6.098 Digital and Computational Photography
6.882 Advanced Computational Photography

Matting & Compositing

Bill Freeman
Frédo Durand
MIT - EECS

How does Superman fly?

- **Super-human powers?**
- **OR**
- **Image Matting and Compositing?**

Slide from Alyosha Efros

Motivation: compositing

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

Motivation: compositing

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

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

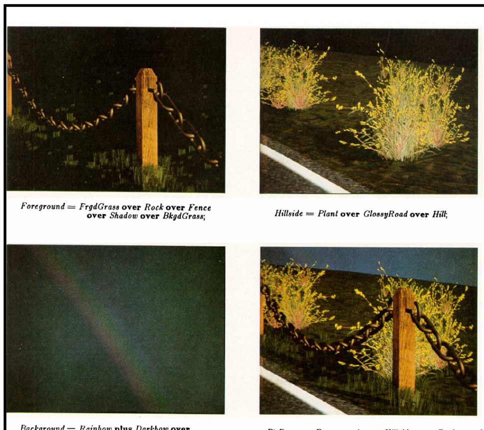

From Porter & Duff 1984

From Cinefex

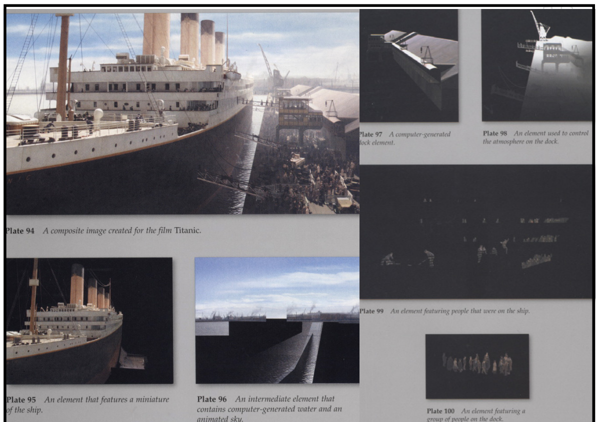

From the Art & Science of Digital Compositing

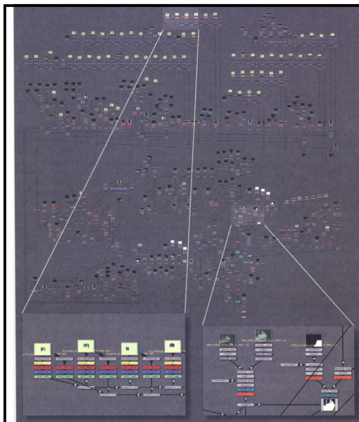

From the Art & Science of Digital Compositing

Slide from Alyosha Efros

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
  - See homework (pyramid splining)
  - See also in a couple weeks (gradient manipulation)
- **Extension to video**
  - Where life is always harder

Alpha

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

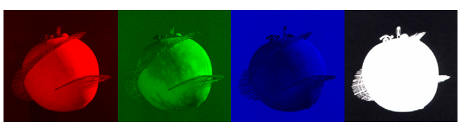

From the Art & Science of Digital Compositing

Why fractional alpha?

- **Motion blur, small features (hair) cause partial occlusion**

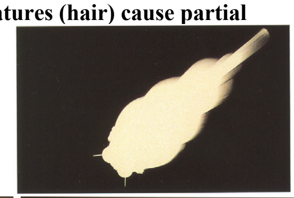

From Digital Domain

With binary alpha

From Digital Domain

With fractional alpha

From Digital Domain

Photoshop layer masks

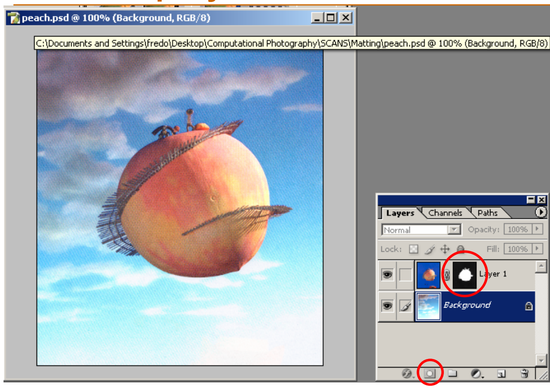

What does R, G, B, α represent?

