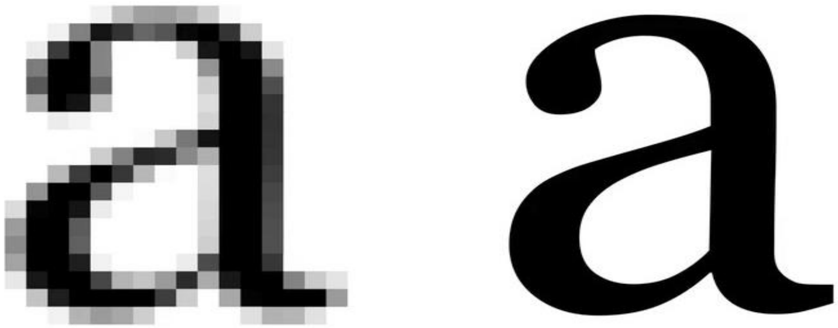

Image Processing & Antialiasing

Part I (Overview and Examples)

Bin Sheng

Image Processing

- IP is fundamental to both computer graphics and computer vision
- Has its own publications and conferences
  - IEEE Transactions on Image Processing (TIP)
  - Image and Vision Computing
  - Journal of Electronic Imaging
  - IEEE International Conference on Image Processing (ICIP)
  - IEEE International Conference on Computational Photography (ICCP)
- Once was closer to **signal theory** and **audio processing** than to graphics
- Image Synthesis in CG
  - model -> image
- Image Processing
  - image ->
    - image
    - measurements
    - model
    - recognition
    - understanding
    - ...
- DSPs and GPUs used in both CG and IP

Outline

- **Overview**
- Example Applications
- Jaggies & Aliasing
- Sampling & Duals
- Convolution
- Filtering
- Scaling
- Reconstruction
- Scaling, continued
- Implementation

What does "image" mean for us?

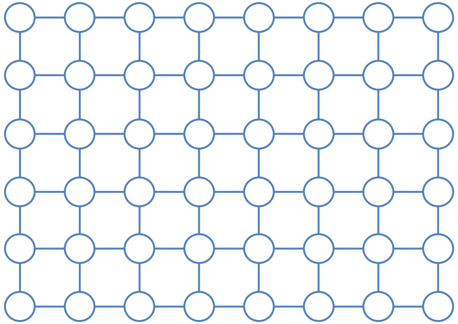

- A 2D domain with samples at regular points (almost always a rectilinear grid)
  - Can have multiple values sampled per point
  - Meaning of samples depend on the application (red, green, blue, opacity, depth, etc.)
- Units also depend on the application
  - e.g., a computed int or float to be mapped to voltage needed for display of a pixel on a screen
  - e.g., as a physical measurement of incoming light (e.g., a camera pixel sensor)
- Introduction to sampling demo

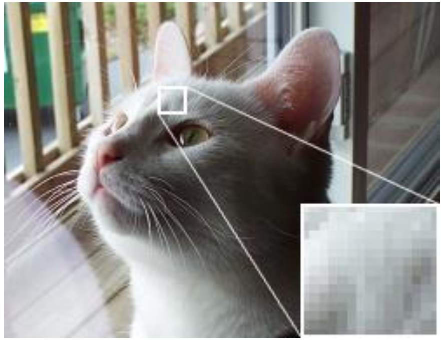

What is a channel?

- A channel is a collection (e.g., array) of all the samples of a particular type
- RGB is a common format for image channels
  - Easy to implement in h/w
  - Corresponds approximately to human visual system anatomy (specialized "R, G, and B" cones)
  - Samples represent the intensity of the light at a point for a given wavelength (red, green, or blue)
- The R channel of an image is an image containing just the red samples

