

The second exam covers all lectures, except that there are not questions specifically about texture synthesis, intelligent scissors, graph cuts, or NeRF. The exam is roughly split between the fundamentals introduced in the first half of the semester (particularly filtering and template matching; light, color, and contrast enhancement; warping and morphing; single view geometry and camera models) and the second half (image-based lighting; keypoint matching, alignment, and stitching; object/image retrieval and other applications). I've also outlined what I consider to be the key ideas of each section, and nearly all of the exam questions are related to these key ideas.

Filtering

Key ideas

- Design a linear filter based on an equation for local neighborhoods, or to perform an operation such as local averaging, or to achieve a desired effect such as smoothing or edge detection
- Compute the response to a linear filter
- Familiarity with definition and use of padding and kernel size
- Method and appropriate use for non-linear filters, such as median and bilateral filters
- How to perform filtering in frequency domain and representation of images and filters in terms of frequencies

1. Design a 3x3 linear filter, such that the output will be

$$im_{out}(i, j) = 4 \cdot im_{in}(i, j) - im_{in}(i - 1, j) - im_{in}(i + 1, j) - im_{in}(i, j - 1) - im_{in}(i, j + 1)$$

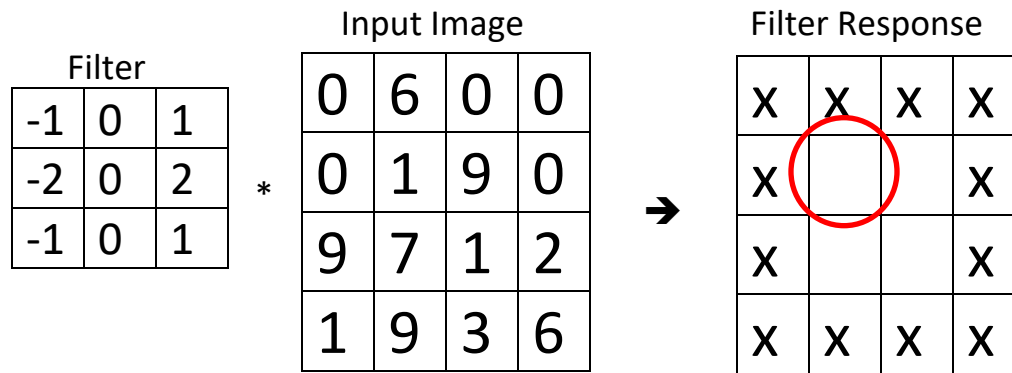
where (i, j) is a pixel coordinate that is not on the border, and im_{in} and im_{out} are the input image and output of the filter.

Filter

2. Which statement about the filter below is most true?


0	0	0
-1	0	1
0	0	0

- a. The filter is a high-pass filter (mostly high frequencies are preserved)
- b. The filter is a low-pass filter (mostly low frequencies are preserved)
- c. The filter preserves both high frequencies and low frequencies about equally
- d. Which frequencies are preserved depends on the input




