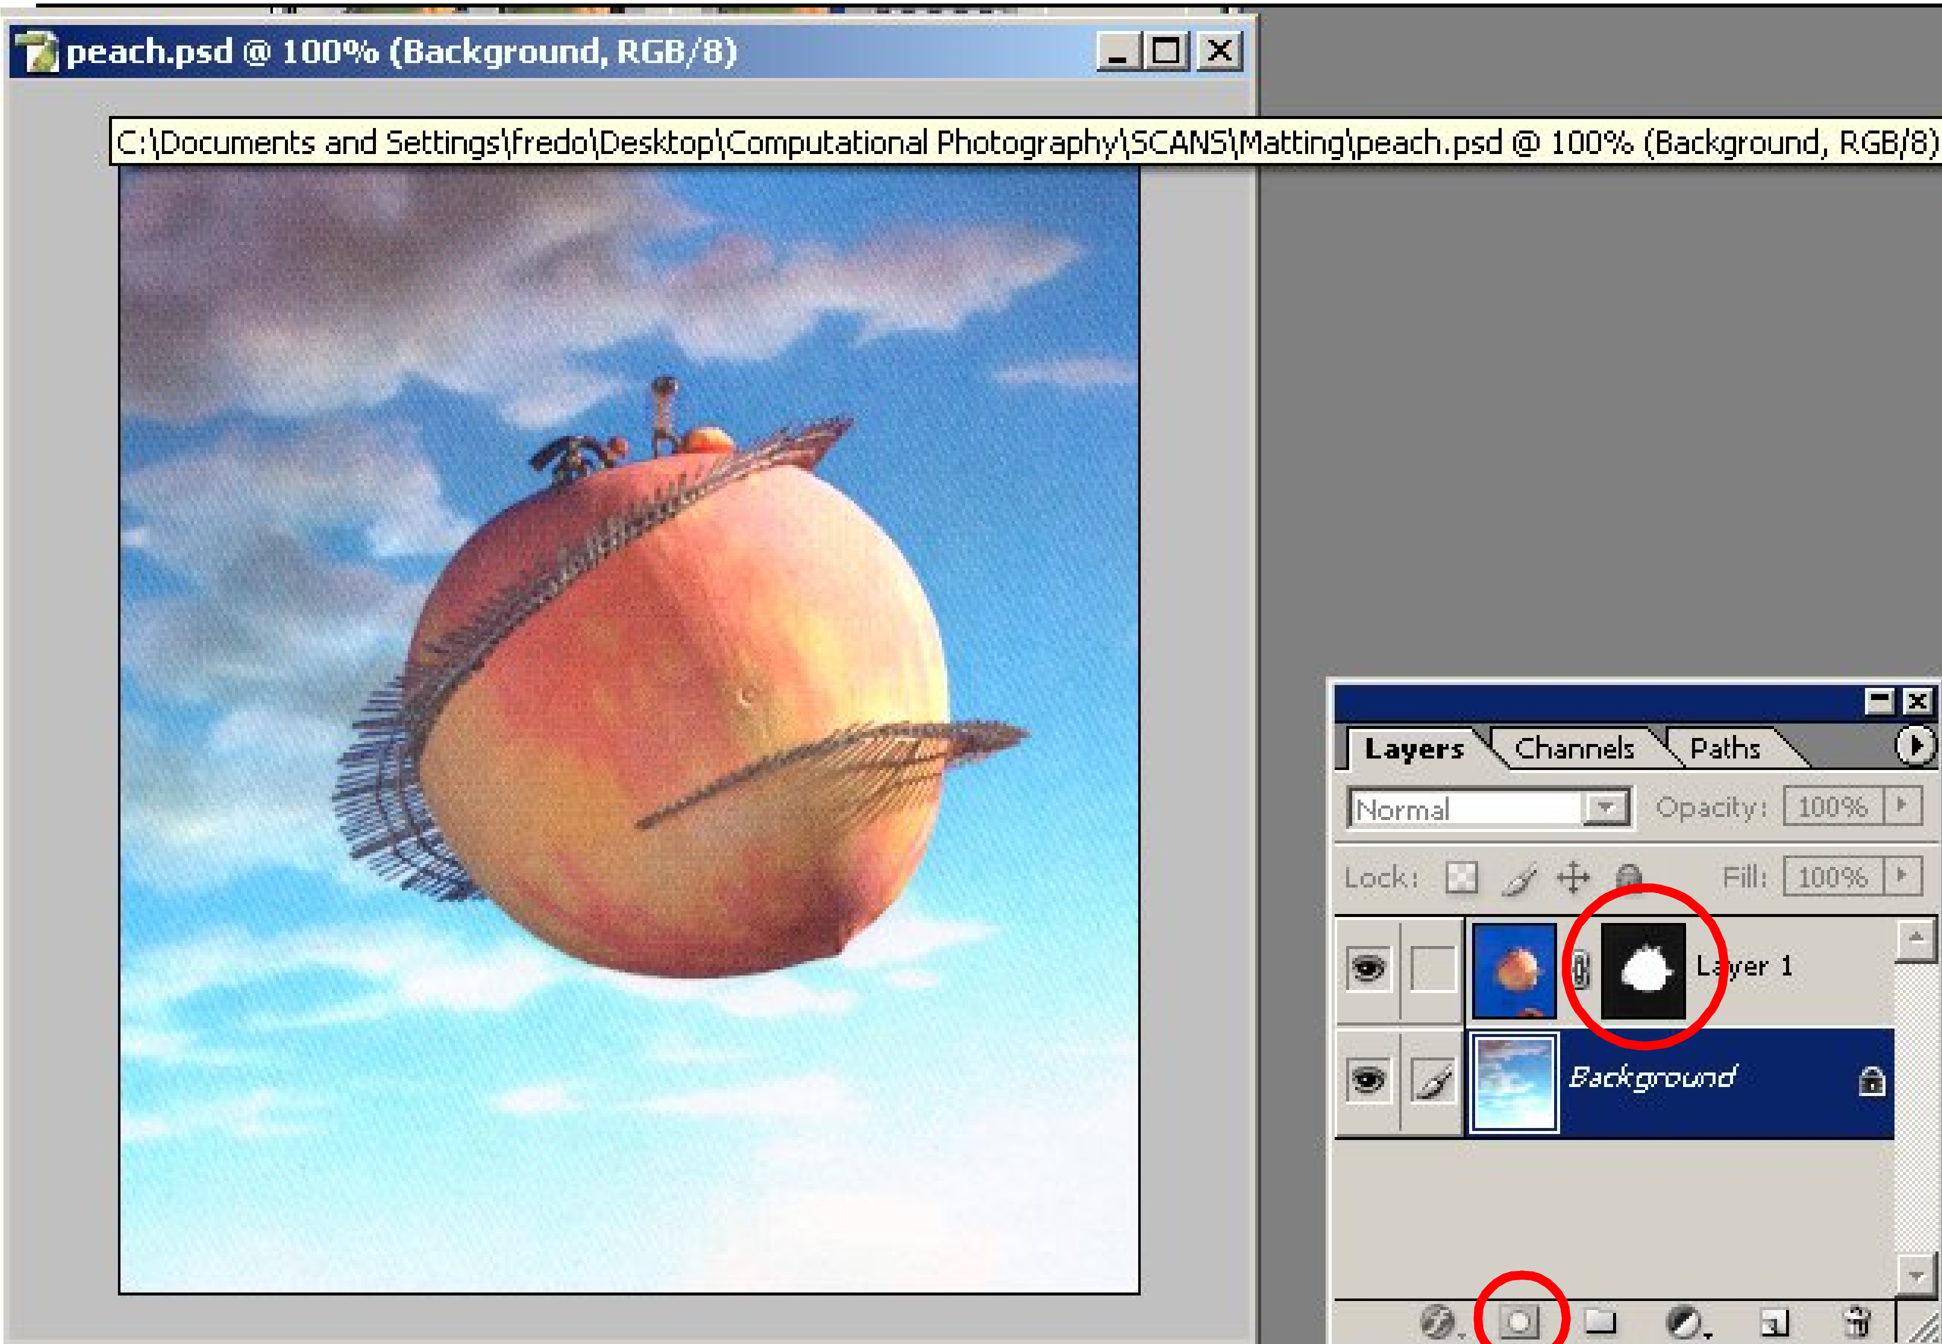

Compositing

F **α** **B** **(1-α)**

 *  + 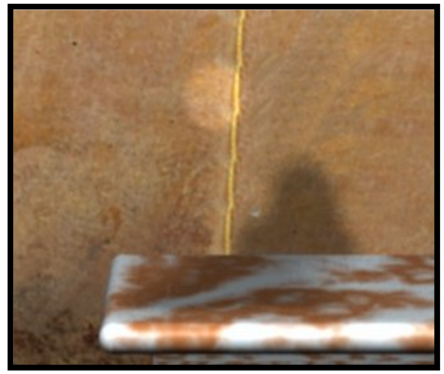 * 

Foreground **Traveling Matte** **Background** **Holdout Matte**

Fundamental equation:

$$C=\alpha F+(1-\alpha) B$$

= 

C

Slide from Pat Hanrahan

Why fractional alpha?

- **Thin features (e.g. hair) cause mixed pixels**

Alpha Matte

Composite

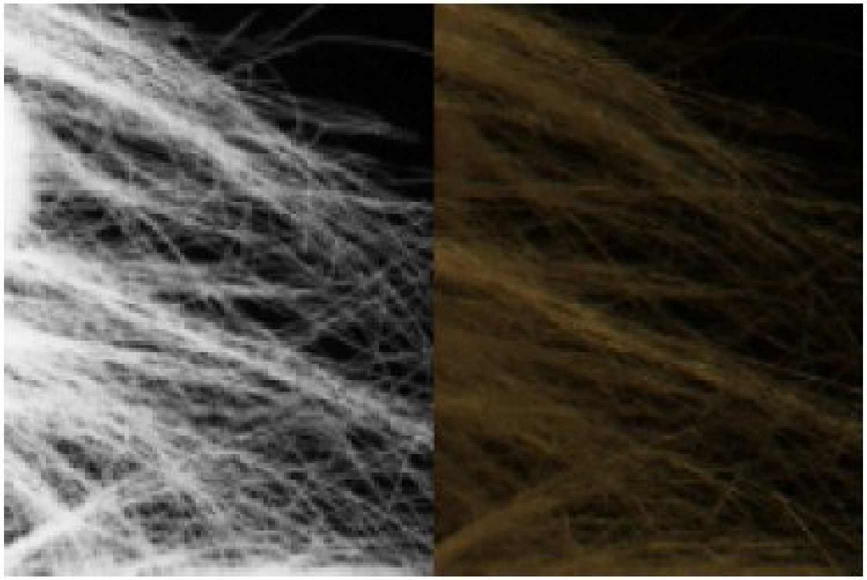

Inset

Why fractional alpha?

- **Motion blur "smears" foreground into background**

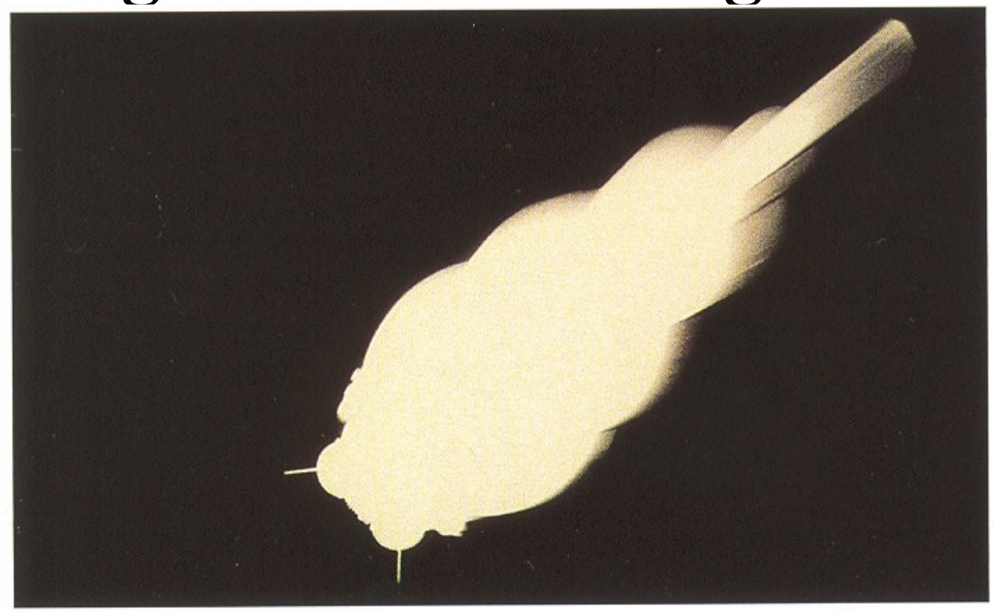

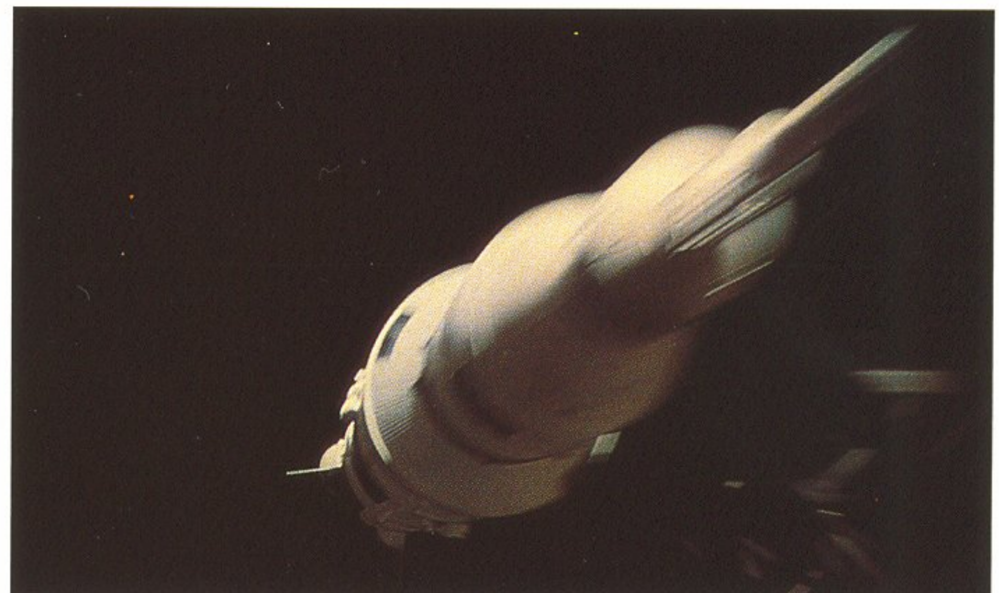

From Digital Domain

With binary alpha

From Digital Domain

With fractional alpha

From Digital Domain

Why fractional alpha?