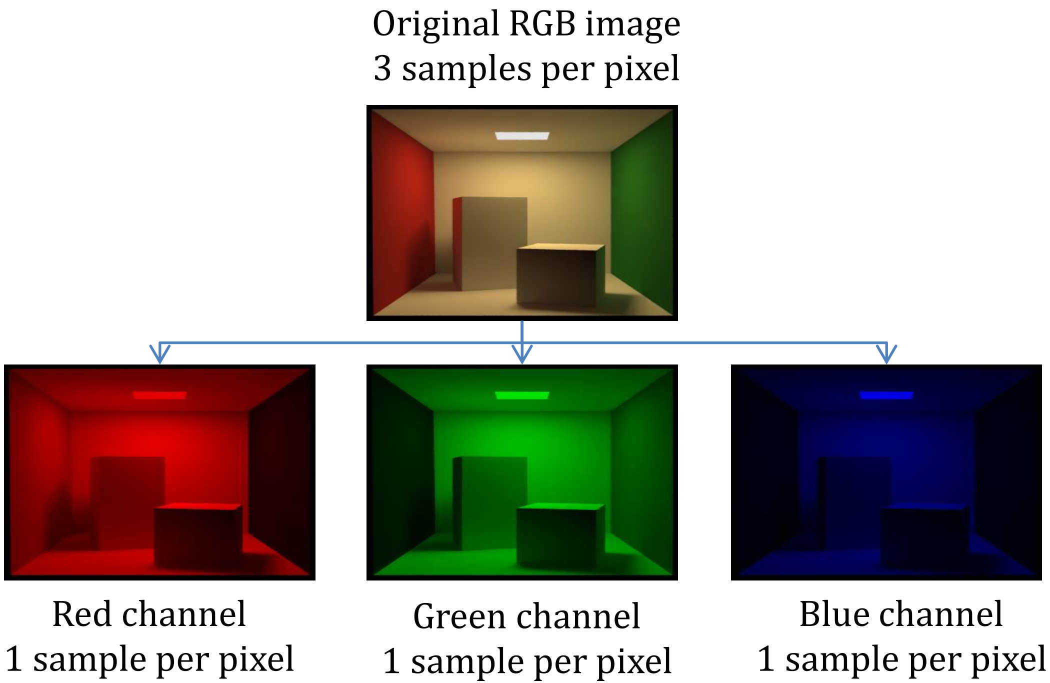

The alpha channel

- In addition to the R, G, and B channels of an image, add a fourth channel called α (transparency/opacity/translucency)
- Alpha varies between 0 and 1
  - Value of 1 represents a completely opaque pixel, one you cannot see through
  - Value of 0 is a completely transparent pixel
  - Value between $0 < \alpha < 1$ determines translucency
- Useful for blending images
  - Images with higher alpha values are less transparent
  - Linear interpolation ($\alpha X + (1 - \alpha)Y$) or full [Porter-Duff compositing algebra]

The orange box is drawn on top of the purple box using $\alpha = 0.8$

Modeling an image

- Model a one-channel m × n image as the function $u(i, j)$
  - Maps pairs of integers (pixel coordinates) to real numbers
  - i and j are integers such that $0 \leq i < \mathrm{m}$ and $0 \leq j < \mathrm{n}$
- Associate each pixel value $u(i, j)$ to small area around display location with coordinates (i, j)
- A pixel here looks like a square centered over the sample point, but it's just a scalar value and the actual geometry of its screen appearance varies by device
  - Roughly circular spot on CRT (Cathode Ray Tube)
  - Rectangular on LCD panel

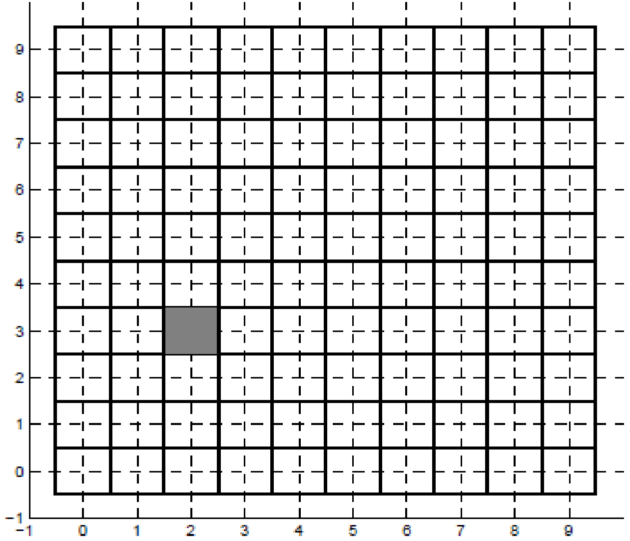

Pixels

- Pixels are point samples, not "squares" or "dots"
- Point samples reconstructed for display (often using multiple subpixels for primary colors)

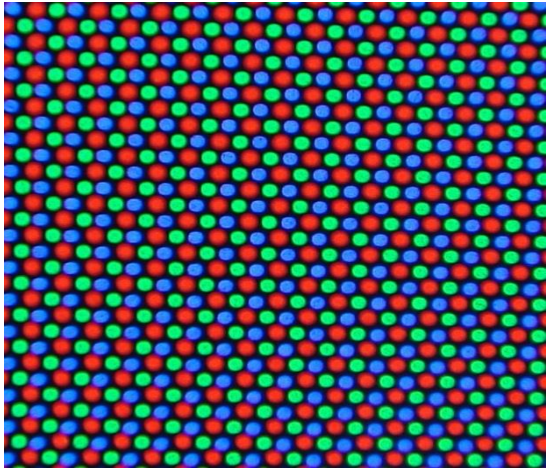

Close-up of a CRT screen

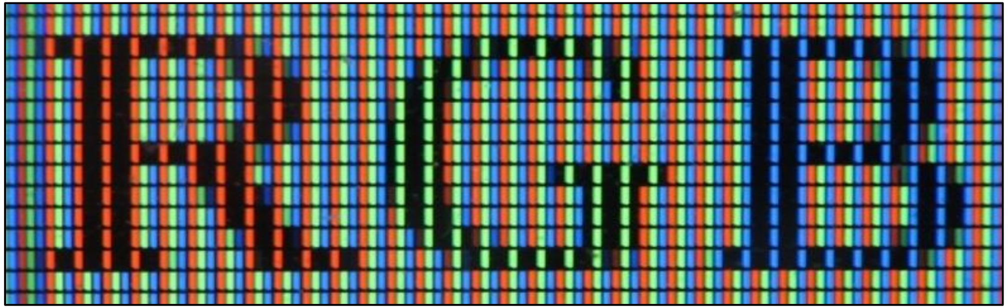

Close-up of an LCD screen

Discrete Images vs. Continuous Images

- Two kinds of images
  - Discrete
  - Continuous
- Discrete image
  - Function from $\mathbb{Z}^2$ to $\mathbb{R}$
  - How images are stored in memory
  - The kind of images we generally deal with as computer scientists