3. What is the response at the circled position?

- a. 19
- b. 10
- c. 1
- d. None of these

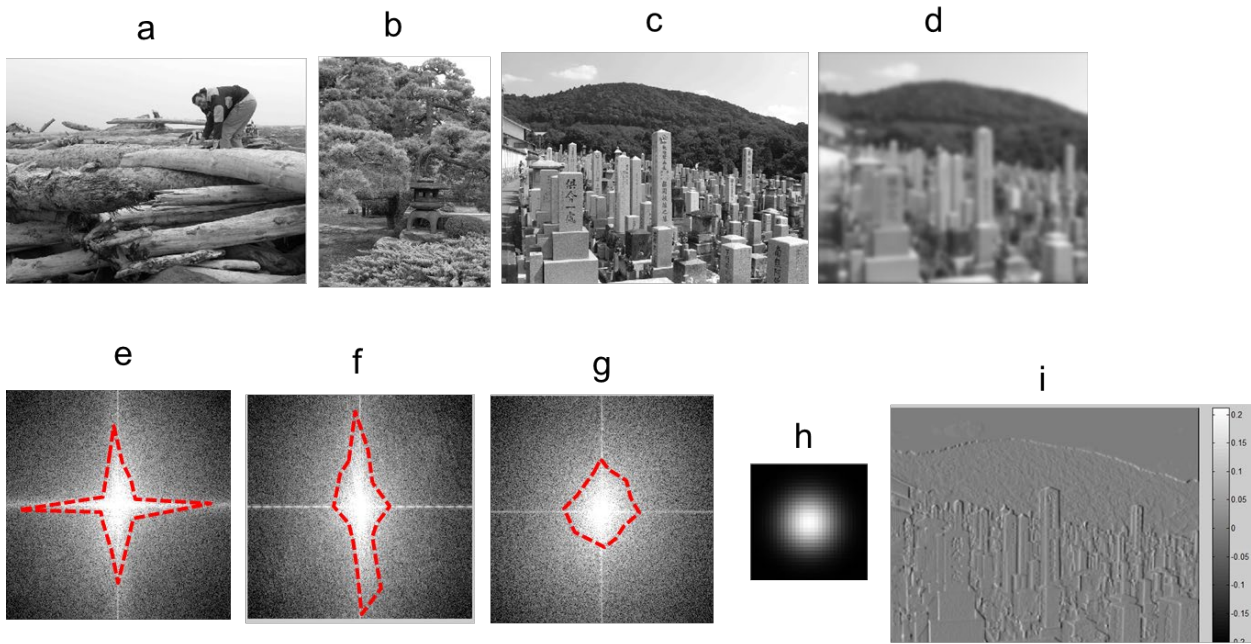


4. What is the response at the circled position?
- 18
 - 16
 - 12
 - None of these



5. What is the response at the circled position?
- 3
 - 23
 - 5
 - None of these

6. Match and answer



FFT Images (center is 0 frequency)

- i. Which FFT image best matches (a)? _____
- ii. Which FFT image best matches (b)? _____
- iii. Object _____ convolved with Object _____ produced (d)
- iv. Below, write values for a 3x3 filter that when applied to (c) would result in something similar to image (i)

7. True or False (2 pts)

$$\text{filter2}(A, k \cdot B) - \text{filter2}(A, k \cdot C) = k \cdot \text{filter2}(A, B - C) \quad (k \text{ is a scalar})$$

8. Which of the following is not true about the Fourier transform?
- Can be used to improve the speed of linear filtering
 - Provides a different view of filters that is sometimes easier to interpret
 - Adds information to the image
 - Preserves the energy in the image

Templates and image pyramids

Key ideas

- *Implementation and applications of SSD and NCC*
 - *Downsampling / upsampling and anti-aliasing*
 - *How to compute Gaussian and Laplacian pyramids*
9. Suppose we have computed the normalized cross correlation (NCC), where higher means a better match (max value of 1), and sum of squared differences (SSD), where lower means a better match (min value of 0), for two non-uniform patches. Which of the following are true?
- NCC=1 if and only if SSD=0
 - If SSD=0, then NCC=1
 - If NCC=1, then SSD=0
 - None of these are necessarily true
10. Is there any case where NCC is undefined, but SSD can be computed?
- No
 - Yes, if both patches are equal
 - Yes, if at least one patch has uniform values
 - Yes, if either patch has negative values
11. Which of these statements are false?
- NCC is invariant to shifts in intensity
 - SSD is usually a better matching cost/score than NCC for calibrated stereo pairs and tracking in video
 - NCC is faster to compute than SSD
 - NCC is usually a better matching cost/score than SSD for comparing patches of images from images taken at different times or different cameras

12. If you want to downsample an image by a factor of 2, what is the problem with directly sampling pixels on every other row and column, and how can you fix it?
- Aliasing, apply a high pass filter before sampling
 - Aliasing, apply a low pass filter before sampling
 - Extrapolation, apply a high pass filter before sampling
 - Extrapolation, apply a low pass filter before sampling
 - There is no problem
13. True or False: An image can be losslessly reconstructed from its Laplacian pyramid