- **Handling (semi)transparent objects**

From Smith & Blinn's SIGGRAPH'96 paper

Compositing

- ***The variables of interest:***
  Given the foreground color F=(Fr, Fg, Fb), the background color (Br, Bg, Bb) and α for each pixel
- **The compositing operation is: $C=\alpha F+(1-\alpha)B$**

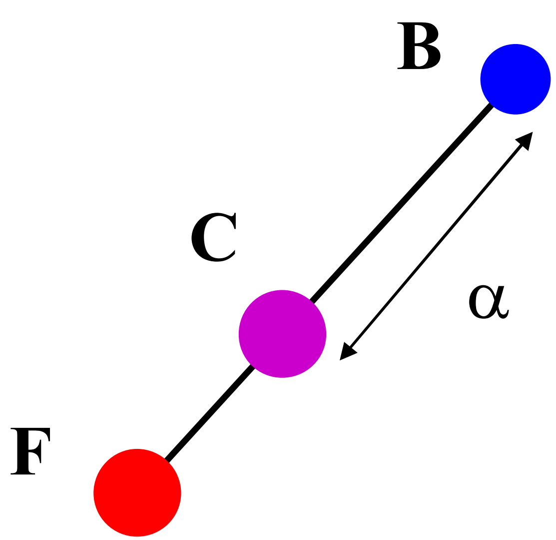

Note: $0 <= \alpha <= 1$ interpolates a color C on the line between F and B

Matting problem

- **Inverse problem:**
  Assume an image is the α–composite of a foreground and a background
- **Given an image color C, find F, B and α so that $C=\alpha F+(1-\alpha)B$**

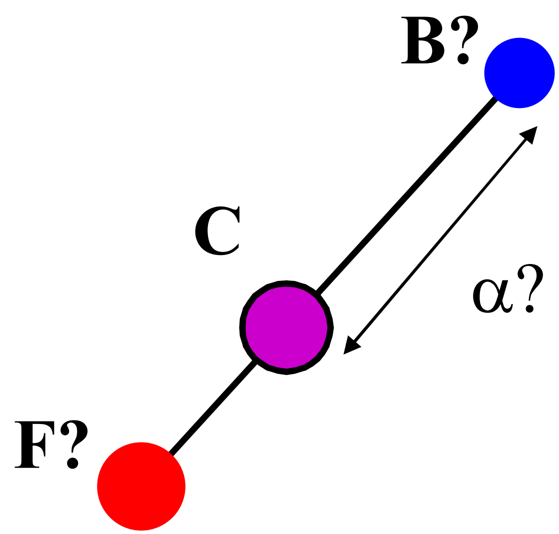

Why Matting is Hard…

- **$C=\alpha F+(1-\alpha)B$**
- **How many unknowns, how many equations?**

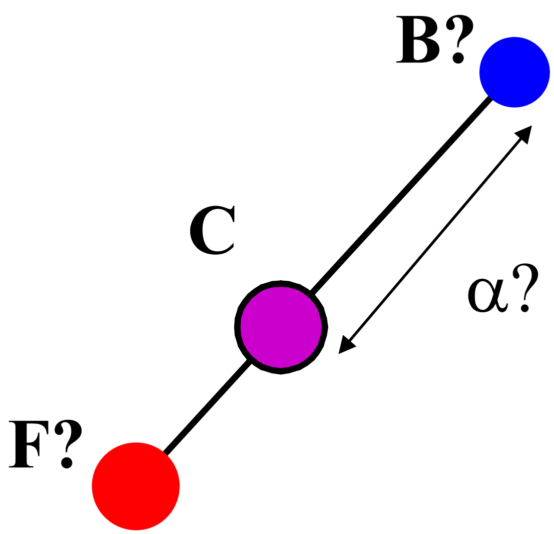

$$C_r = \alpha F_r + (1-\alpha) B_r$$
$$C_g = \alpha F_g + (1-\alpha) B_g$$
$$C_b = \alpha F_b + (1-\alpha) B_b$$

- **7 unknowns, 3 equations**
- **Bottom line: we need fewer unknowns (or more equations)**

Traditional blue screen matting

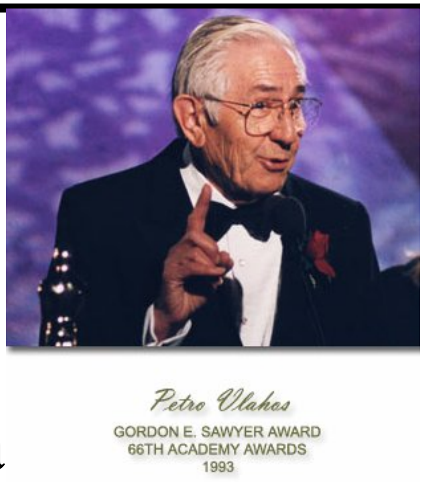

- Invented by Petro Vlahos (Technical Academy Award 1995)
- formalized by Smith & Blinn
- Initially for film, then video, then digita
- Assume that the foreground has no blue
- Assume background is mainly blue

From Cinefex

Example

Example

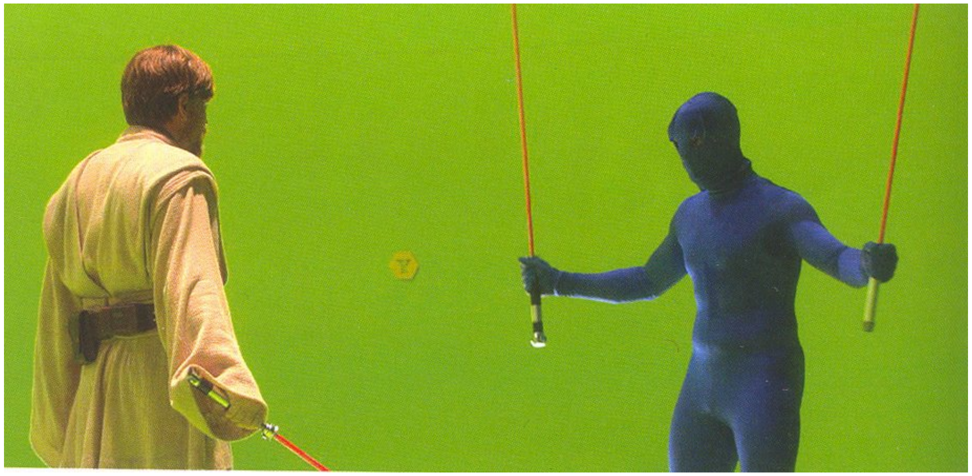

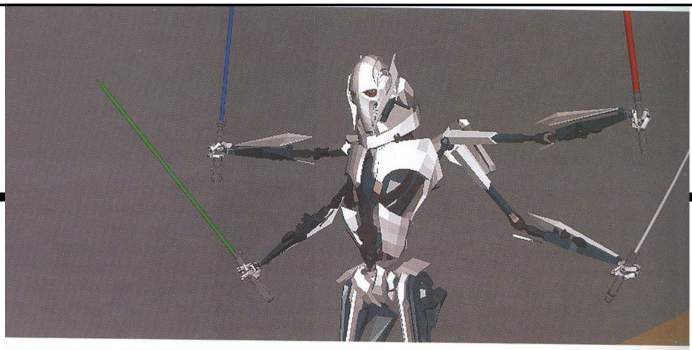

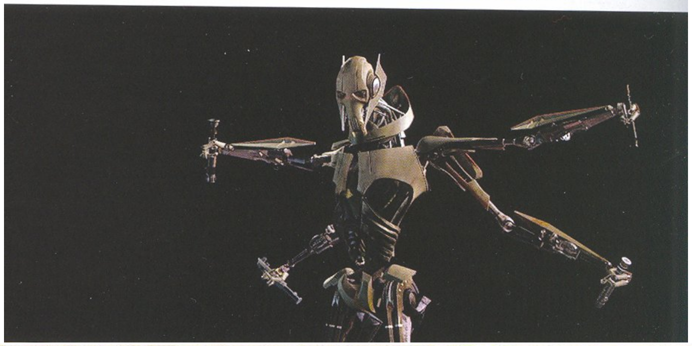

From Cinefex

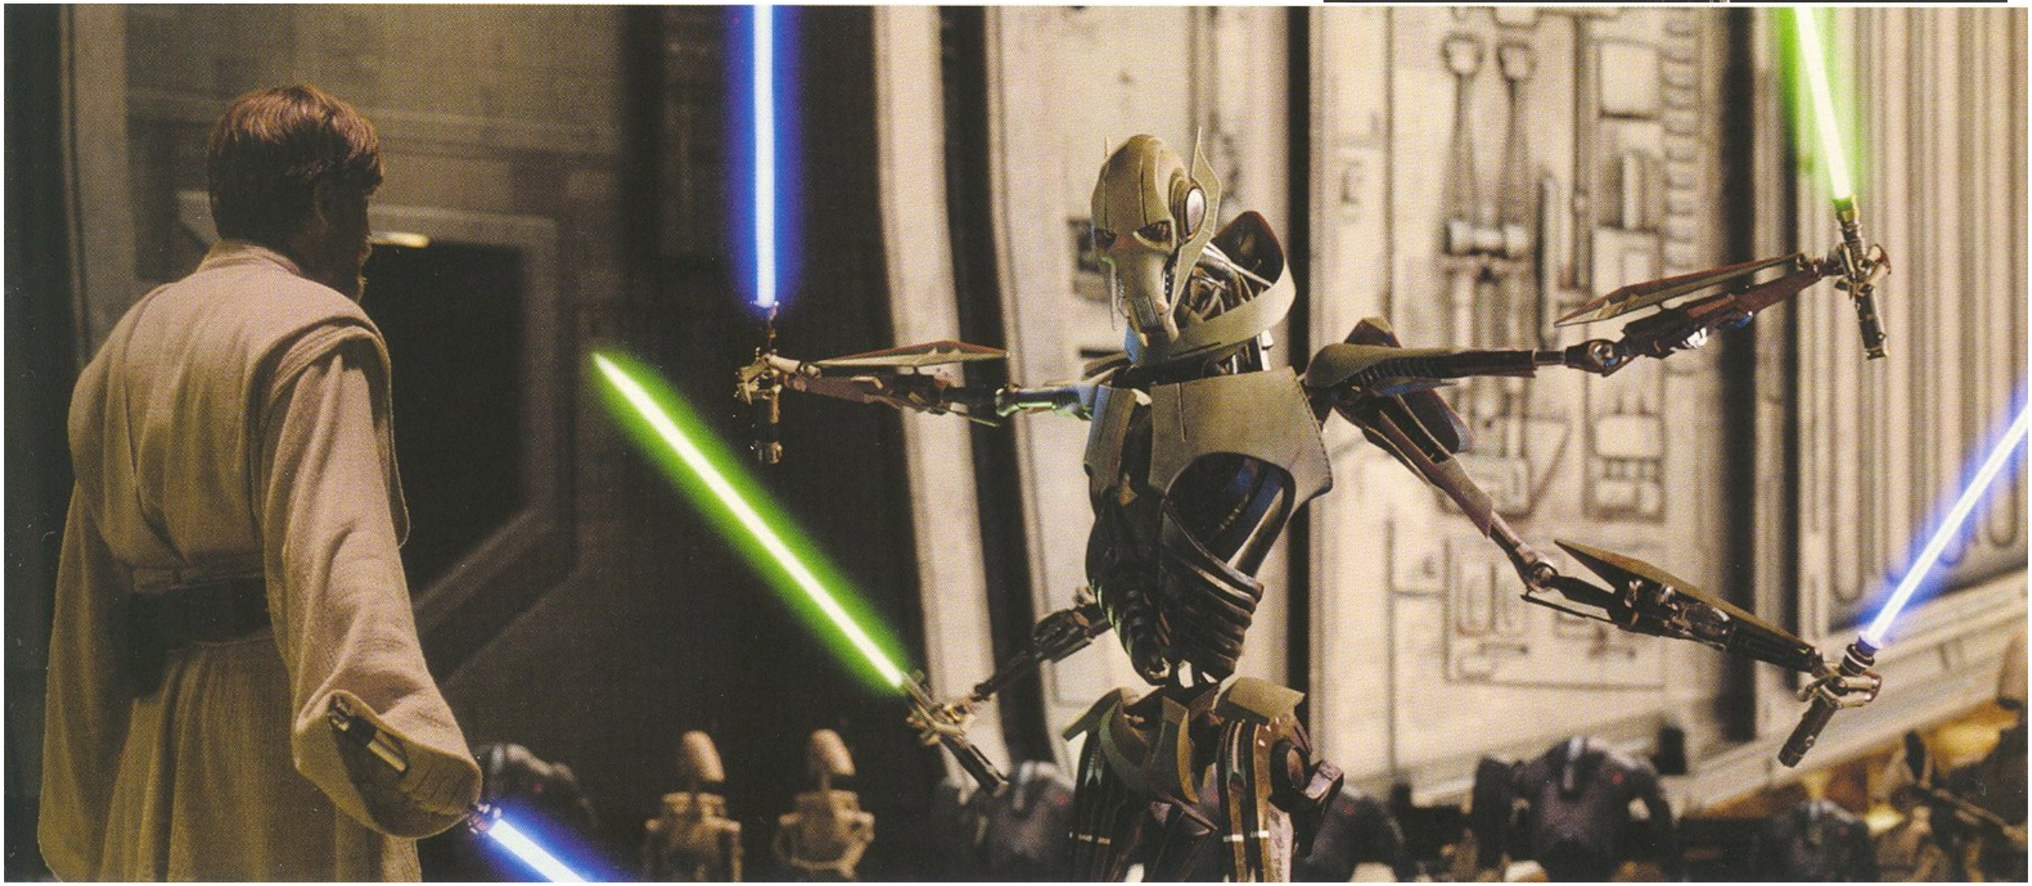

How blue screen works

$$C_r = \alpha F_r + (1-\alpha) B_r$$
$$C_g = \alpha F_g + (1-\alpha) B_g$$
$$C_b = \alpha F_b + (1-\alpha) B_b$$

- **Idealized version:**
  no blue in foreground. Only blue in background

$$F_b = 0\,,\ B_r = B_g = 0$$

- **Equations simplify to**

$$C_r = \alpha F_r$$
$$C_g = \alpha F_g$$
$$C_b = (1-\alpha) B_b$$

- **3 equations in 3 unknowns**

Grey Object or Skin

$$C_r = \alpha F_r + (1-\alpha) B_r$$
$$C_g = \alpha F_g + (1-\alpha) B_g$$
$$C_b = \alpha F_b + (1-\alpha) B_b$$

- **Generalize a little**
  If we assume object is grey:

$$F_r = F_g = F_b = F \; , \; B_r = B_g = 0$$

- **Equations simplify to**

$$C_r = \alpha F$$
$$C_g = \alpha F$$
$$C_b = \alpha F + (1-\alpha) B_b$$

- **Similar simplification if skin color:**
  F ~ (k , k/2 , k/2)

Blue/Green screen matting issues

- **Color limitation**
  - Annoying for blue-eyed people
  - ➔ adapt screen color (in particular green)
- **Blue/Green spilling**
  - The background illuminates the foreground, blue/green at silhouettes
  - Modify blue/green channel, e.g. set to min (b, a_2g)
- **Shadows**
  - How to extract shadows cast on background

Blue/Green screen matting issues

Plate 52 *(b) The element placed into the scene without spill suppression. Note the blue fringes on the subject, particularly in the hair.*

From the Art & Science of Digital Compositing

- **http://www.digitalgreenscreen.com/figure3.html**

Figure 3. *Firefox Blue Spill Matte Series 1, original shot. Note blue reflected on wing surfaces from bluescreen -- undesirable but unavoidable on such surfaces.*

Extension: Chroma key

- **Blue/Green screen matting exploits color channels**
- **Chroma key can use an arbitrary background color**
- **See e.g.**
  - http://www.cs.utah.edu/~michael/chroma/
  - Keith Jack, "Video Demystified", Independent Pub Group (Computer), 1996

What about adding more equations?