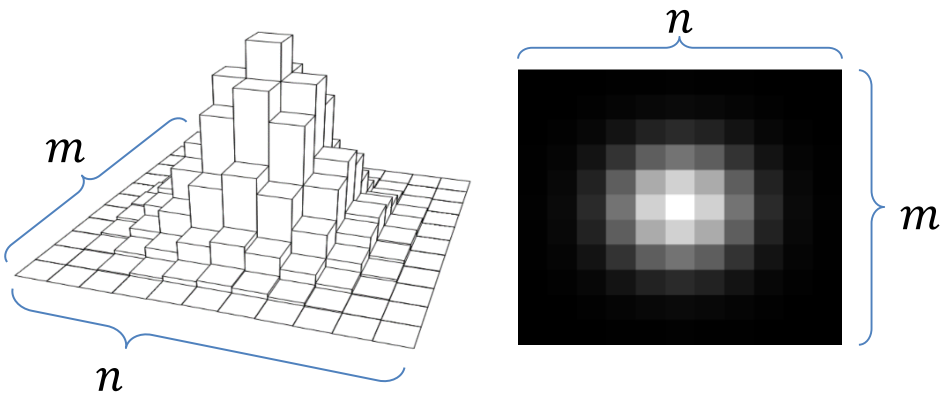

Discrete image $u(i, j)$

Discrete Images vs. Continuous Images

- Continuous image
  - Function from $\mathbb{R}^2$ to $\mathbb{R}$
  - Images in the real world
  - "Continuous" refers to the domain, not the values (discontinuities could still exist)
- Example: Gaussian distribution
  - i_0 and j_0 are the center of the Gaussian
  - $u: \mathbb{Z}^2 \rightarrow \mathbb{R}, u(i,j) = e^{-(i-i_0)^2-(j-j_0)^2}$
  - $v: \mathbb{R}^2 \rightarrow \mathbb{R}, v(i,j) = e^{-(i-i_0)^2-(j-j_0)^2}$
  - $i_0 = (n-1)/2$ and $j_0 = (k-1)/2$ (n odd)
  - Here $n = 11$ and $m = 11$

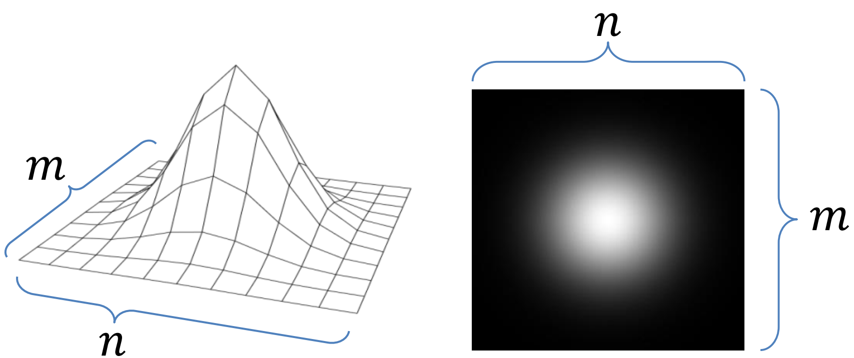

Continuous image $v(i,j)$

Idealized Five Stage Pipeline of Image Processing

- The stages are
  - Image acquisition – how we obtain images in the first place
  - Preprocessing – any effects applied before mapping (e.g., crop, mask, filter)
  - Mapping – catch-all stage involving image transformations or image composition
  - Post processing – any effects applied after mapping (e.g., texturizing, color remapping)
  - Output – printing or displaying on a screen
- In practice, stages may be skipped
- Middle stages are often interlaced

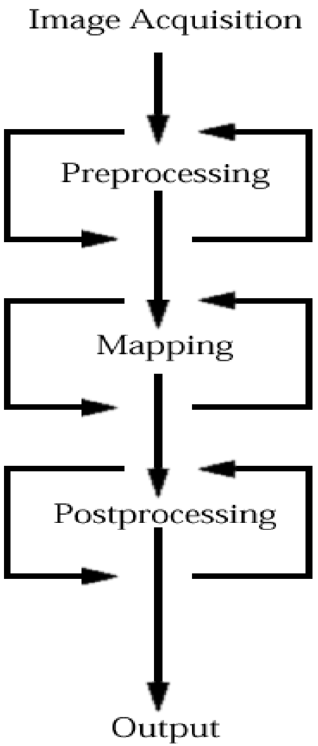

Stage 1: Image Acquisition

- Image Synthesis
  - Images created by a computer
  - Painted in 2D
    - Corel Painter (website)
    - Photoshop (website)
  - Rendered from 3D geometry
    - Pixar's RenderMan (website)
    - Autodesk's Maya (website)
    - Your CS123 projects
  - Procedurally textured
    - Generated images intended to mimic their natural counterparts
    - e.g., procedural wood grain, marble
- Image Capture
  - Images from the "real world"
  - Information must be digitized from an analog signal
  - Common capture methods:
    - Digital camera
    - Satellite data from sensors (optical, thermal, radiation,...)
    - Drum scanner
    - Flatbed photo scanner
    - Frames from video