– α: 1 means opaque, 0 means transparent
- **But what about R, G, and B?**

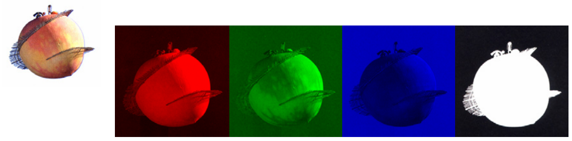

Two possible answers:

- ***Premultiplied***
  the color of the object is R/α, G/α, B/α
- **or not**
  the color of the object is R, G, B, and these values need to be multiplied by α for compositing

Pre-multiplied alpha

- **(R, G, B, α) means that the real object color is (R/α , G/α, B/α) and transparency is α.**
- **Motivated by supersampling for antialiasing in CG**

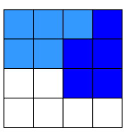

$\{R_i, G_i, B_i, \alpha_i\}$

supersampled pixel

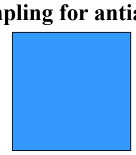

$1/n\ \Sigma\ R_i,$
$1/n\ \Sigma\ G_i,$
$1/n\ \Sigma\ B_i,$
$1/n\ \Sigma\ \alpha_i$

resampled (averaged value)

- **If I combine multiple subpixels, the same operations apply to the four channels**
  – In particular if I transform the image for scale/rotate

The compositing equation

Porter & Duff Siggraph 1984

- **Given Foreground F_A and Background F_B images**
- **For premultiplied alpha:**
  Output = $F_A+(1-\alpha_A)\ F_B$
- **For non-premultiplied:**
  Output = $\alpha\ F_A+(1-\alpha_A)\ F_B$

operation	quadruple	diagram	F_A	F_B
A over B	(0,A,B,A)	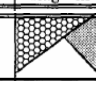	1	$1-\alpha_A$

Composing Two Elements

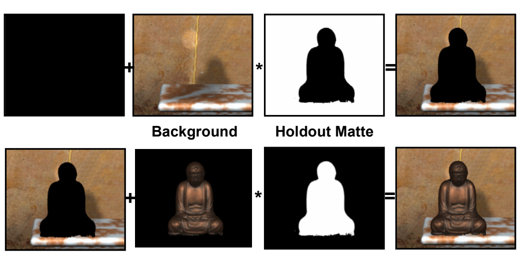

Slide from Pat Hanrahan

Optical Printing

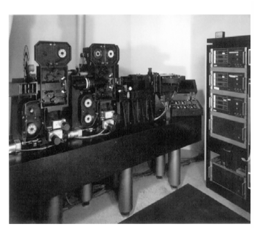

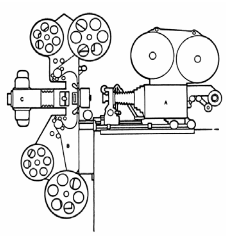

From: "Industrial Light and Magic," Thomas Smith (p. 181)

From: "Special Optical Effects," Zoran Perisic

Slide from Pat Hanrahan

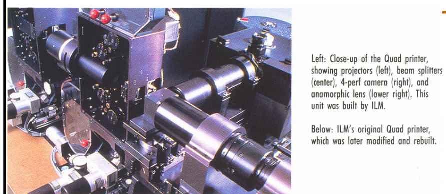

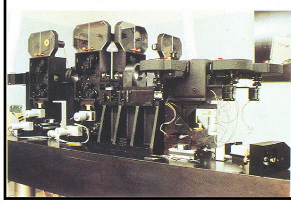

From: "Industrial Light and Magic," Thomas Smith

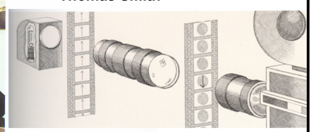

Limitations of alpha

- **Hard to represent stainglasses**
  – It focuses on subpixel occlusion (0 or 1)
- **Does not model more complex optical effects**
  – e.g. magnifying glass

Questions?

