Final Exam Study Guide CSE/EE 486 Fall 2007

Lecture 2 Intensity Sufaces and Gradients
Image visualized as surface. Terrain concepts.
Gradient of functions in 1D and 2D
Numerical derivatives. Taylor series. Finite differences.
Image gradients.
Functions of gradients: Magnitude. Orientation.

Lecture 3 Linear Operators
Linear filters.
Convolution vs Cross Correlation
How to perform convolution. Border handling issues.
Properties of convolution (e.g. commutative, associative, linear)
Finite difference filters

Lecture 4 Smoothing
Gaussian image noise model
How smoothing reduces noise
When is smoothing good? When is it bad?
Box filter (Averaging filter)
Gaussian smoothing filter
Separability. Cascading.
Formula for sigma of cascaded Gaussians.

Lecture 5 Gradients and Edge Detection
Smoothing and differentiation (combine filters)
Prewitt filter. Sobel Filter.
Derivative of Gaussian filter.
Edges are intensity discontinuities.
Types of edges: step, ramp, ridge
Relationship between step/ramp edge and first derivative
Finding edges by thresholding on gradient magnitude
Canny criteria: good detection, good localization, low false positives
Optimal filter well approximated by deriv of Gaussian filter
Non-maximum suppression for edge thinning
Hysteresis Thresholding definition

Lecture 6 Harris Corner Detectors
Corners have good localization in all directions
Visualization of small window shifting over uniform region, edge, and corner
Harris mathematics: sum of squared differences of shifted region
First order approx (Taylor series for 2D functions)
Harris detector second moment matrix
Classification via eigenvalues
Corner response measure R

Lecture 7 Correspondence Matching
The Correspondence Problem.
Dense correspondence vs sparse correspondence
Correlation-based methods
SSD (sum of squared differences)
Relation between SSD and correlation
NCC (Normalized Cross Correlation)
NCC response range is -1 to 1

Lecture 8 Introduction to Stereo
Can recover depth from two images (and why you can't from just one)
Parallax
Anaglyphs and red/cyan glasses. Basic idea of why it works.
Julesz random dot experiment. What it is. What it proves.
Geometry of a simple stereo system (right camera displaced along X axis)
Stereo disparity. Equation relating depth, disparity, baseline and focal length.
Meaning of epipolar constraint and epipolar line.
Why are matches constrained to lie along an epipolar line?

Lecture 9 Stereo Algorithms
Concept of Disparity space image (DSI)
Lowest cost path through the DSI
Scanline consistency: Ordering constraint. Right and left occlusions.
Cox and Hingorani solution to interscanline matching
  Dynamic programming
  Path can proceed in three directions (match, left occlude,right occlude)
  Costs for each choice of direction to proceed

Lecture 10 Image Pyramids
Cascaded Gaussians revisited
Multiresolution Pyramid Data Structure
Basic operations: smooth, downsample, upsample
Application: making thumbnail images for web pages
  subsampling leads to high-frequency artifacts (aliasing)
  smoothing before subsampling prevents aliasing
Scale space: different scales emphasize features of different sizes

Lecture 11 LoG Blob Detection
Numerical second derivatives (Taylor series)
Finite difference Laplacian filter
Second derivative of Gaussian;
Laplacian of Gaussian (LoG) filter
Finding edges as Zero-crossings of LoG (or DoG)
Approx LoG by difference of Gaussian (DoG)
  Leads to efficient implementation due to Gaussian separability, cascadability
Applications of LoG: blob finding; compression.

Lecture 12 Camera Projection (Extrinsics)
Forward Projection Model
Relation between World Coords ; camera cords; image (film) cords; pixel coords
Homogeneous coordinates
Perspective Projection
Matrix equations relating world to camera coordinate systems
Extrinsic Parameters
  Camera Offset
  Camera Rotation
  Relation between offset/rotation and rotation/translation

Lecture 13 Camera Projection (Intrinsics)
Intrinsic Parameters
  relating film coordinates and pixel coordinates
  Focal length, scalex, scaley, offsetx, offsety
  Representation as a matrix
Geometric Image mappings.
  Linear mappings can be written as matrices
  2D planar transformations
  Translation, scale, rotation
  Euclidean, similarity, affine, and projective mappings
  Effects of each type of mapping
  What properties do they preserve (e.g. orientation, length, angles, parallelism, …)

Lecture 14 Parameter Estimation
Fitting parameterized models
Determining how many point correspondences needed
Least squares line fitting
Be able to solve simple least squares problems
  (sum of square error; take derivatives and set to zero; form matrix equation to solve)
Algebraic distance vs orthogonal distance
Image Warping
  Forwards vs Backward Warping
  Bilinear Interpolation

Lecture 15 Robust Estimation and RANSAC
concept of inliers vs outliers
General idea behind RANSAC procedure
  Sample minimal set of points; fit entity; count inliers supporting that entity; repeat

Lecture 16 Planar Homographies
Homographies
  Nonlinear transformation of 2D coords
  Linear transformation if you use homogeneous coords
Projection equations for points on a planar surface
  Frontal plane, calibrated camera: Similarity transformation
  Frontal plane, uncalibrated camera: Affine transformation
  Arbitrary plane, calibrated camera: Homography
  Arbitrary plane, uncalibrated camera: Homography

Lecture 17 Stabilization and Mosaicing

Video Stabilization

  Match all sequence frames to one reference frame (all must have some overlap)

  Frame to frame chaining of transformations (only need pairwise overlap)

Ghosting (out of plane pixels don't map correctly due to parallax)

Images from rotating camera

  Images are related by homography, regardless of scene structure!!!!