Stage 2: Preprocessing

- Each source image is adjusted to fit a given tone, size, shape, etc., to match a desired quality or to match other images
- Can make a set of dissimilar images appear similar (if they are to be composited later), or make similar parts of an image appear dissimilar (such as contrast enhancement)

Original

Adjusted grayscale curve

Stage 2: Preprocessing (continued)

- Preprocessing techniques include:
  - Adjusting color or grayscale curve

  - Cropping

  - Masking (cutting out part of an image)

  - Blurring and sharpening

  - Edge detection/enhancement

  - Filtering and antialiasing

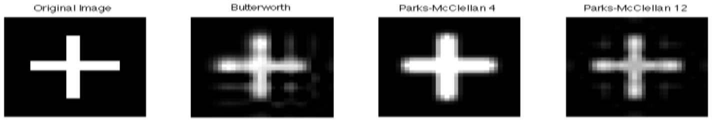

  - Scaling up (super sampling) or scaling down (sub sampling)

Stage 2: Preprocessing (continued)

- Notes:
  - Blurring, sharpening, and edge detection can also be postprocessing techniques
  - Some preprocessing algorithms are not followed by mapping, others that involve resampling the image may be interlaced with mapping: filtering is done this way

Stage 3: Mapping

- Mapping is a catch-all stage where several images are combined, or geometric transformations are applied
- Transformations include:
  - Rotating
  - Scaling
  - Shearing
  - Warping
  - Feature-based morphing
- Compositing:
  - Basic image overlay
  - Smooth blending with alpha channels
  - Poisson image blending
    - Seamlessly transfers "details" (like edges) from part of one image to another

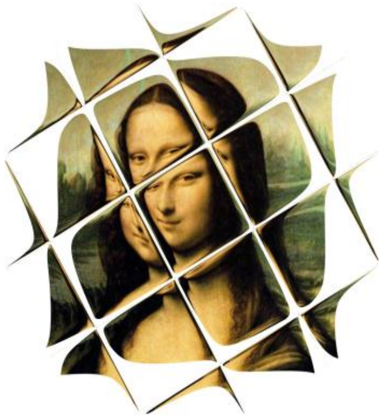

Image Warping

Poisson Image Blending

Image credit: © Evan Wallace 2010

Stage 4: Postprocessing

- Creates global effects across an entire image or selected area
- Art effects
  - Posterizing
  - Faked "aging" of an image
  - Faked "out-of-focus"
  - "Impressionist" pixel remapping
  - Texturizing
- Technical effects
  - Color remapping for contrast enhancement
  - Color to B&W conversion
  - Color separation for printing (RGB to CMYK)
  - Scan retouching and color/contrast balancing

Posterizing

Aging

Impressionist

Stage 5: Output (Archive/Display)

- Choice of display/archive method may affect earlier processing stages
  - Color printing accentuates certain colors more than others
  - Colors on the monitor have different **gamuts** and **HSV** values than the colors printed out
    - Need a mapping
  - HSV = hue, saturation, value, a cylindrical coordinate system for the RGB color model
  - Gamut = set of colors that can be represented by output device/printer

- Display Technologies
  - Monitor (CRT → LCD/LED/OLED/Plasma panel)

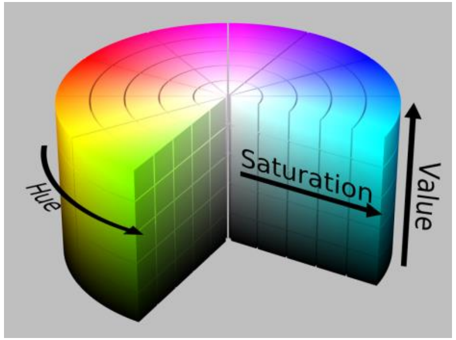

An HSV cylinder

D

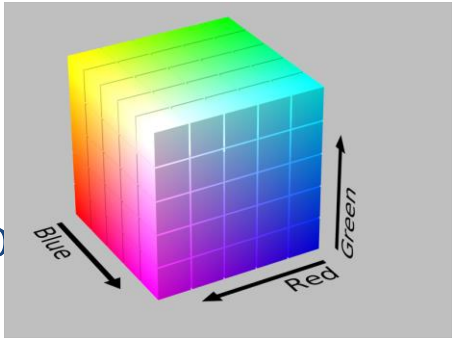

An RGB cube

Outline

- Overview
- **Example Applications**
- Jaggies & Aliasing
- Sampling & Duals
- Convolution
- Filtering
- Scaling
- Reconstruction
- Scaling, continued
- Implementation

Example 1: Edge Detection Filtering

- Edge detection filters measure the difference between adjacent pixels
- A greater difference means a stronger edge
- A threshold is sometimes used to remove weak edges