- **Any ideas?**

What about adding more equations?

- **Any ideas?**
- **Smith and Blinn, Siggraph 1996**
  take pictures in front of two different backgrounds!
- **Triangulation Matting**

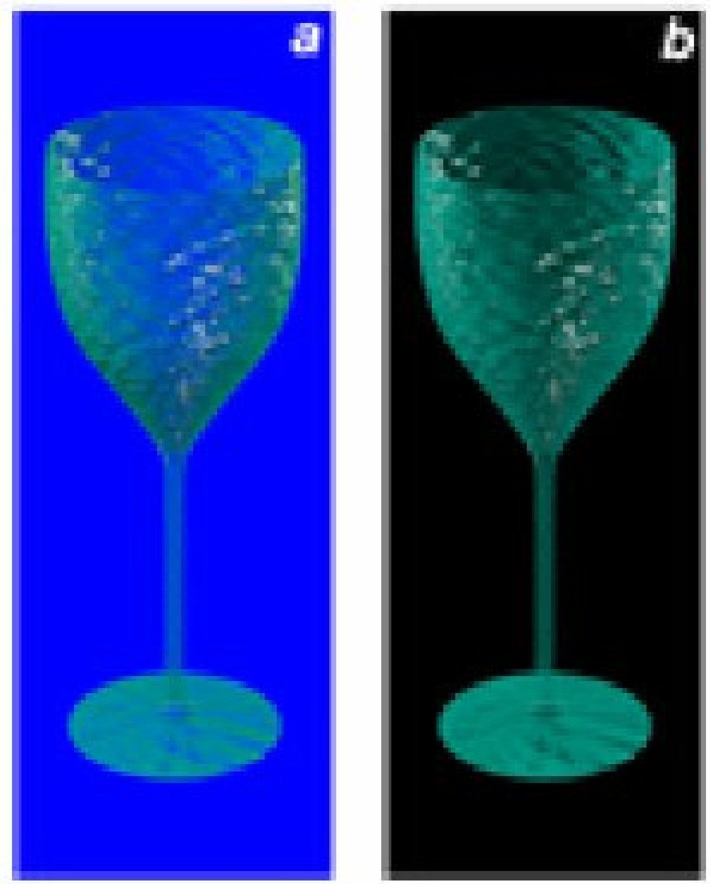

Triangulation Matting

$$C_{r1} = \alpha F_r + (1-\alpha) B_{r1}$$
$$C_{g1} = \alpha F_g + (1-\alpha) B_{g2}$$
$$C_{b1} = \alpha F_b + (1-\alpha) B_{b2}$$
$$C_{r2} = \alpha F_r + (1-\alpha) B_{r2}$$
$$C_{g2} = \alpha F_g + (1-\alpha) B_{g2}$$
$$C_{b2} = \alpha F_b + (1-\alpha) B_{b2}$$

6 equations in 4 unknowns

Triangulation Matting Examples

From Smith & Blinn's
SIGGRAPH'96 paper

More Examples

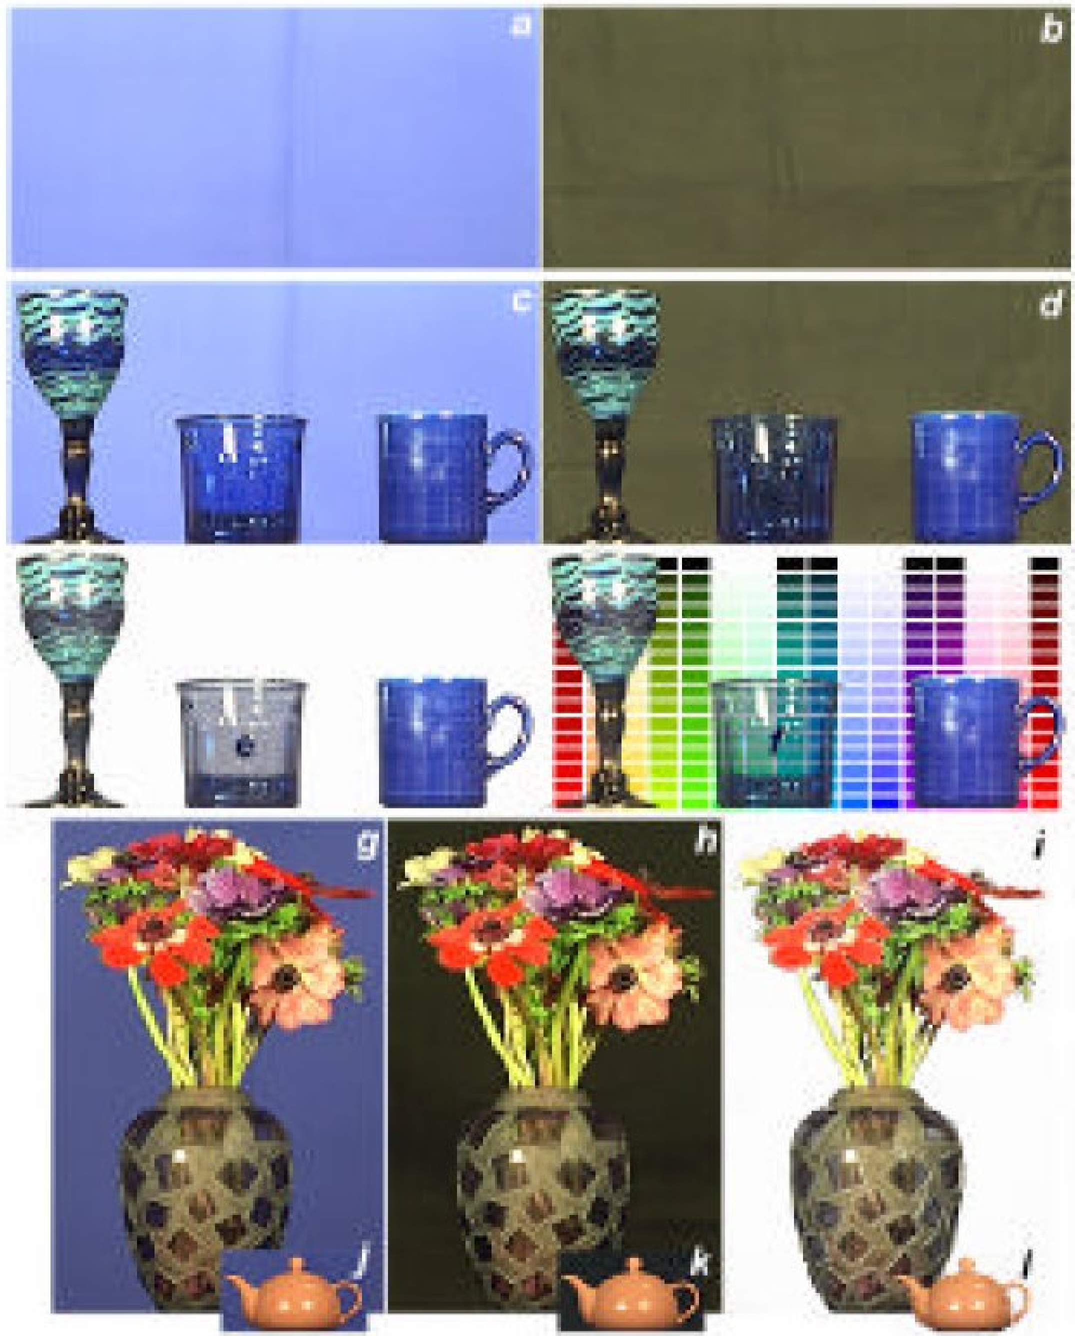

More examples

Side note:

- **Smith&Blinn triangulation approach is used to compute ground truth mattes for comparison in recent matting papers (e.g. Bayesian matting).**

Difference Matting

- e.g. Qian and Sezar
- If we are willing to use two pictures, why don't we take one without the object in it, and take another one with it. Then compare the two.
- Related to background subtraction
- Very useful for video

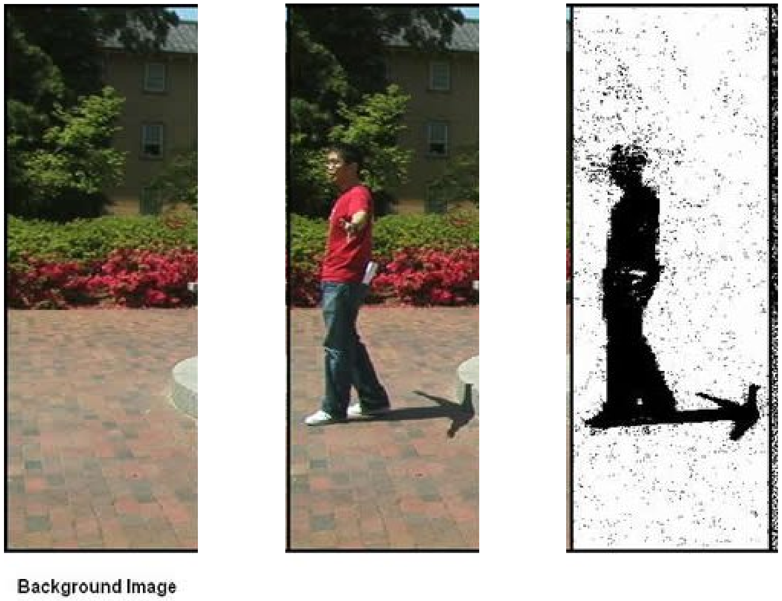

http://www.cs.unc.edu/~lguan/Research.files/backgroundSubtractionResult.JPG

Zhaozheng Yin, CSE Dept, PSU

Shape Constrained Figure-Ground Segmentation

The conditional probability is

$$P(s|I) = \frac{1}{Z}\exp\left\{\sum_{i\in A}\sum_{k}\lambda_k \underline{\Phi_k(s_i, I)} + \sum_{(i,j)\in B}\underline{\Psi(s_i, s_j, I)}\right\}$$

$$\propto \frac{1}{Z}\exp\{-E\}$$

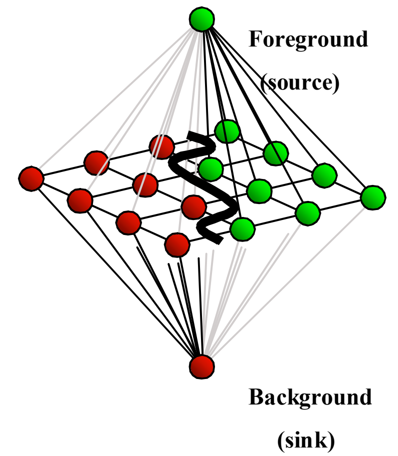

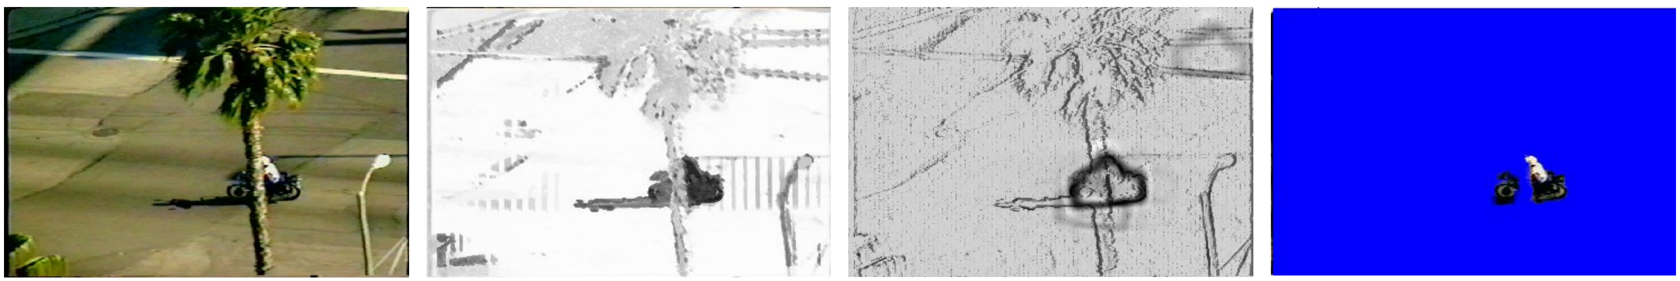

Data term (motion, contrast, color, temporal consistency etc):

$$E_k(s_i, I) = -\Phi_k(s_i, I) = -\log p(s_i|f_k(I))$$

link terms are based on edge gradients, as well as previously learned object shape information

Solution (foreground mask) computed using graph cuts.