Panoramic mosaics (and Quicktime VR)

Intensity/Color Blending: averaging; feathering

Lecture 18 Generalized Stereo: Epipolar Geometry

Epipolar Geometry Concepts

  Epipoles

  Epipolar Lines

  Conjugate Epipolar lines

Lecture 19 The Essential/Fundamental Matrix

Basic math behind the Essential matrix E

  Plane formed by camera centers and 3D scene point

  That plane intersects image planes in conjugate epipolar lines

Longuet-Higgens equation

Fundamental Matrix F vs Essential matrix E

  Essential matrix – cameras calibrated

  Fundamental matrix – cameras can be uncalibrated

How to use F (or E) to map points in one image to lines in another

How to map points from im1 into lines in im2

How to map points from im2 into lines in im1

Lecture 20 The Eight-Point Algorithm

How to compute F from point correspondences

  Set up least squares equations

  Homogeneous linear system of equations (solve using SVD)

  Needs at least 8 points

Practical issues

  8 points chosen should be well distributed in 3D (e.g. not coplanar)

  Also they should be well distributed in the image (e.g. not all in one little patch)

Stereo Rectification

  Warp two stereo images so that point matches lie along corresponding scan lines

Algorithm1

  Apply homographies to map epipoles to infinity along X axis, such that conjugate epipolar lines are corresponding rows in the two warped images

  Problem: can't handle epipoles in the image

Algorithm 2 : polar rectification

  Nonlinear (straight lines don't remain straight), but handles ALL configurations, including epipoles in the image

Lecture 21 Stereo Reconstruction

Stereo Reconstruction

Calibrated cameras: can recover full scene structure

Only intrinsic params known: can recover up to unknown scale factor

Uncalibrated cameras: can recover up to unknown 3D projective transformation

Lecture 22 Camera Motion (T&V 8.1&8.2)
Motion Field: projection of 3D relative velocity onto 2D image plane
Optic Flow: observed motion of brightness patterns in image
Motion field and optic flow are not equivalent.
  No motion, but brightness changes (e.g. changes due to illumation)
  Motion, but not observable in image (e.g. no texture)
Motion field equations
  Translation vs rotation components
  Rotation component does not depend on depth
Pure translational motion
  Focus of expansion/contraction
  Vector length inversely proportional to scene depth

Lecture 23 Flow Estimation (T&V 8.3&8.4)
Brightness constancy equation
The aperature problem
  Only component of flow perpendicular to gradient is observable
Computing optic flow: differential vs matching methods
Differential methods: Lucas-Kanade (LK) algorithm
  Assume flow constant over patch
  Form least squares system with more equations than unknowns
  When is the LK system ill-conditioned vs well-conditioned?
Matching approaches
  Methods based on normalized cross correlation of corners

Lecture 24 Video Change Detection (a little is in T&V 8.6)
Pros and Cons of different approaches
  Background subtraction
  Frame Differencing
  Adaptive backgvround subtraction
  Persistent frame differencing
Two-frame differencing vs Three-frame differencing
Appearance-based tracking
Blob merge and split

Lecture 25 Structure From Motion
What is "Match move" application?
Perspective vs weak perspective vs orthographic projection
Factorization
  Forming observation matrix W by tracking points
  W = M * S
  Rank theorem (rank of W is at most 3)
  Decomposing W using SVD, and solving for M and S

Lecture 26 Color and Light
Light Transport: Illumination + material + sensor
Spectral power distribution of light source
Surface reflection: Specular vs Lambertian surfaces
Spectral albedo
Human Vision: Rods and Cones
How SPD + albedo + sensor response combine to determine color value
Metamers: different colors that look the same
  under certain illumination conditions
  for certain observers
Why is color constancy hard?

Lecture 27 Skin Color / Color Classification
Physics of Skin Color
  Spectral Albedo of Skin (biology and photometry)
  Illuminant SPD and Camera Response
  Skin color locus
Jones and Rehg
  Color histograms
  Likelihood ratios
  Mixture of Gaussian models (jrmogskin.m program)
Sample Applications: Human-Computer Interaction
Sample Applications: Adult image classification

Lecture 28 Intro to Tracking (some in T&V 8.4.2 and Appendix A.8)
Data association - assign M new observations to N existing trajectories
  Motion prediction (e.g. Kalman filtering)
  Gating (simple geometric filtering)
  Similarity scores / affinity matri
Simple methods for data association
  Global Nearest Neighbor (GNN)
  Linear Assignment Problem
  Softassign

Lecture 29 Video Tracking: Mean-Shift
Mean-shift algorithm to find mode of a sample of points
Adapting mean-shift for image tracking.
Color histograms

Lecture 30 Video Tracking: Lucas-Kanade
Generalizes computation of optical flow
  from translational flow to affine flow over patch
Definition of Jacobian of a transformation

Lecture 31 Object recognition: SIFT keys

Local feature-based approach to object recognition
- Simple example: correlation of collects of Harris corner patches

SIFT = scale invariant feature transform
- Accounts for differences in appearance due to scale, rotation and translation
- Scale space for selecting scale (radius of feature)
- Orientation histogram for selecting rotation
- Spatial-orientation histograms of gradients to creating feature vector

Lecture 32 Object recognition: PCA-based methods

Concept of Eigenspace
- Useful method of data compression
- Discovery of a lower dimensional linear subspace that spans the data

General idea of computing an eigenspace from points
- Subtract center of mass; form scatter matrix
- Perform Singular Value Decomposition (SVD)
- Use eigenvectors associated with largest magnitude eigenvalues
  - These form a basis that 1) rotates and 2) projects into low-dimensional subspace

Examples: Murase and Nayar

Examples: Eigenfaces