***Sobel** edge detection filter*

Example 1: Edge Detection Filtering (Continued)

- Used with MRI scans to reveal boundaries between different types of tissues
- MRI scan is image where gray level represents tissue density
- Used same filter as previous slide

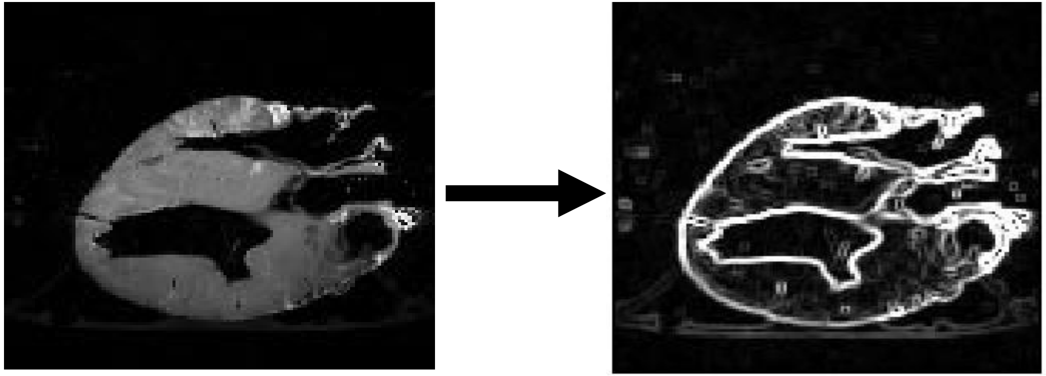

Original MRI image of a dog heart

Image after edge detection

Example 2: Image Enhancement for Forensics

- Extract evidence from seemingly incomprehensible images
- Normally, image enhancement uses many filtering steps, and often no mapping at all
- Former Prof. Michael Black and his class, CS296-4, received a commendation for helping Virginia police in a homicide case

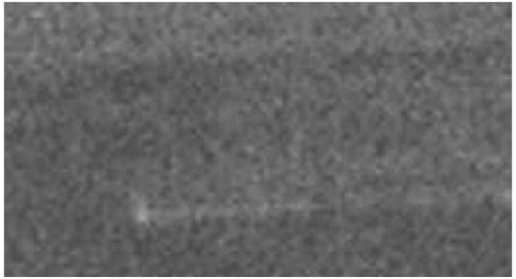

Before enhancement

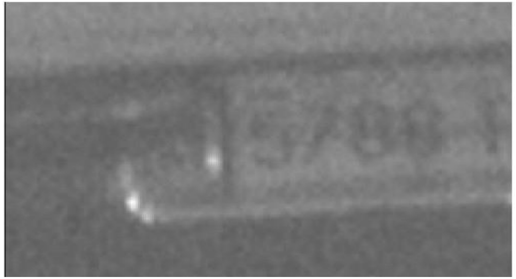

After enhancement

Example 2: Image Enhancement for Forensics

- We have a security camera video of the back of a car that was used in a robbery
- The image is too dark and noisy for the police to pull a license number
- Though humans can often discern an image of poor quality, filtering can make it easier for a pattern-recognition algorithm to decipher embedded symbols
  - Optical Character Recognition
- Step 1: Get the frame from the videotape digitized with a frame-grabber

Example 2: Image Enhancement for Forensics

- Step 2: Crop out stuff that appears to be uninteresting (outside plate edges)
- This step can speed process by doing image processing steps on fewer pixels
- Can't always be done, may not be able to tell which sections are interesting without some processing

Example 2: Image Enhancement for Forensics

- Step 3: Use edge-sharpening filter to add contrast to plate number
- This step enhances edges by raising discontinuities at brightness gaps in image

Example 2: Image Enhancement for Forensics

- Step 4: Remap colors to enhance contrast between numbers and plate itself
- Now, can make a printout for records, or just copy plate number down: YNN-707!
- Note that final colors do not even resemble real colors of license plate—enhancement techniques have seriously distorted the colors!

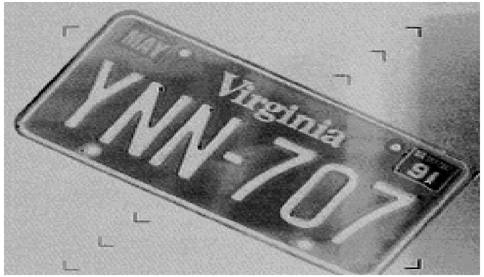

Multipart Composition

- Image composition is popular in the art world, as well as in tabloid news
- Takes parts of several images and creates single image
  - Hard part is making all images fit together naturally
- Artists can use it to create amazing collages and multi-layered effects
- Tabloid newspaper artists can use it to create "News Photos" of things that never happened – "Fauxtography".
  - There is no visual truth in media!