Shape Constrained Figure-Ground Segmentation

In natural images, the transition between foreground and background usually happens gradually, we use Random-Walk matting (Grady 2005) to assign foreground opacity to those uncertain pixels.

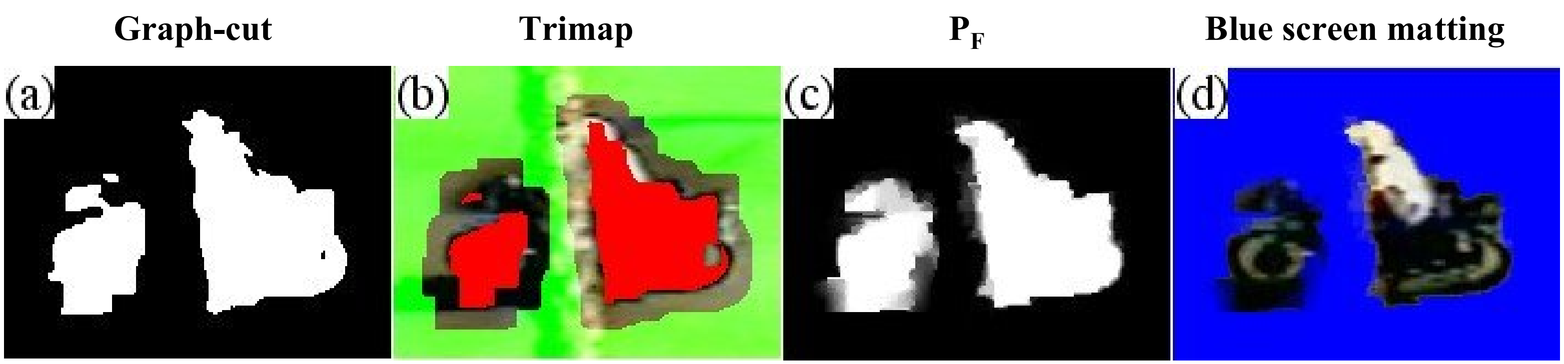

Shape Constrained Figure-Ground Segmentation

Results

Shape Constrained Figure-Ground Segmentation

Human Body Shape Learning

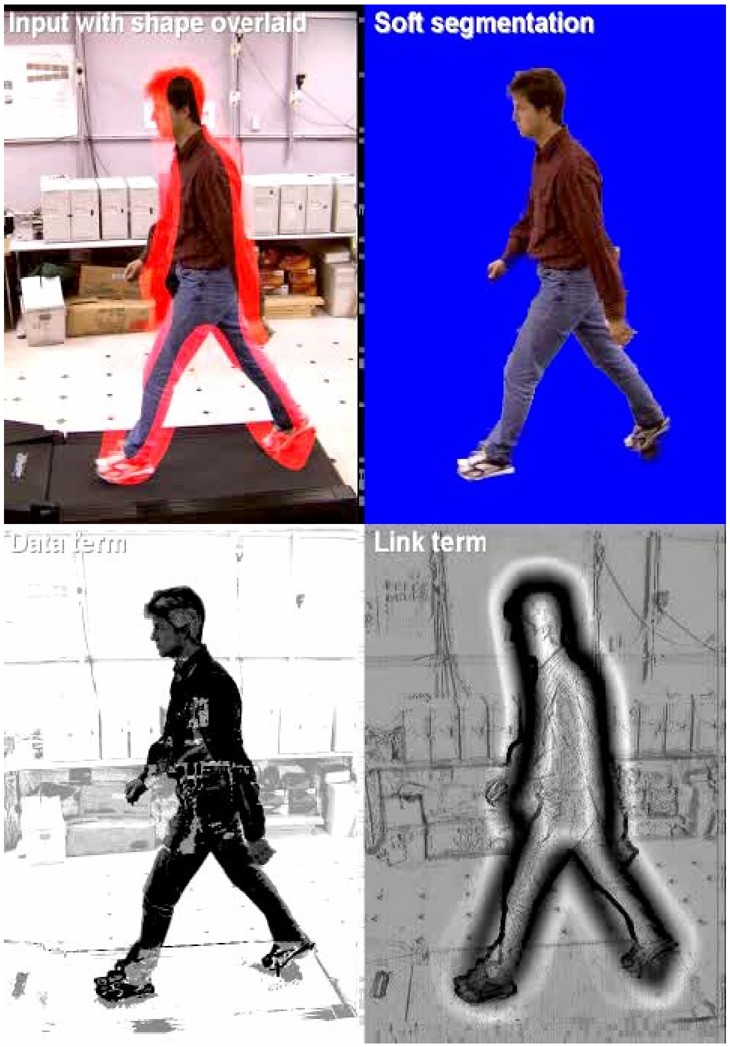

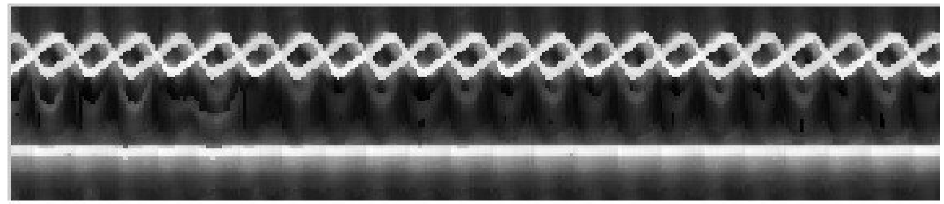

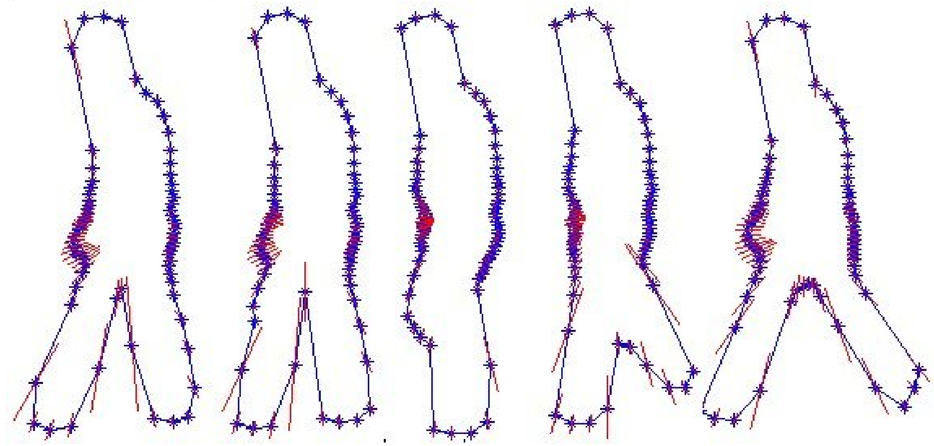

Blue: mean; Red: variance

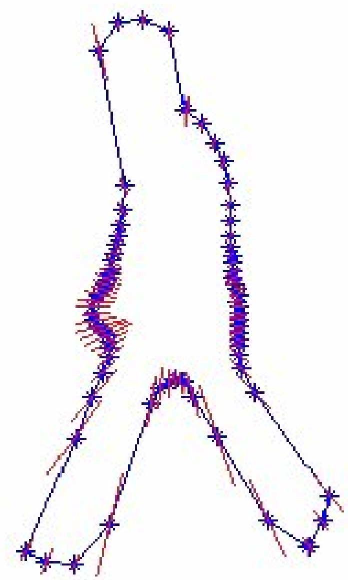

Shape Constrained Figure-Ground Segmentation

Apply the learned shape for segmentation

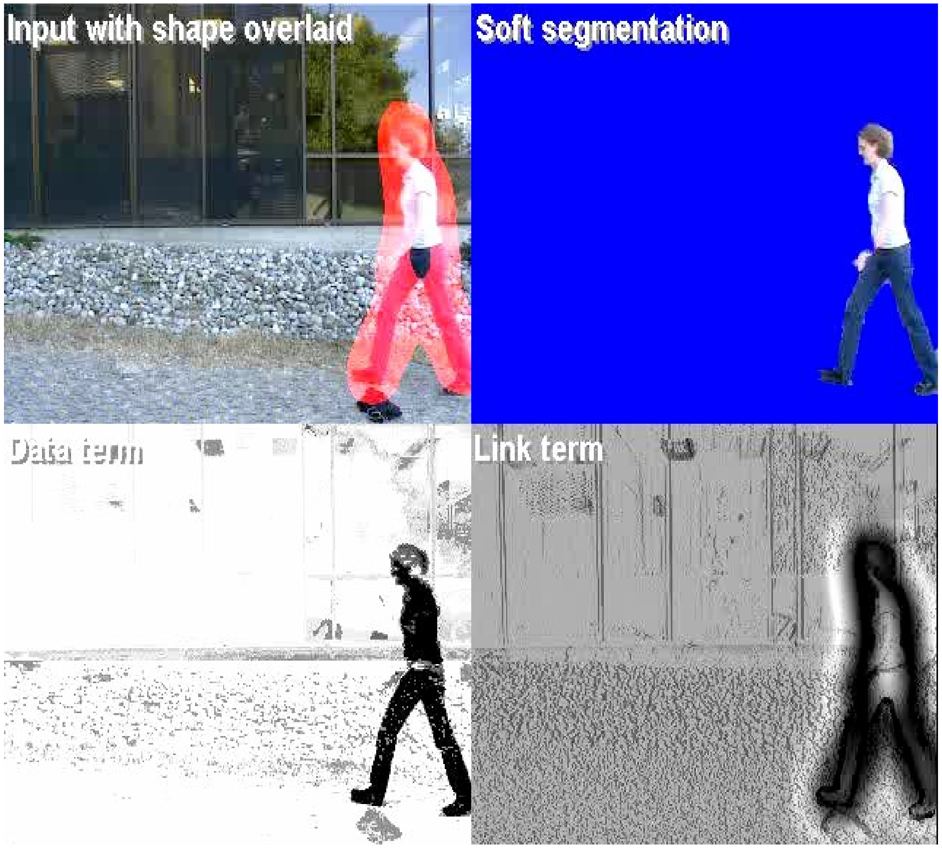

Video editing:

Shape Constrained Figure-Ground Segmentation

Real-time demo, using color, edges and stereo (depth)

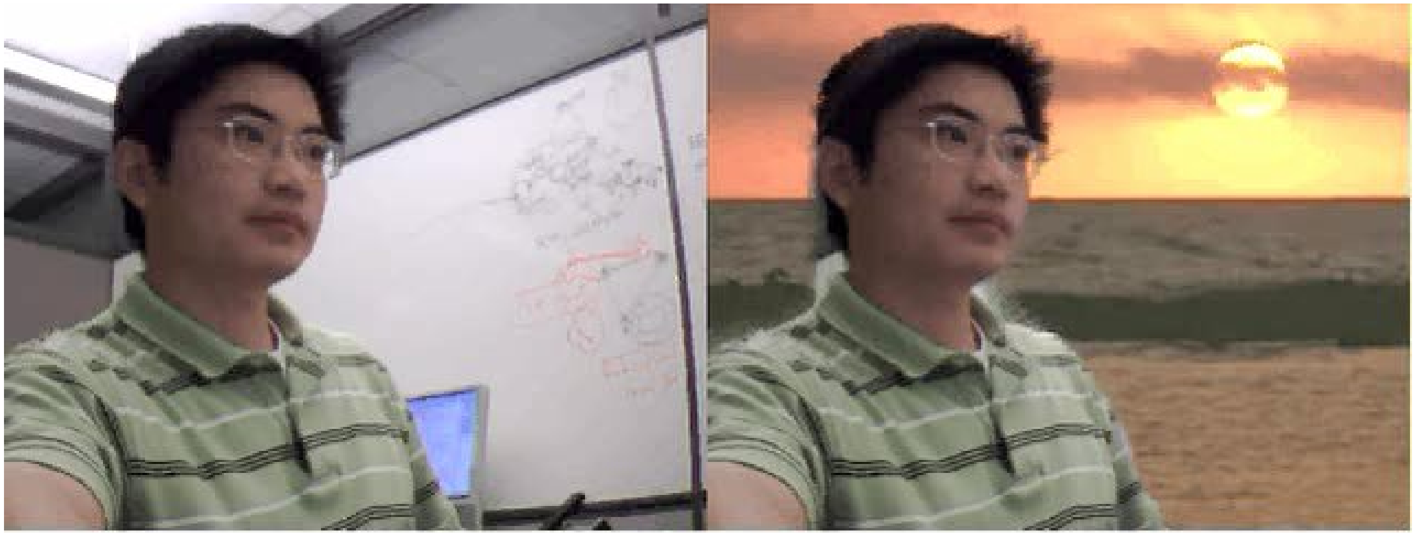

Natural Image Matting

- **Works for single image**
- **Background/foreground not known in advance**
- **Need "hints" from the user, in form of a <u>trimap</u>**

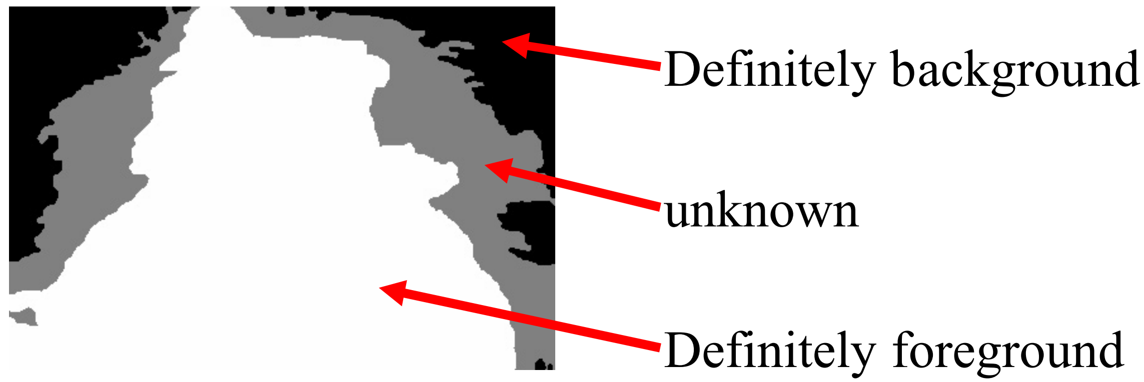

- **General idea: compute probability distributions of foreground and background color near unknown points and use them to determine alpha, F and B.**

Collecting Fg/Bg Samples

- For each unknown pixel, collect samples of nearby labeled foreground and background pixels