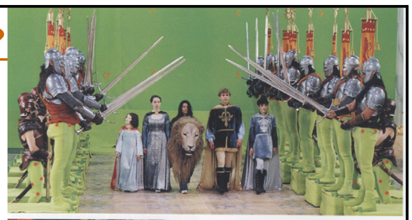

From Cinefex

Compositing

- ***Non* premultiplied version:**
  Given the foreground color $F=(R_F, G_F, B_F)$, the background color (R_B, G_B, B_B) and α for each pixel
- **The over operation is: $C=\alpha F+(1-\alpha)B$**
  – (in the premultiplied case, omit the first α)

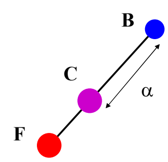

Matting problem

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

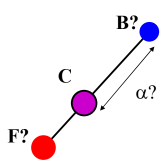

Matting ambiguity

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

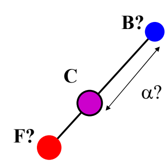

Matting ambiguity

- C=α F+(1-α)B
- 7 unknowns: α and triplets for F and B
- 3 equations, one per color channel

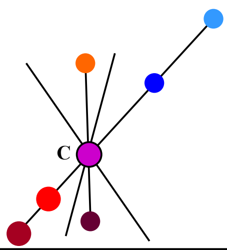

Matting ambiguity

- C=α F+(1-α)B
- 7 unknowns: α and triplets for F and B
- 3 equations, one per color channel
- With known background (e.g. blue/green screen): 4 unknowns, 3 equations

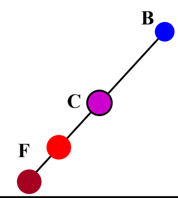

Questions?

From Cinefex

Traditional blue screen matting

- Invented by Petro Vlahos (Technical Academy Award 1995)
- Recently formalized by Smith & Blinn
- Initially for film, then video, then digital
- Assume that the foreground has no blue
- Note that computation of α has to be analog, needs to be simple enough

From Cinefex

Traditional blue screen matting

- Assume that blue *b* and green *g* channels of the foreground respect $b \leq a_2 g$ for a_2 typically between 0.5 and 1.5
- $\alpha = 1 - a_1(b - a_2 g)$
  - clamped to 0 and 1
  - a_1 and a_2 are user parameters
  - Note that $\alpha = 1$ where assumption holds

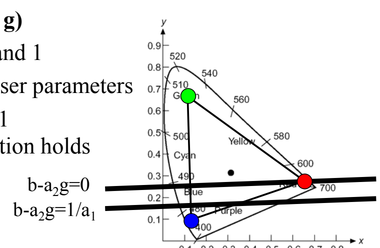

Traditional blue screen matting

- Assume that blue and green channels of the foreground respect $b \leq a_2 g$ for a_2 typically between 0.5 and 1.5
- $\alpha = 1 - a_1(b - a_2 g)$
  - clamped to 0 and 1
  - where a_1 and a_2 are user parameters
  - Note that $\alpha = 1$ where assumption holds
- Lots of refinements (see Smith & Blinn's paper)

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

Questions?

Hint: PSet 2 solution is in next slides

- **Hint 2: start problem set 2 early!**

Recall: Matting ambiguity

- $C=\alpha F+(1-\alpha)B$
- 7 unknowns: α and triplets for F and B
- 3 equations, one per color channel

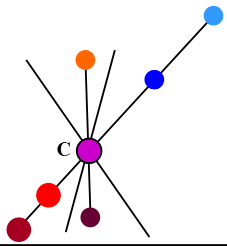

Natural matting

[Ruzon & Tomasi 2000, Chuang et al. 2001]

- **Given an input image with arbitrary background**
- **The user specifies a coarse *Trimap* (known Foreground, known background and unknown region)**
- **Goal: Estimate F, B, alpha in the unknown region**
  - We don't care about B, but it's a byproduct/unkown

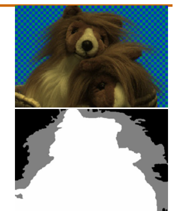

images from Chuang et al

Now, what tool do we know to estimate something, taking into account all sorts of known probabilities?

Who's afraid of Bayes?