Famous Faked Photos

Chinese press photo of Tibet railway

Tom Hanks and JFK

Example image composition (1/5)

- Lars Bishop, former CS123 Head TA, created a news photo of himself “meeting” with former Russian President Boris Yeltsin
  - post-Gorbachev and Perestroika. He served 10 July 1991 – 31 December 1999, resigned in favor of Putin)
- Needless to say, Lars Bishop never met Mr. Yeltsin
- Had to get the images, cut out the parts he wanted, touch them up, paste them together, and retouch the end result

Image of Boris (from Internet)

Image of Lars (from video camera)

Example image composition (2/5)

- Cut the pictures we want out of the original images
  - Paint a region around important parts of images (outline of people) using Photoshop
  - Continue touching up this outline until no background at edge of people
  - Use a smart lasso tool that grows until it hit the white background, thus selecting subject. (“Magic Wand” tool in photoshop can accomplish this)

Example image composition (3/5)

- Filter the images to make them appear similar, and paste them together
  - Boris is blurred and brightened to get rid of the halftoning lines (must have been a magazine photo)
  - Lars is blurred and noise is added to match image quality to that of Boris
  - Images are resized so Boris and Lars are at similar scales

Example image composition (4/5)

- Finalize image
  - Created a simple, two-color background and added noise so it fit with the rest of the image, placed cutout of the two subjects on top of background
  - This left a thin white halo around the subjects, so used a “Rubber Stamp” tool to stamp background noise patterns over halo, making seams appear less obvious

Example image composition (5/5)

- Final Image (with retouching at edges)

BISHOP AND YELTSIN TALK PEACE

BISHOP: “I couldn’t understand a single word he said!”

Image Composition – Frankenface

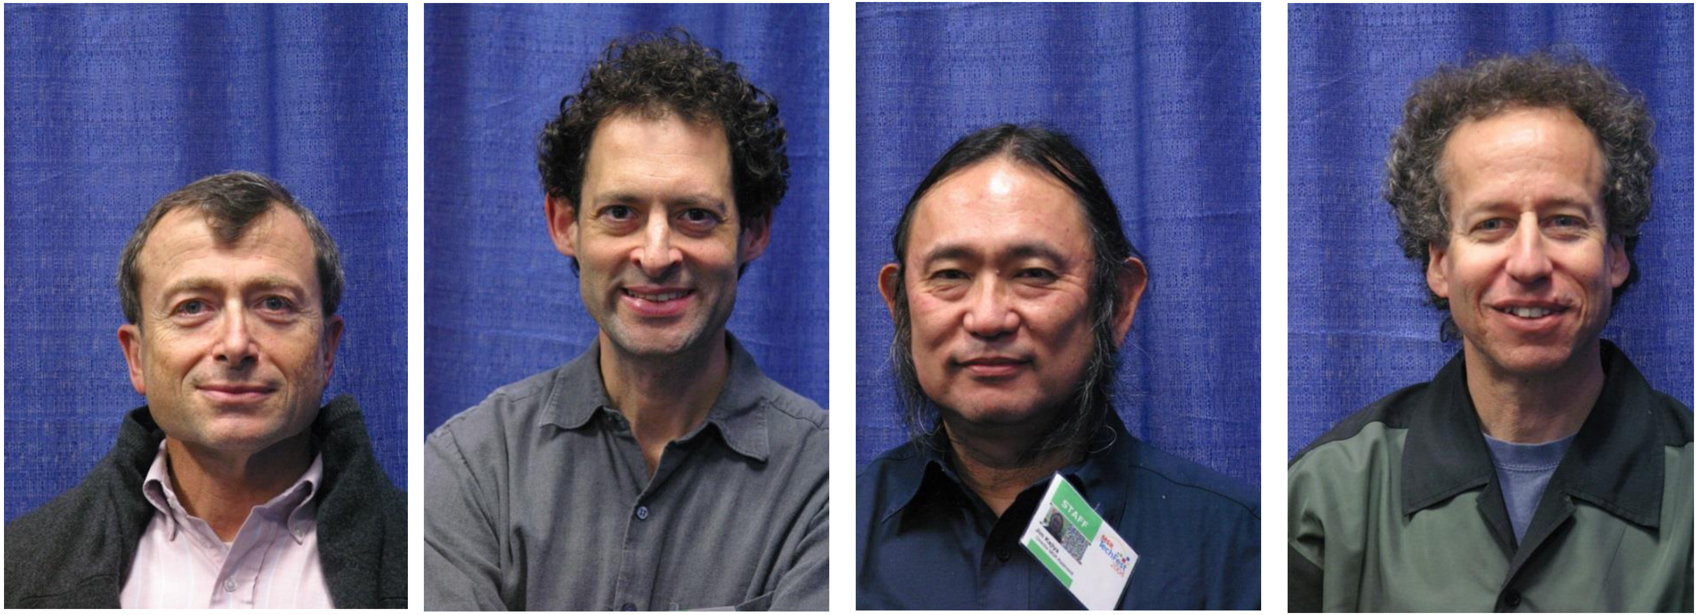

Aseem Agarwala, Mira Dontcheva, Maneesh Agrawala, Steven Drucker, Alex Colburn, Brian Curless, David Salesin, Michael Cohen. **Interactive Digital Photomontage**. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2004)*, 2004. http://grail.cs.washington.edu/projects/photomontage/