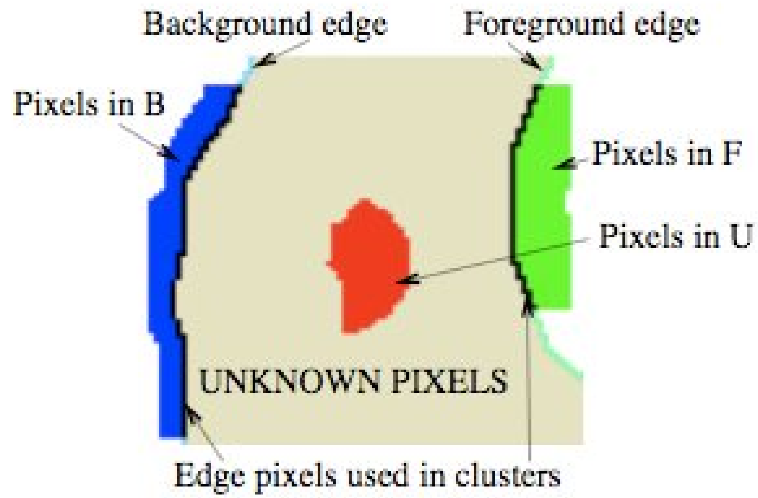

- Estimate distributions P(F) and P(B) using your favorite parameteric or nonparametric method

Ruzon and Tomasi

- **Estimate distributions as mixtures of Gaussians with spherical covariance matrices**
- **Group Gaussian clusters into pairs (pi,qi) where pi is from P(F) and qi is from P(B). Some unlikely pairs are removed using heuristic constraints.**

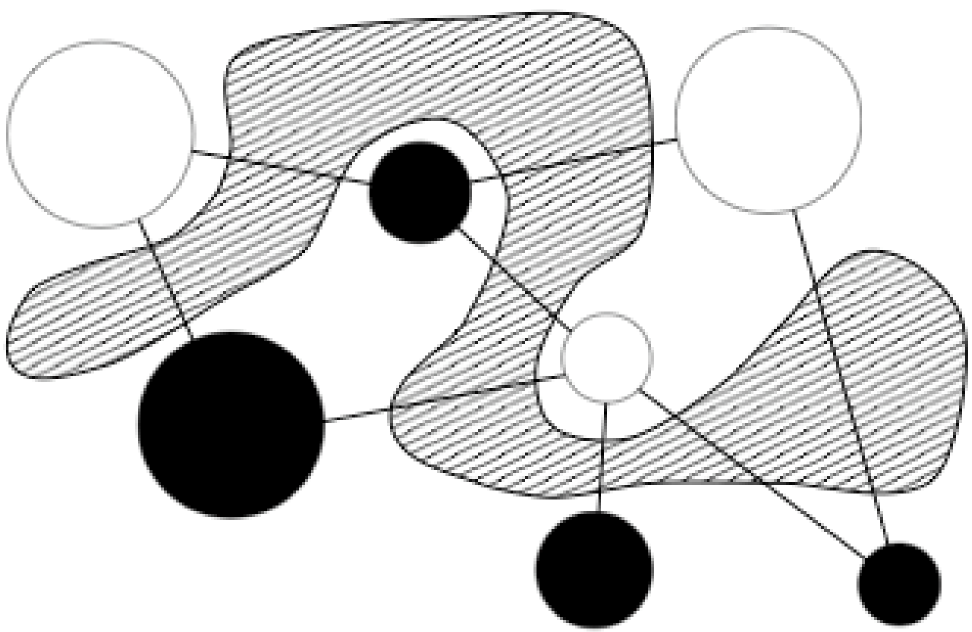

For instance:
White are foreground components
Black are background components

The line segments connect pairs of clusters that can “go together”

Ruzon and Tomasi

- **For an unknown color C, we'd like to figure out its alpha value, by aggregating information across the pairs of clusters**
- **Insight: C is drawn from a distribution that represents a "morph" between a foreground and background color cluster pair.**

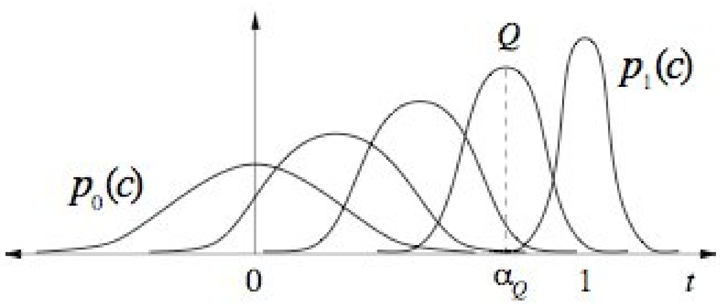

Figure 5. A 1-D example of interpolation. As t varies, the mean and variance of a Gaussian interpolates between $p_0(c)$ and $p_1(c)$. The value of t that maximizes the value at Q is the alpha value α_Q.

So given a cluster pair, interpolate means/variances between the two with parameter 0<=t<=1.

The interpolated Gaussian that yields the highest likelihood of color C is chosen, and argmax(t) becomes our estimate of alpha!

Ruzon and Tomasi

- **Since we don’t know which Fg/Bg color cluster pair to use, we combine results for alpha across all feasible pairs and take argmax of that function instead.**
- **After computing alpha, F and B are determined by weighted combination of cluster pairs.**

examples

examples

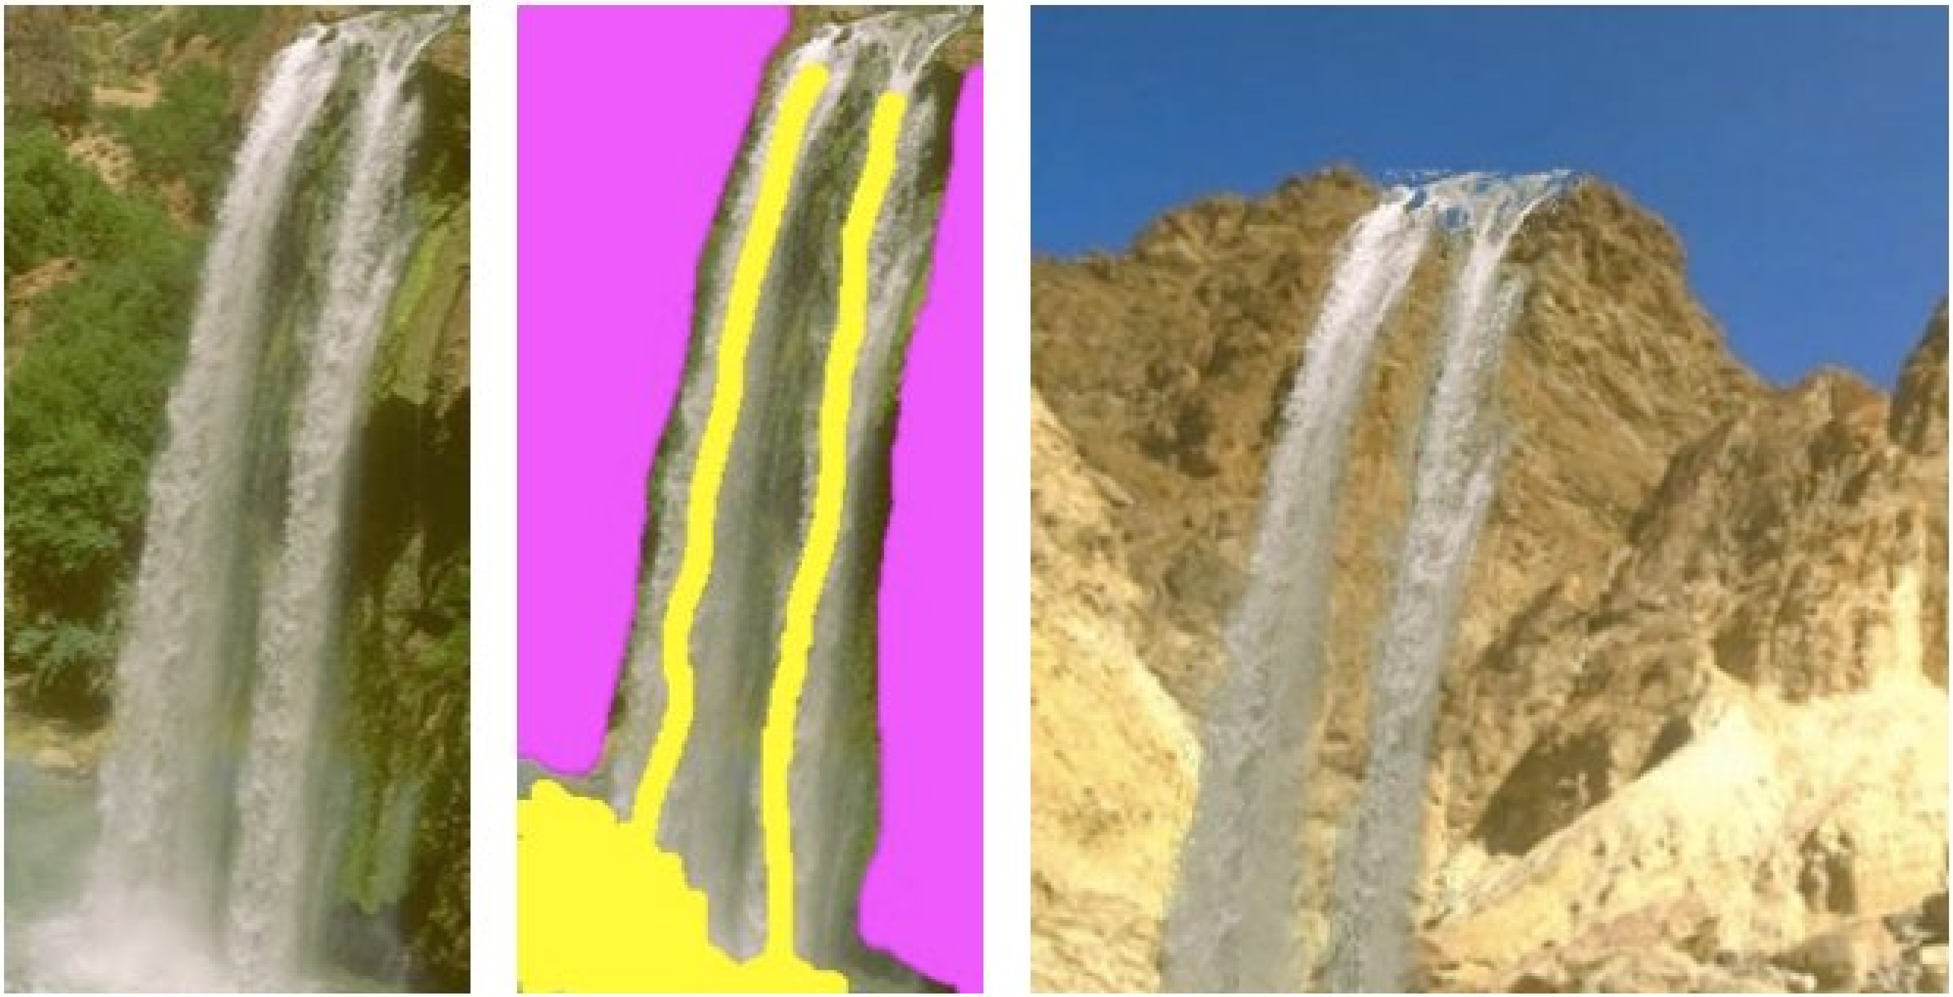

Hillman et.al. 2001

- **Note that color clusters tend not to be spherical**

  - for instance, same hue but diff intensity leads to elongated clusters along the rgb "diagonal"
- **Define cluster by a line segment in color space**

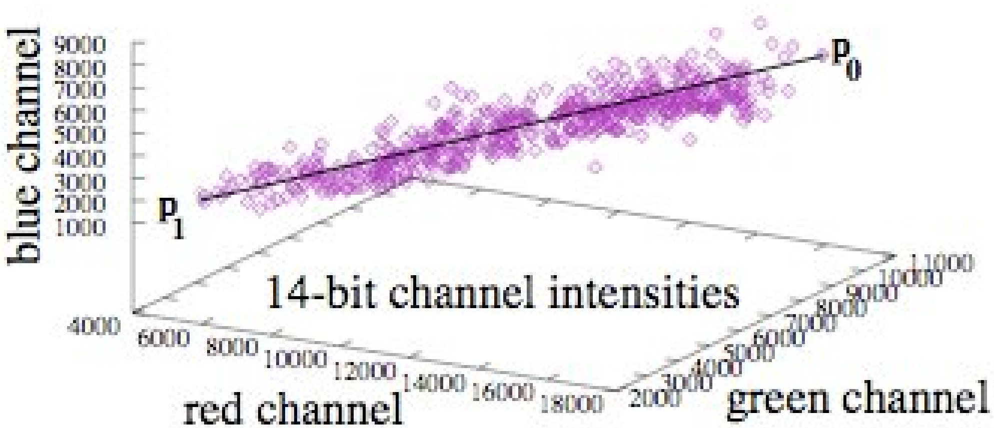

Figure 3. Cluster of points in RGB space with the line $\overline{p_0 p_1}$

Hillman et.al. 2001

- One line segment represents foreground colors and another line segment represents background colors
- Given unknown color C, find colors F and B that lie closest to it on the two lines
- Project C onto segment F-B, and compute its alpha.