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

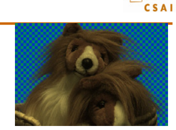

$$P(x|\mathbf{C}) = P(\mathbf{C}|x)\, P(x) \,/\, P(\mathbf{C})$$

The parameters you want to estimate — **Color** you observe — Likelihood function — Prior probability — Constant w.r.t. parameters x.

Matting and Bayes

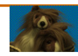

- What do we observe: Color C
- **What are we looking for?**

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

The parameters you want to estimate — Color you observe — Likelihood function — Prior probability — Constant w.r.t. parameters x.

Matting and Bayes

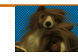

- What do we observe: Color C
- **What are we looking for: F, B, α**

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

Foreground, background, transparency you want to estimate — Color you observe — Likelihood function — Prior probability — Constant w.r.t. parameters x.

Matting and Bayes

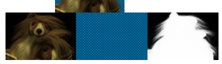

- What do we observe: Color C
- What are we looking for: F, B, α
- **Likelihood probability?**
  - Given F, B and Alpha, probability that we observe C

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

Foreground, background, transparency you want to estimate — Color you observe — **Likelihood function** — Prior probability — Constant w.r.t. parameters x.

Matting and Bayes

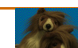

- What do we observe: Color C
- What are we looking for: F, B, α
- **Likelihood probability?**
  - Given F, B and Alpha, probability that we observe C
  - If measurements are perfect, non-zero only if $C=\alpha F+(1-\alpha)B$
  - But assume Gaussian noise with variance σ_C

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

Foreground, background, transparency you want to estimate — Color you observe — **Likelihood function** — Prior probability — Constant w.r.t. parameters x.

Matting and Bayes

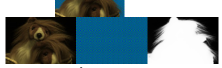

- What do we observe: Color C
- What are we looking for: F, B, α
- Likelihood probability: Compositing equation + Gaussian noise with variance σ_C
- **Prior probability:**
  - How likely is the foreground to have color F? the background to have color B? transparency to be α?

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

Foreground, background, transparency you want to estimate — Color you observe — Likelihood function — **Prior probability** — Constant w.r.t. parameters x.

Matting and Bayes

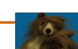

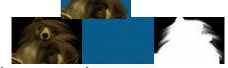

- What do we observe: Color C
- What are we looking for: F, B, α
- Likelihood probability: Compositing equation + Gaussian noise with variance σ_C
- **Prior probability:**
  Build a probability distribution from the known regions
  - ***This is the heart of Bayesian matting***

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

Foreground, background, transparency you want to estimate — Color you observe — Likelihood function — **Prior probability** — Constant w.r.t. parameters x.

Questions?

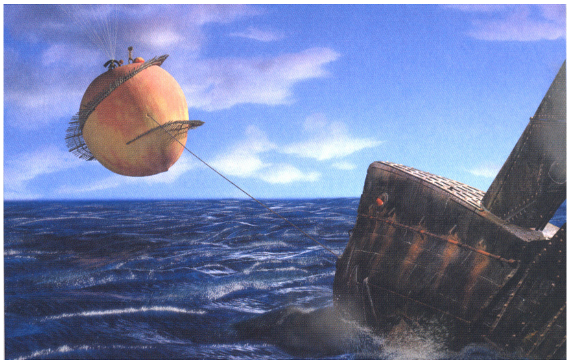

From the Art & Science of Digital Compositing

Let's derive

- **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,α)

- **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)**

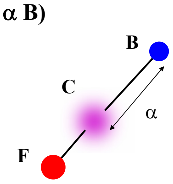

Prior probabilities L(F) & L(B)

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

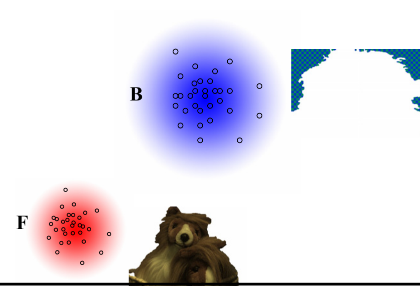

Prior probabilities L(F) & L(B)