Image Composition - Frankenface

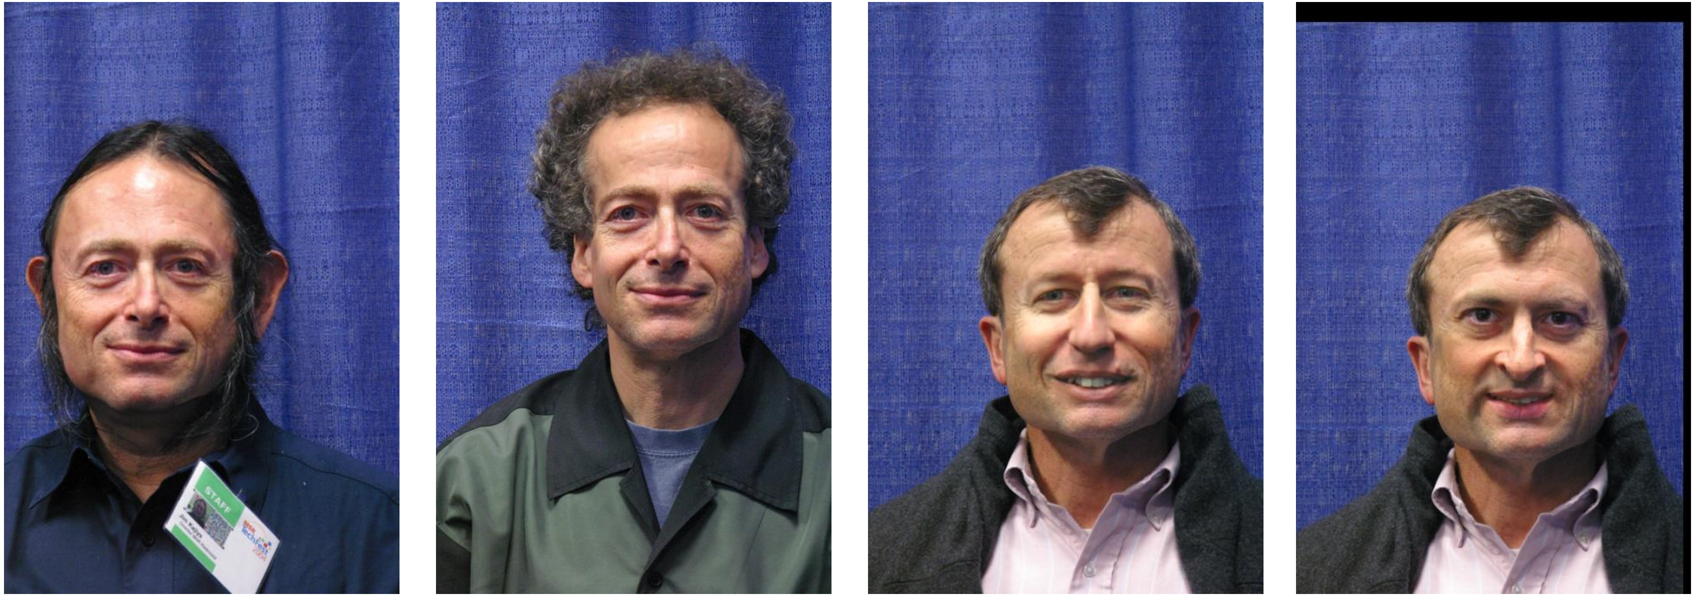

Aseem Agarwala, Mira Dontcheva, Maneesh Agrawala, Steven Drucker, Alex Colburn, Brian Curless, David Salesin, Michael Cohen. **Interactive Digital Photomontage**. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2004)*, 2004. http://grail.cs.washington.edu/projects/photomontage/

3D Image Processing

- 3D images
  - 3D image volumes from MRI scans need image processing
  - 2D image processing techniques often have 3D analogs
  - Display becomes more difficult: **voxels** replace pixels (**volumetric rendering)**
  - Increases time and space complexity:
    - 4 channel 1024x1024 image = 4 megs
    - 4 channel 1024x1024x1024 image = 4 **gigs**!
    - N^2 processing algorithms become N^3

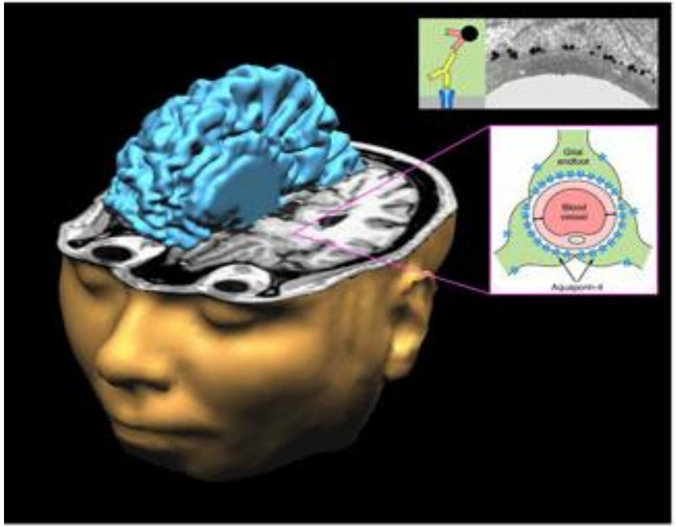

Illustration: Erlend Nagelhus and Gunnar Lothe.
3D MRI: Kyrre Eeg Emblem, Rikshospitalet, and Inge Rasmussen, Nidelven Hjerneforskningslaboratorium.