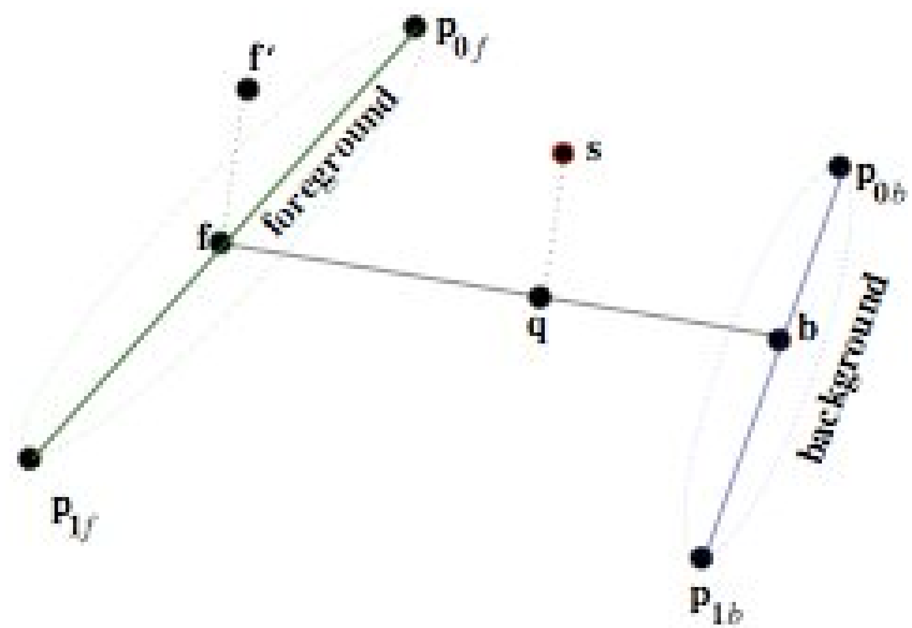

Bayesian Matting

- **Chuang et.al. 2001**
- **More principled method than Hillman**
- **Similar to Ruzon+Tomasi, but allows for elongated clusters**

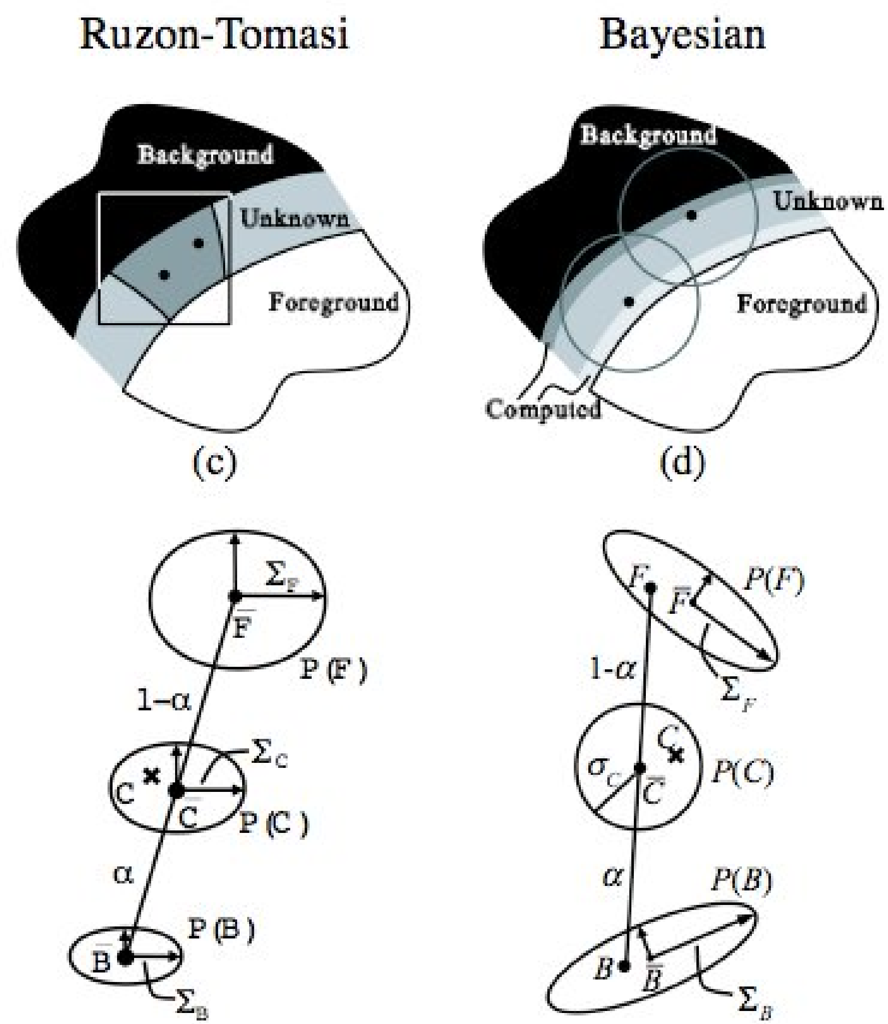

Bayes theorem

$$P(x|y) = P(y|x)\, P(x) \,/\, P(y)$$

The parameters you want to estimate

What you observe

Likelihood function

Prior probability

Constant w.r.t. parameters x.

Matting and Bayes

- **What do we observe?**

$$P(x|\mathbf{y}) = P(\mathbf{y}|x)\, P(x) \,/\, P(\mathbf{y})$$

The parameters you want to estimate

What you observe

Likelihood function

Prior probability

Constant w.r.t. parameters x.

Matting and Bayes

- **What do we observe?**
  - Color C at a pixel

$$P(x|\mathbf{C}) = P(\mathbf{C}|x)\, P(x) \,/\, P(\mathbf{C})$$

The parameters you want to estimate

Color you observe

Likelihood function

Prior probability

Constant w.r.t. parameters x.

Matting and Bayes

- What do we observe: Color C
- **What are we looking for?**

$$P(\mathbf{x}|C) = P(C|\mathbf{x})\, P(\mathbf{x}) \,/\, P(C)$$

The parameters you want to estimate

Color you observe

Likelihood function

Prior probability

Constant w.r.t. parameters x.

Matting and Bayes

- What do we observe: Color C
- **What are we looking for: F, B,** α

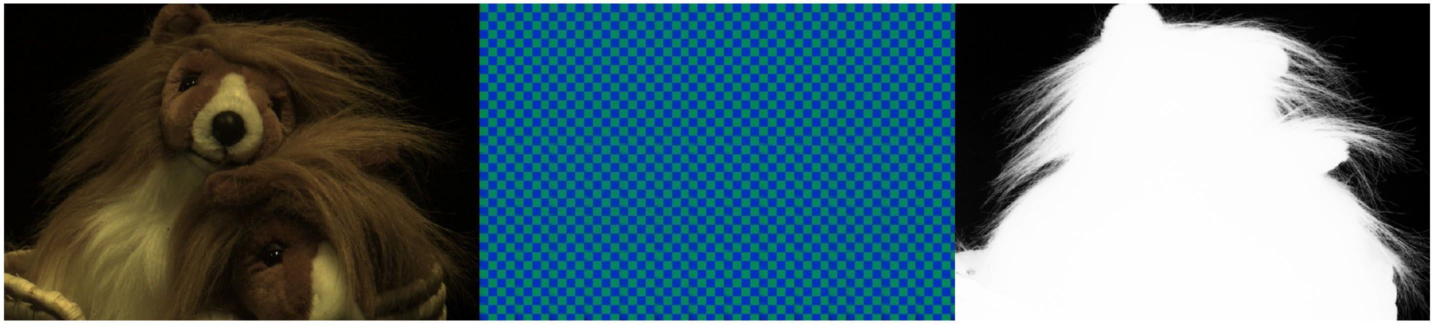

$$P(\mathbf{F},\mathbf{B},\alpha|C) = P(C|\mathbf{F},\mathbf{B},\alpha)\ P(\mathbf{F},\mathbf{B},\alpha)\ /\ P(C)$$

Foreground, background, transparency you want to estimate

Color you observe

Likelihood function

Prior probability

Constant w.r.t. parameters x.

Matting and Bayes

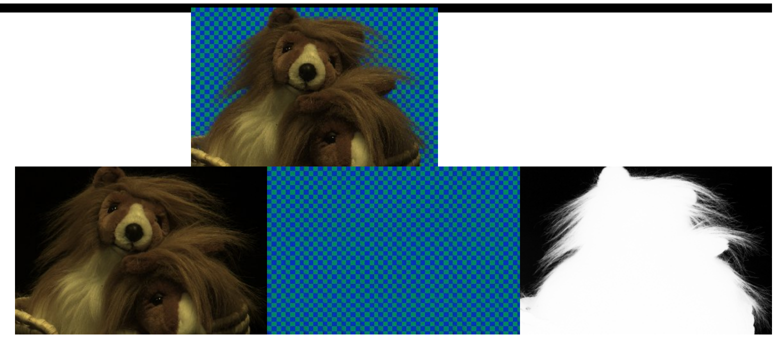

- What do we observe: Color C
- What are we looking for: F, B, α
- **Likelihood probability?**
  - Given F, B and Alpha, probability that we observe C

$$P(F,B,\alpha|C) = \mathbf{P(C|F,B,}\alpha\mathbf{)}\ P(F,B,\alpha) / P(C)$$

Foreground, background, transparency you want to estimate

Color you observe

Likelihood function

Prior probability

Constant w.r.t. parameters x.

Matting and Bayes

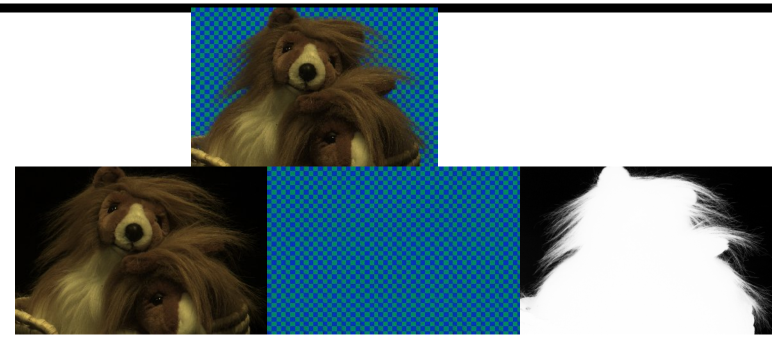

- What do we observe: Color C
- What are we looking for: F, B, α
- **Likelihood probability?**
  - Given F, B and Alpha, probability that we observe C
  - If measurements are perfect, non-zero only if $C=\alpha F+(1-\alpha)B$
  - But assume Gaussian noise with variance σ_C

$$P(F,B,\alpha|C) = \mathbf{P(C|F,B,}\alpha\mathbf{)}\ P(F,B,\alpha)\ /\ P(C)$$

Foreground, background, transparency you want to estimate

Color you observe

Likelihood function

Prior probability

Constant w.r.t. parameters x.

Matting and Bayes

- What do we observe: Color C
- What are we looking for: F, B, α
- Likelihood probability: Compositing equation + Gaussian noise with variance σ_C
- **Prior probability:**
  - How likely is the foreground to have color F? the background to have color B? transparency to be α?

$$P(F,B,\alpha|C) = P(C|F,B,\alpha)\ \mathbf{P(F,B,\alpha)}\ /\ P(C)$$

Foreground, background, transparency you want to estimate

Color you observe

Likelihood function

Prior probability

Constant w.r.t. parameters x.

Matting and Bayes

- What do we observe: Color C
- What are we looking for: F, B, α
- Likelihood probability: Compositing equation + Gaussian noise with variance σ_C
- **Prior probability:**
  Build a probability distribution from the known regions
  - ***This is the heart of Bayesian matting***

$$P(F,B,\alpha|C) = P(C|F,B,\alpha)\, \mathbf{P(F,B,\alpha)} \,/\, P(C)$$

Foreground, background, transparency you want to estimate

Color you observe

Likelihood function

Prior probability

Constant w.r.t. parameters x.

Let's derive it

- **Assume F, B and α are independent**

$$P(F,B,\alpha|C) = P(C|F,B,\alpha)\, P(F,B,\alpha) \,/\, P(C)$$
$$= P(C|F,B,\alpha)\, P(F)\, P(B)\, P(\alpha)/P(C)$$

- **But multiplications are hard!**
- **Make life easy, work with log probabilities**
  L means log P here:

$$L(F,B,\alpha|C) = L(C|F,B,\alpha) +$$
$$L(F) + L(B) + L(\alpha) - L(C)$$

- **And ignore L(C) because it is constant**

Log Likelihood: $L(C|F,B,\alpha)$

- **Gaussian noise model:** $$e^{\frac{-\text{color difference}^2}{\sigma_C^2}}$$
- **Take the log:** $$L(C|F,B,\alpha) = -\|C - \alpha F - (1-\alpha) B\|^2 / \sigma^2_C$$
- **Unfortunately not quadratic in all coefficients (product α B)**

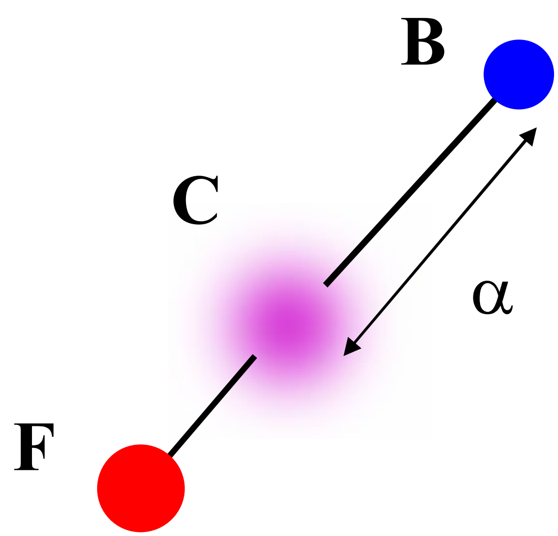

Prior probabilities L(F) & L(B)