- **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

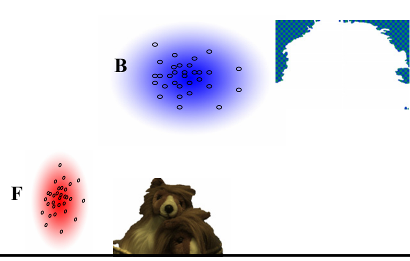

Prior probabilities L(F) & L(B)

- **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**

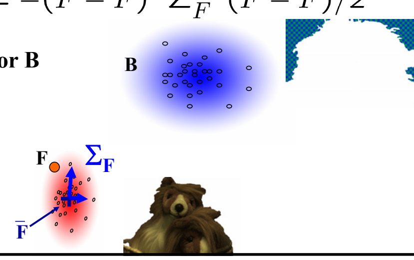

Prior probabilities L(α)

- **What about alpha?**
- **Well, we don't really know anything**
- **Keep L(α) constant and ignore it**
  - But see coherence matting for a prior on α

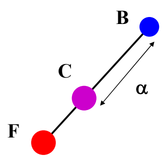

Questions?

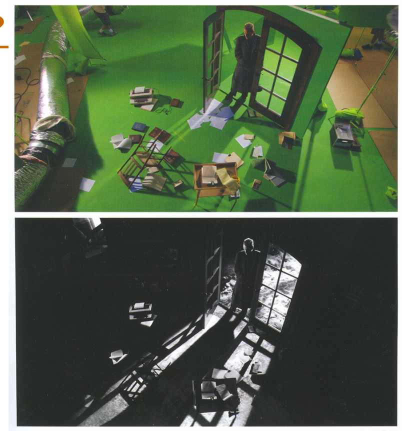

Recap: Bayesian matting equation

- **Maximize L(C|F,B,α) + L(F) +L(B)+L(α)**

L(C|F,B,α)= - ||C - α F – (1-α) B||2 / σ^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-α) B**
  ➔ iteratively solve for F,B and for α

For α constant

- **Derive 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

- **Derive 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

Questions?

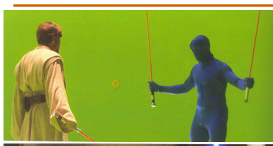

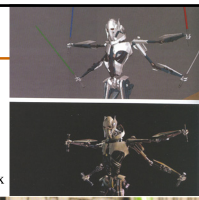

From Cinefex

Additional gimmicks (not on p-set!)