University of Oslo, 1999

Computer Vision (1/2)

- Computer graphics is the business of using models to create images; computer vision solves the opposite problem—deriving models from images
- Computer must do all the processing without human intervention
- Often, processing techniques must be fast
  - Slow processing will add to camera-to-reaction latency (lag) in system
- Common preprocessing techniques for computer vision:
  - Edge enhancement
  - Region detection
  - Contrast enhancement
  - Feature point detection

Computer Vision (2/2)

- Image processing makes information easier to find
- Pattern detection and pattern recognition are separate fields in their own right
  - Pattern detection: looking for features and describing the image's content at a higher level
  - Pattern recognition: classifying collections of features and matching them against library of stored patterns. (e.g., alphanumeric characters, types of abnormal cells, or human features in the case of biometrics)
  - Pattern detection is one important component of pattern recognition
- Computer vision can be used to recreate 3D scenes from 2D color/depth images
- Computational photography combines computer vision and computer graphics (see next slide)
- For more on computer vision:
  - Professor James Tompkins: CSCI 1430 (Introduction to Computer Vision, Spring), CSCI 2951I (Computer Vision for Graphics and Interaction, Fall)
  - Other departments: CLPS 1520 (Computational Vision, Fall), CLPS 1590 (Visualizing Vision, Spring), ENGN 2560 (Computer Vision, Spring)

Example: Style transfer for headshot portraits

- Computational vision
  - Matches points in the input image and the example image
  - Keypoint detection and correspondence
- Image processing
  - Match local statistics
    - Local contrast
    - Tone
    - Detail
- Allow amateurs to easily produce great photos!

(a) *Input: a casual face photo* (b) *Outputs: new headshots with the styles transferred from the examples. The insets show the examples.*

Example: Style transfer for general photos

Input image

Example image: Ansel Adams

Output image

Microsoft Kinect

- Uses computer vision to "see" your body's shape
  - Extract multiple "skeletons" from depth image
    - Body as a controller
    - Gesture recognition
  - Facial recognition
- Works with cheap hardware
  - RGB camera
  - CMOS depth sensor
    - Projected infrared pattern to see in darkness
  - Total cost around $100
- Current research uses Kinects to construct 3D models
  - [DynamicFusion - Kinect Fusion]

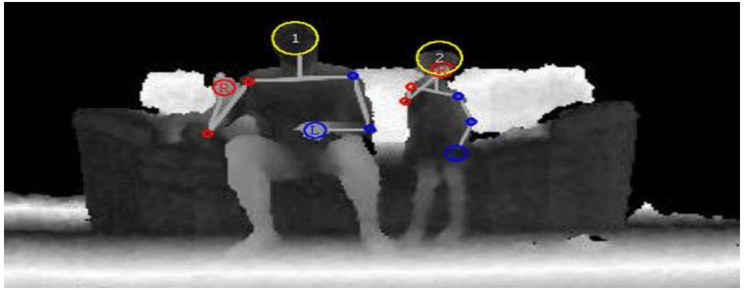

Joints of skeletons on top of depth map

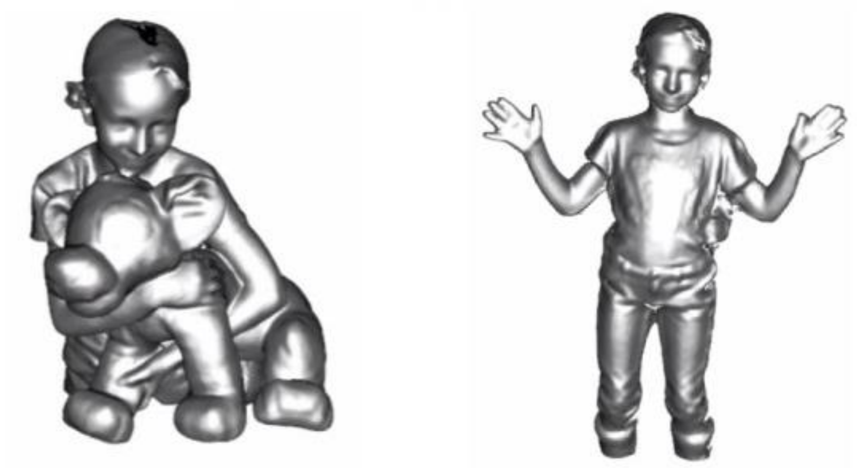

DynamicFusion - using Kinect

3D Mapping and Augmented Reality

- Extensive research capturing 3d information from color and depth images
- Can be used for many purposes
  - Paleontology/Archeology
  - Performance capture (movies, games)
  - Architects and interior design
  - Augmented reality
  - Engineering

Augmented reality with Microsoft Hololens

Outline

- Overview
- Example Applications
- **Jaggies & Aliasing** (next class)
- Sampling & Duals
- Convolution
- Filtering
- Scaling
- Reconstruction
- Scaling, continued
- Implementation