- **Gaussians based on pixel color from known regions**

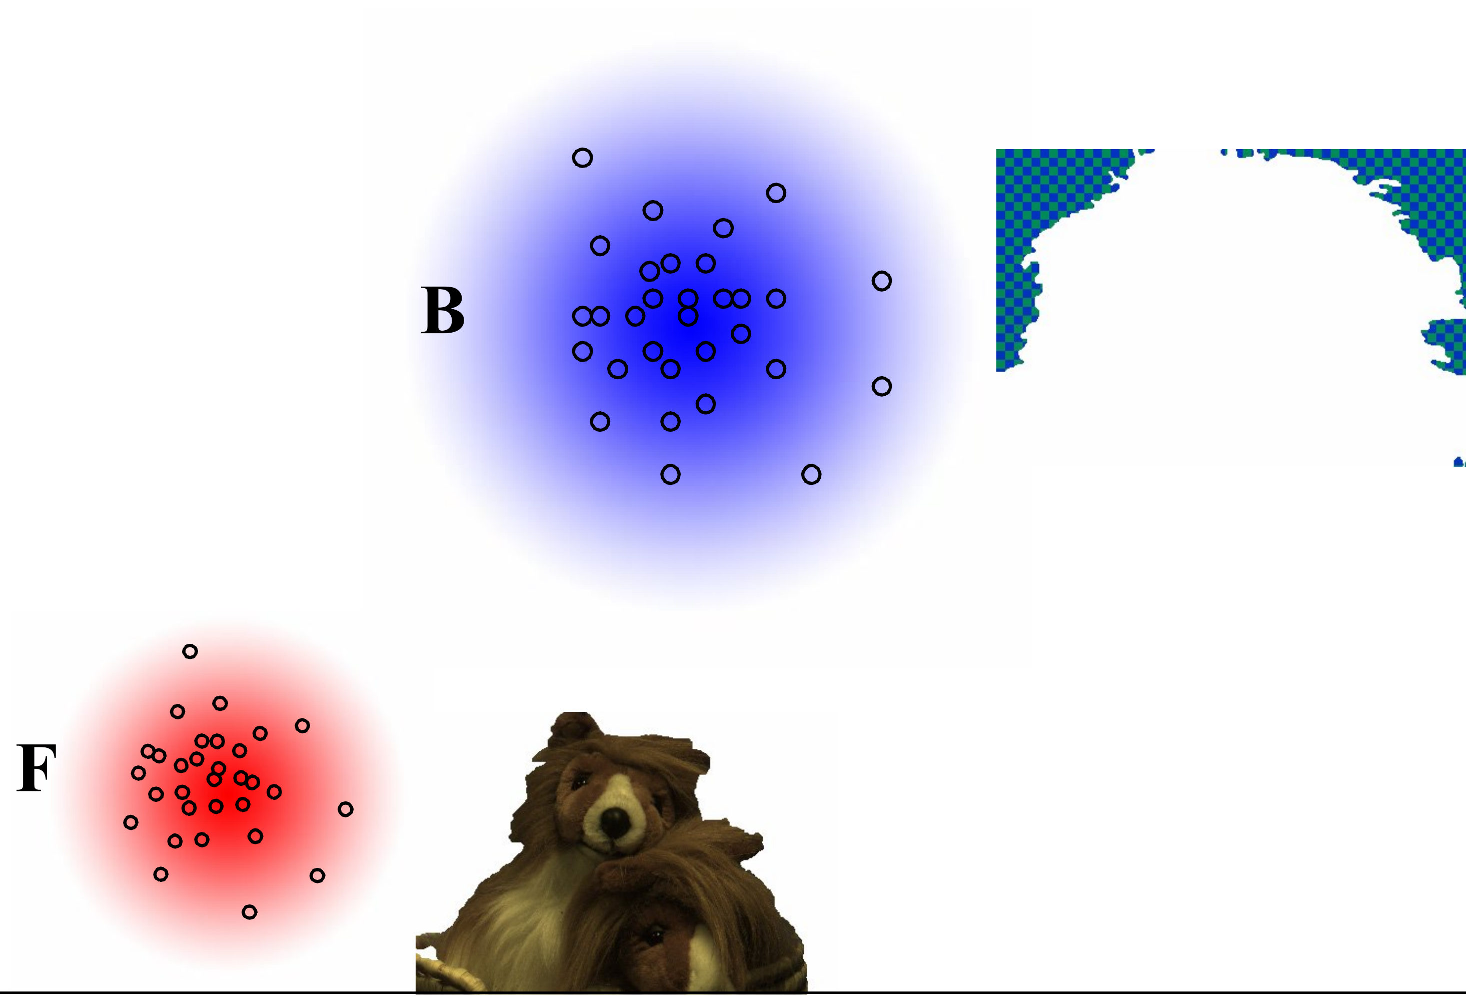

Prior probabilities L(F) & L(B)

- **Gaussians based on pixel color from known regions**
  - Can be anisotropic Gaussians
  - Compute the means $\overline{F}$ and $\overline{B}$ and covariance Σ_F, Σ_B

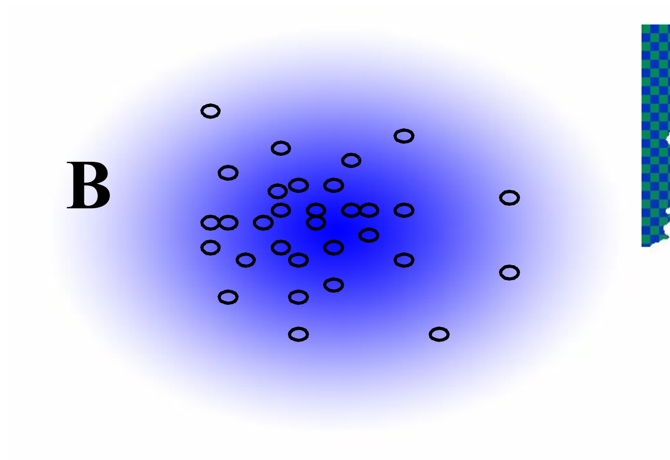

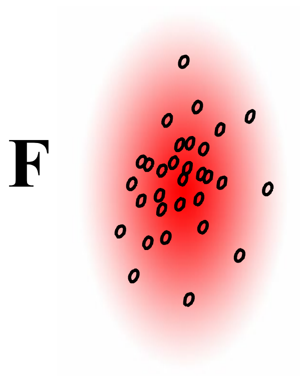

Prior probabilities L(F) & L(B)

- **Gaussians based on pixel color from known regions**

$$\bar{F} = \frac{1}{N_F} \sum F_i \qquad \Sigma_F = \frac{1}{N_F} \sum (F_i - \bar{F})(F_i - \bar{F})^T$$

$$L(F) = -(F - \bar{F})^T \Sigma_F^{-1} (F - \bar{F})/2$$

- **Same for B**

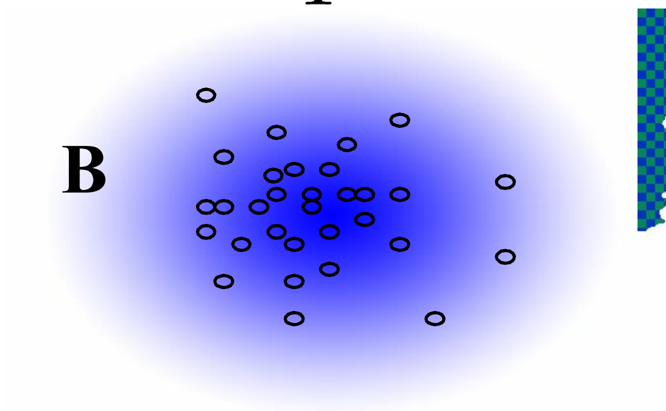

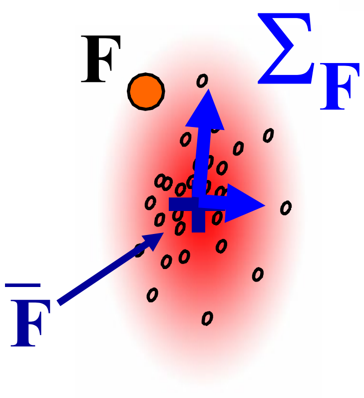

Prior probabilities L(α)

- **What about alpha?**
- **Well, we don't really know anything**
- **Keep L(α) constant and ignore it**
  - But … if we were labeling video frames, we could make prior predictions on value of α using temporal coherence (previous frames)

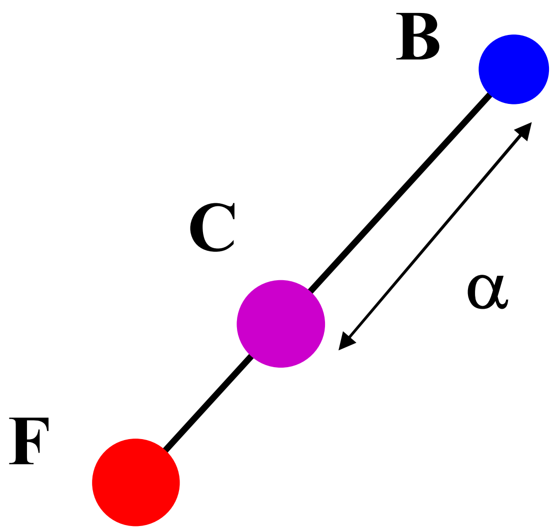

Recap: Bayesian matting equation

- **Maximize $L(C|F,B,\alpha) + L(F) + L(B) + L(\alpha)$**

$$L(C|F,B,\alpha) = - ||C - \alpha F - (1-\alpha) B||2 / \sigma^2_C$$

$$L(F) = -(F - \bar{F})^T \Sigma_F^{-1} (F - \bar{F})/2$$

$$L(B) = -(B - \bar{B})^T \Sigma_B^{-1} (B - \bar{B})/2$$

- **Unfortunately, not a quadratic equation because of the product $(1-\alpha)$ B**

  ➔ iteratively solve for F,B and for α

For α constant

- **Derivative of L(C|F,B,α) + L(F) +L(B)+L(α) wrt F & B, and set to zero gives**

$$\begin{bmatrix} \Sigma_F^{-1} + I\alpha^2/\sigma_C^2 & I\alpha(1-\alpha)/\sigma_C^2 \\ I\alpha(1-\alpha)/\sigma_C^2 & \Sigma_B^{-1} + I(1-\alpha)^2/\sigma_C^2 \end{bmatrix} \begin{bmatrix} F \\ B \end{bmatrix} = \begin{bmatrix} \Sigma_F^{-1}\overline{F} + C\alpha/\sigma_C^2 \\ \Sigma_B^{-1}\overline{B} + C(1-\alpha)/\sigma_C^2 \end{bmatrix},$$

For F & B constant

- **Derivative of L(C|F,B,α) + L(F) +L(B)+L(α) wrt α, and set to zero gives**

$$\alpha = \frac{(C-B)\cdot(F-B)}{\|F-B\|^2}$$

Recap: Bayesian matting

- **The user specifies a trimap**
- **Compute Gaussian distributions $\overline{F}$, Σ_F and $\overline{B}$, Σ_B for foreground and background regions**
- **Iterate**
  - Keep α constant, solve for F & B (for each pixel)
  - Keep F & B constant, solve for α (for each pixel)

Note that pixels are treated independently

Recap: Bayes cookbook

- **Express everything you know as probabilities**
  - Use Gaussians everywhere. Maybe multiple of them.
  - Learn from examples when you have them
  - Hack a noise model when you don't
  - Leave constant when desperate
  - More precisely, use Gaussian noise to express the likelihood to observe the input given any parameter in the solution space
    - Soft consistency constraint
- **Work in the log domain where everything is additive**
- **Find the maximum**

Additional Details

- **Use multiple Gaussians**
  - Cluster the pixels into multiple groups
  - Fit a Gaussian to each cluster
  - Solve for all the pairs of F & B Gaussians
  - Keep the highest likelihood
- **Use local Gaussians**
  - Not on the full image
- **Solve from outside-in**