Lighting and Color

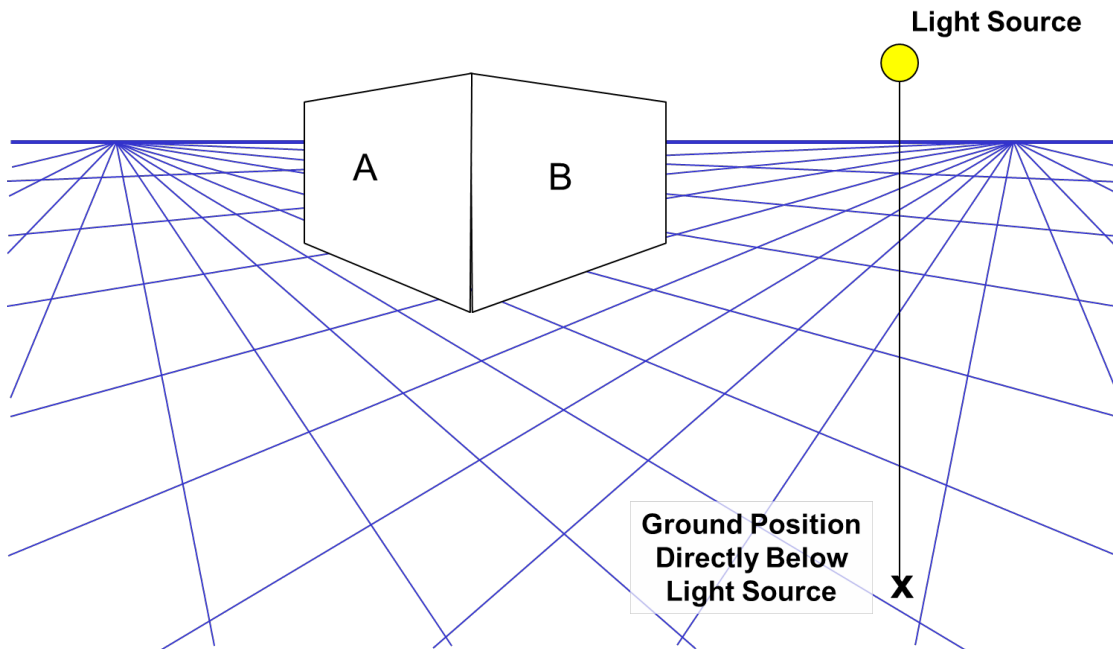
Key ideas

- *Diffuse vs. specular reflection*
- *Effects of material and geometry on reflection (albedo, surface orientation, viewing direction, etc.)*
- *Use and advantages/disadvantages of color spaces: RGB, Lab, HSV, YCrCb*
- *How to compute histograms*
- *How to improve contrast or balance intensities/color*

14. Under which type of reflection does incoming light scatter from the surface?
- Diffuse reflection
 - Specular reflection
15. Which of the following factors impact the observed intensities of diffusely reflected light? (can choose more than one)
- The intensities of the light coming into the reflecting surface
 - The albedo of the reflecting surface
 - The orientation of the reflecting surface compared to the direction of incoming light
 - The orientation of the camera/observer compared to the surface normal
16. Which of the following factors impact the observed intensities of specularly reflected light? (can choose more than one)
- The intensities of the light coming into the reflecting surface
 - The albedo of the reflecting surface

- c. The orientation of the reflecting surface compared to the direction of incoming light
- d. The orientation of the camera/observer compared to the surface normal

17. In the illustration below, assume that each side of the depicted block has the same albedo. Which side will look brighter: side A or side B
- a. A will look brighter
 - b. B will look brighter



18. Imagine that a one-eyed robot is in a dark world, and the only light source is emanating from its eye. Is it possible for the robot to see the shadows that its light source directly creates?
- a. No
 - b. Yes