- **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**

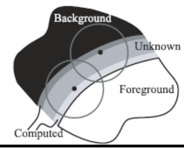

See Chuang et al.'s paper

Results

- **From Chuang et al. 2001**

Bayesian approach

Ground truth

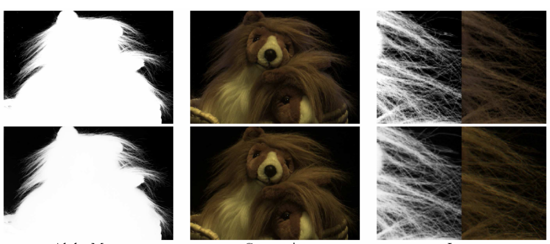

Alpha Matte | Composite | Inset

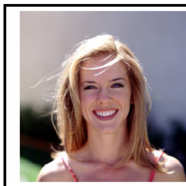

Bayesian approach

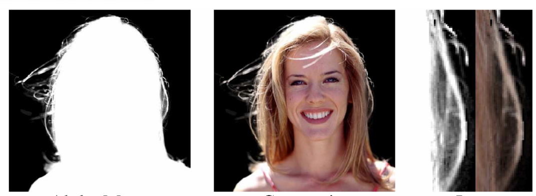

Alpha Matte | Composite | Inset

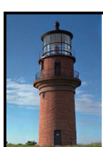

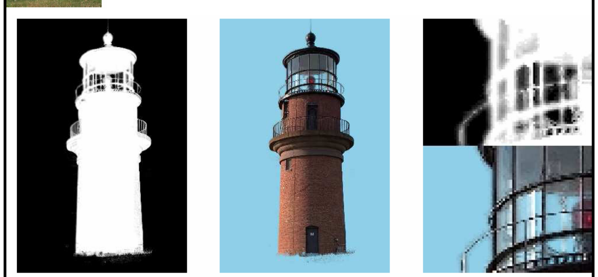

Alpha Matte | Composite | Inset

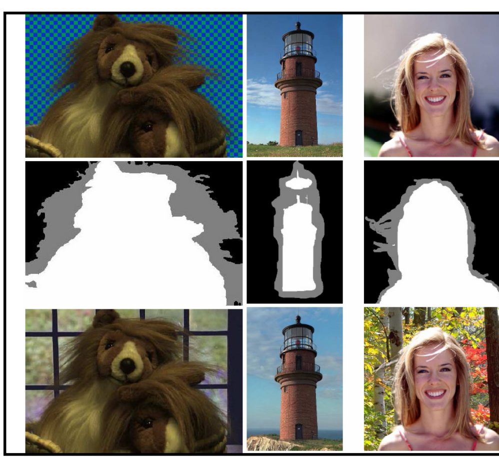

From Chuang et al 2001

Questions?

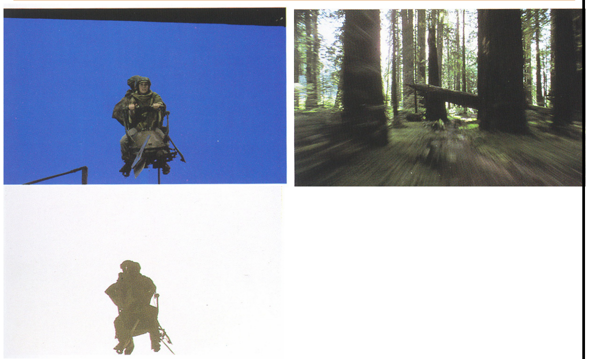

From Industrial Light & Magic, Smith

Extensions: Video

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

Environment matting

Model complex optical effects

Each pixel can depend on many background pixels

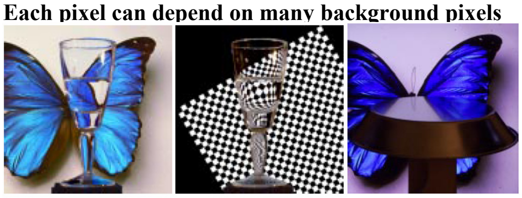

Figure 1 A water goblet, digitally composited onto background images, preserving the effects of refraction.

http://grail.cs.washington.edu/projects/envmatte/

Environment Matting and Compositing
Zongker, Werner, Curless, and Salesin. *SIGGRAPH 99, August 1999.*

Environment Matting Extensions: Towards Higher Accuracy and Real-Time Capture
Chuang, Zongker, Hindorff, Curless, Salesin, and Szeliski. *SIGGRAPH 2000.*

Environment matting

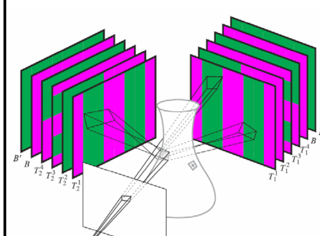

Figure 2 The environment matting process uses structured textures to capture how light is reflected and refracted from a backdrop (right shaft), as well as from various sidedrops (left shaft). The process also captures light coming from the backdrop that is seen through uncovered portions of a pixel (center shaft).

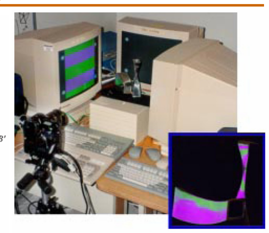

Figure 3 A photograph of the experimental setup used to capture environment mattes. The camera in the foreground photographs objects surrounded by structured light patterns displayed on the computer monitors. The image is processed to extract only the area covered by the backdrop pattern, resulting in an image as shown in the inset.

From Zongker et al.

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

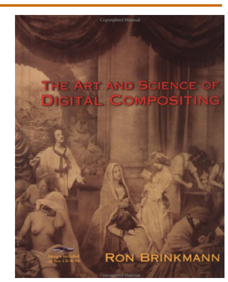

More Refs

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

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 a 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**