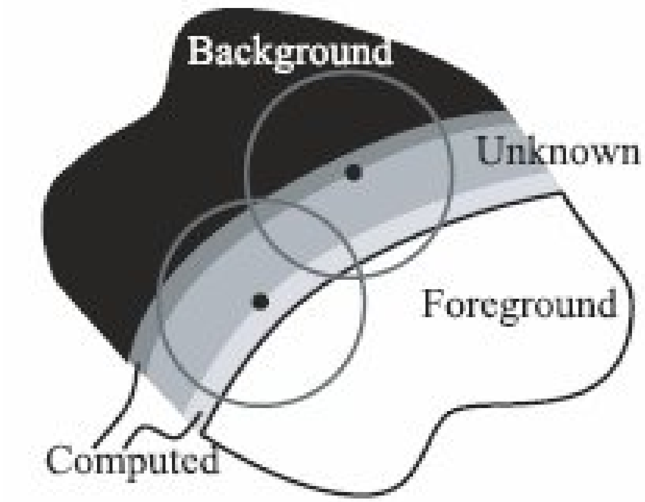

See Chuang et al.'s paper

Results

- **From Chuang et al. 2001**

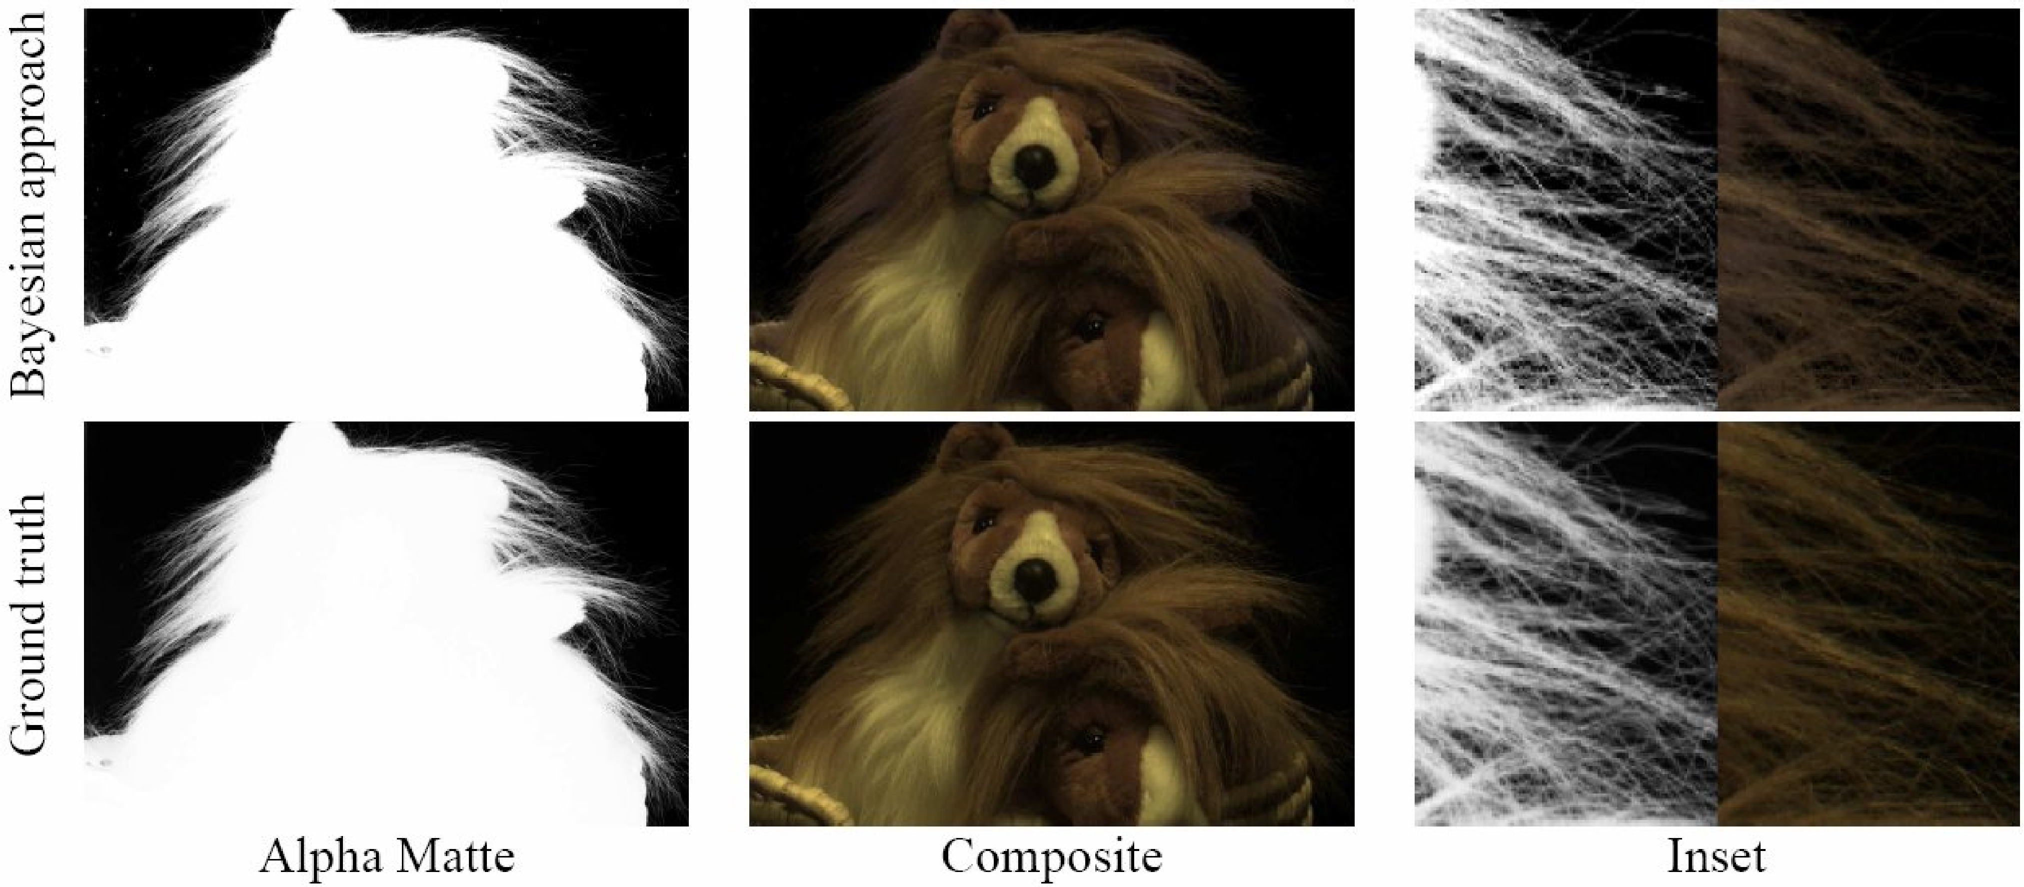

Bayesian approach

Alpha Matte

Composite

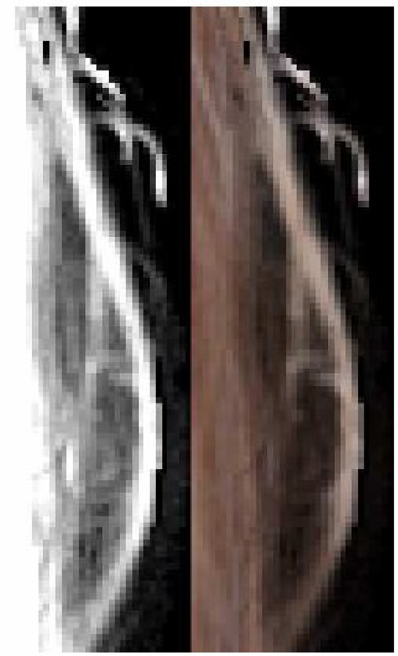

Inset

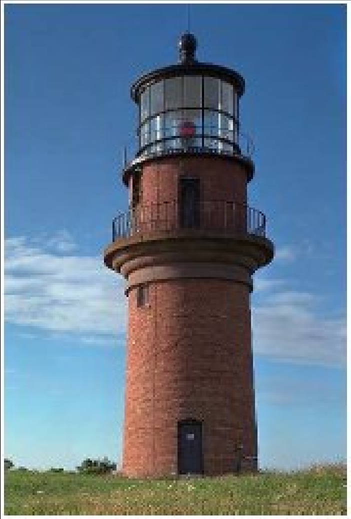

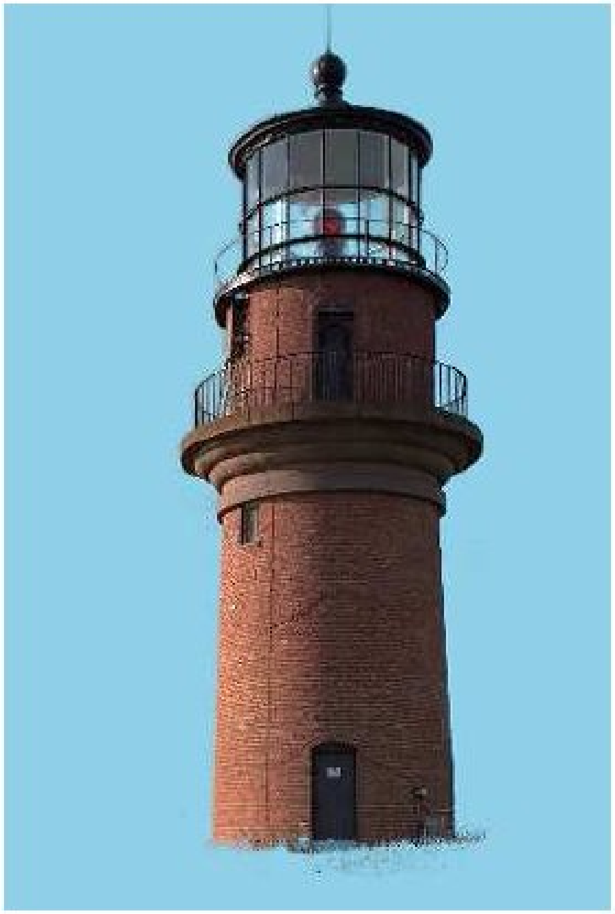

Alpha Matte Composite Inset

Extensions: Video

- Interpolate trimap between frames
- Exploit the fact that background might become visible
- http://grail.cs.washington.edu/projects/digital-matting/video-matting/

Video Matting of Complex Scenes

Yung-Yu Chuang[1] Aseem Agarwala[1] Brian Curless[1] David Salesin[1,2] Richard Szeliski[2]

[1]University of Washington [2]Microsoft Research

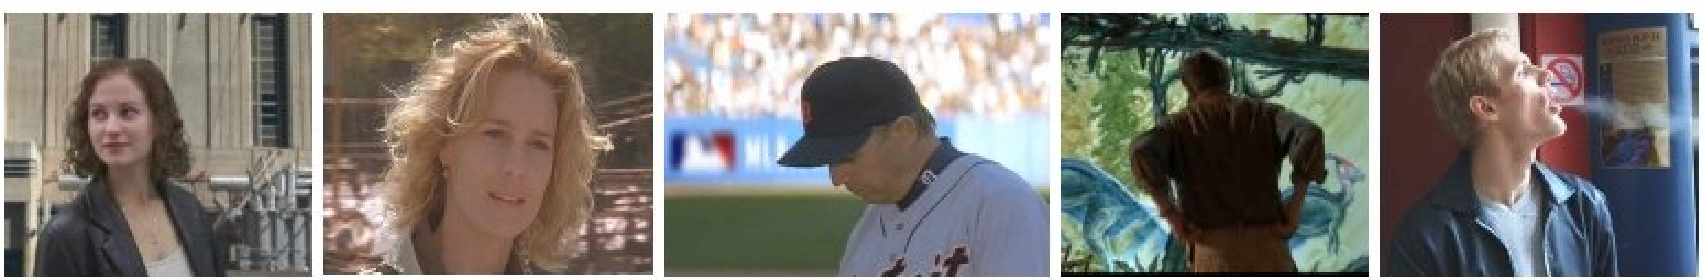

Questions?

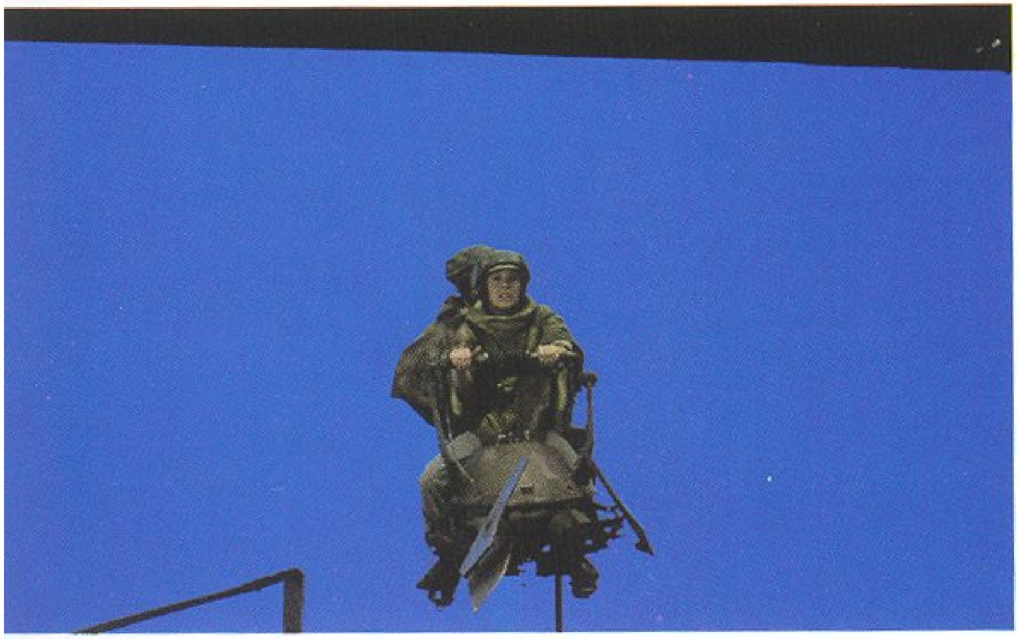

From Industrial Light & Magic, Smith

References

- **Smith & Blinn 1996**
  http://portal.acm.org/citation.cfm?id=237263
  Formal treatment of Blue screen
- **Ruzon & Tomasi 2000**
  http://ai.stanford.edu/~ruzon/alpha/
  The breakthrough that renewed the issue
  (but not crystal clear)
- **Chuang et al. 2001**
  http://research.microsoft.com/vision/visionbasedmodeling/publications/Chuang-CVPR01.pdf
- **Brinkman's Art & Science of Digital Compositing**
  – Not so technical , more for practitioners

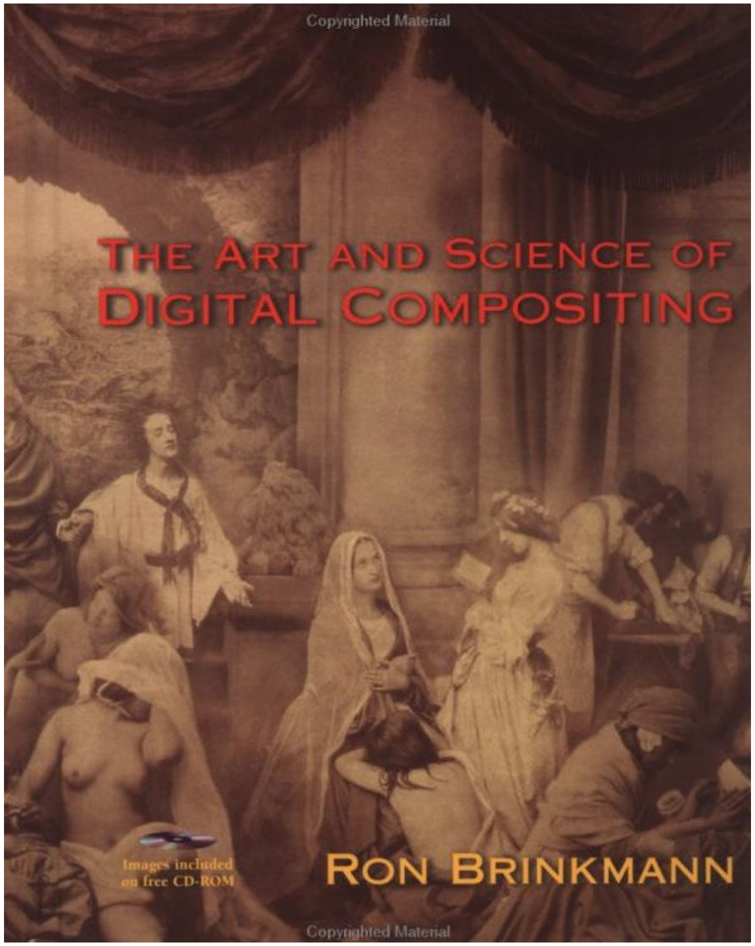

More Refs

Matting:

- http://graphics.cs.cmu.edu/courses/15-463/2004_fall/www/Lectures/matting.pdf
- http://www.csie.ntu.edu.tw/~cyy/publications/papers/Chuang2004Phd.pdf
- http://www.cse.ucsd.edu/classes/wi03/cse291-j/lec10-compositing.pdf
- http://graphics.stanford.edu/courses/cs248-99/comp/hanrahan-comp-excerpt.ppt

Chroma Key

- http://www.cs.utah.edu/~michael/chroma/

Blue screen:

- http://www.sut.ac.th/emdp/VisualEffect/The%20Blue%20Screen%20-%20Chroma%20Key%20Page.htm
- http://www.cs.princeton.edu/courses/archive/fall00/cs426/papers/smith95c.pdf
- http://www.seanet.com/Users/bradford/bluscrn.html
- http://en.wikipedia.org/wiki/Bluescreen
- http://www.neopics.com/bluescreen/
- http://entertainment.howstuffworks.com/blue-screen.htm
- http://www.vce.com/bluescreen.html
- http://www.pixelpainter.com/NAB/Blue_vs_Green_Screen_for_DV.pdf

Petro Vlahos (inventor of blue screen matting)

- http://theoscarsite.com/whoswho4/vlahos_p.htm
- http://en.wikipedia.org/wiki/Petro_Vlahos

To buy a screen:

http://shop.store.yahoo.com/cinemasupplies/chromkeyfab.html

Superman & blue screen:

- http://supermancinema.co.uk/superman1/the_production/the_crew/fx_bios/index.shtml
- http://home.utm.utoronto.ca/~kin/bluescreen.htm