19. For each material, indicate whether its reflection can be modeled as Lambertian, specular, or mixed:
- a. Balloon
 - b. Cotton shirt
 - c. Polished wood

- d. Mirror
 - e. Cement block
20. Which color space(s) have separate luminance (brightness) and chrominance (color) channels?
(can select more than one)
- a. RGB
 - b. HSV
 - c. L*a*b
 - d. YCbCr
21. Which color space is based on perceptual studies of just noticeable differences, so that Euclidean distance between nearby color values corresponds to human perception?
- a. RGB
 - b. HSV
 - c. L*a*b
 - d. YCbCr
22. What is a disadvantage of the RGB color space compared to others?
- a. The channels are highly correlated, which makes compression harder
 - b. People don't intuitively think of color in terms of red, green, and blue values
 - c. Both (a) and (b)
 - d. None of these
23. If most of the image intensity values are in the mid-range, such that there are few dark or bright pixels, which of these techniques are best to improve contrast?
- a. Gamma adjustment, $\gamma > 1$
 - b. Gamma adjustment, $\gamma < 1$
 - c. Histogram equalization
24. If an image looks washed out with many high intensities but not many low, which of these techniques is best to improve contrast?
- a. Gamma adjustment, $\gamma > 1$
 - b. Gamma adjustment, $\gamma < 1$
 - c. Histogram equalization

25. If an image is too dark with many low intensities but not many high, which of these techniques is best to improve contrast?
- Gamma adjustment, $\gamma > 1$
 - Gamma adjustment, $\gamma < 1$
 - Histogram equalization

Pasting: Compositing and Blending

Key ideas

- Method and use of alpha matting, poisson blending, and laplacian pyramid blending



26. Mary wants to cut out the flower region and overlay it on the image of the sunset. She wants to avoid distorting colors or making it look pixelated. What blending method is most appropriate?
- Use an alpha matte to determine whether the color for each pixel should come from the foreground or background image
 - Use Laplacian pyramid blending
 - Poisson blending (or gradient domain editing)



27. Ann wants to see what her arm will look like if she adds this cat drawing as a tattoo. What blending method is most appropriate to overlay the cat on the arm?
- Use an alpha matte to determine whether the color for each pixel should come from the foreground or background image
 - Use Laplacian pyramid blending
 - Poisson blending (or gradient domain editing)

28. Which of the following statements is false?
- Poisson blending preserves gradient of the source region without changing the background
 - Pixels are pretty blocky if you use cut and paste
 - Foreground colors stay the same when we do Poisson blending
 - One should feather when doing alpha compositing


Image Warping and Morphing

Key ideas

- *Use/conversion of homogeneous vs Cartesian coordinates*
- *Matrix form of transformations: rotation, shear, scale, translation, affine, perspective, similarity*
- *Invariances of transformations*
- *How to solve for transformations from corresponding points*
- *Forward vs. inverse warping/texturing*
- *How to morph shape/texture from one image to another given corresponding points*

29. Which homogeneous coordinate is different from others?
- (9, 6, 3)
 - (12, 8, 4)
 - (48, 32, 16)
 - (4, 2, 1)
30. What properties do affine transformations and projective transformations have in common?
- Parallel lines remain parallel under projection
 - Origin does not necessarily map to origin
 - Ratios are preserved

31. The figure is the output of applying one of the transformations below to a square with vertices $(-1, 1)$, $(-1, -1)$, $(1, 1)$, $(1, -1)$. Which transformation is it (pick most restrictive possible)?




- a. Affine
- b. Projective
- c. Rigid
- d. Similarity

32. What transformations can you model with $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$? (can choose multiple)

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

- a. Translation
- b. Scale
- c. Rotation about the Origin
- d. Shear
- e. Affine
- f. Perspective

33. For (i-iii): Suppose that the image of a face is scaled uniformly by factor s and translated in either direction by t_x and t_y .



i. Write down the equation for a transformed point (x', y') as a function of the original point (x, y) in terms of t_x , t_y , and s .

ii. Write the transformation in matrix form in terms of t_x , t_y , s , x , y , x' , and y' :

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} & \\ & \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

e. If you are given two pairs of corresponding points: (x_1, y_1) to (x_1', y_1') and (x_2, y_2) to (x_2', y_2') , how do you solve for the transformation parameters t_x , t_y , s ? Write down the system of equations in a matrix form (you don't need to solve it):


$$\begin{bmatrix} & & \\ & & \\ & & \end{bmatrix} \begin{bmatrix} s \\ t_x \\ t_y \end{bmatrix} \cong \begin{bmatrix} \\ \\ \end{bmatrix}$$

34. Put these operations in the correct order to perform morphing: (op1) determine corresponding points; (op2) compute average shape; (op3) define triangulation; (op4) warp both images toward the average shape; (op5) cross dissolve.
- (op2) -> (op1) -> (op3) -> (op4) -> (op5)
 - (op1) -> (op2) -> (op4) -> (op3) -> (op5)
 - (op1) -> (op3) -> (op2) -> (op4) -> (op5)
 - (op3) -> (op1) -> (op2) -> (op4) -> (op5)
35. What is the lowest parameter transformation that can map from any point in any one triangle to a corresponding point in any other triangle?
- Translation
 - Similarity
 - Affine
 - Perspective

Single View Metrology and Camera Model

Key ideas

- *Definition of vanishing points/lines and use for comparing object sizes*
- *How to project a 3D point onto the image in a perspective camera model, using equations or matrix operations*
- *Definition and parameterization of intrinsic and extrinsic matrices*
- *Effects of camera parameters such as aperture, lens, focal length, sensor size, image resolution, shutter time*



36. The column is 10 ft tall with the lines evenly marking 1 ft increments, and Yoshi (dinosaur) is on the ground. What is the approximate (within 0.5 ft) height of Yoshi?

- a. 4.5 ft
- b. 6.0 ft
- c. 6.5 ft
- d. 7.0 ft
- e. None of these are within 0.5 ft

37. The column is 10 ft tall with the lines evenly marking 1 ft increments, and Yoshi (dinosaur) is on the ground. What is the approximate (within 0.5 ft) height of the camera?

- a. 4.5 ft
- b. 6.0 ft
- c. 6.5 ft
- d. 7.0 ft
- e. None of these are within 0.5 ft

38. Draw a diagram of the projection of a 3D point onto an image pixel and label: vertical component of 3D position Y ; distance from camera along camera axis Z ; pixel column u ; pixel row v ; principal point u_0, v_0 ; focal length f ; center of projection (aka camera center).

39. Suppose a point's position is X, Y, Z in the camera's 3D coordinates. Given focal length f and principal point (u_0, v_0) , what is the pixel row v of the projected point.

- a. $v = \frac{fY}{X} + v_0$
- b. $v = \frac{X}{fY} + u_0 + v_0$
- c. $v = fY + Zv_0$
- d. $v = \frac{fX}{Z} + u_0$
- e. None of these are correct

40. Complete the intrinsic parameter matrix that projects from a 3D point to a 2D homogenous image coordinate. Assume zero-skew and unit aspect ratio.

$$w \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \begin{bmatrix} & & \\ & & \\ & & \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

41. Let $(X=2, Y = 0, Z = 10)$, $f=2$, $u_0 = 0.5$, $v_0=0.5$. Solve for u .

- a. 0.9
- b. 0.7
- c. 0.4
- d. 0.5

42. Suppose you are taking a picture in low-light conditions and your first photo looks grainy. Which of the following can increase the amount of light coming into the camera without changing the field of view? (can select more than one)
- Increase aperture
 - Decrease aperture
 - Increase focal length
 - Decrease focal length
 - Increase shutter time
 - Decrease shutter time
43. How can you increase the field of view?
- Increase aperture
 - Decrease aperture
 - Increase focal length
 - Decrease focal length
 - Increase shutter time
 - Decrease shutter time

Single-view 3D Reconstruction

Key ideas

- *How to solve for camera parameters from vanishing points*
 - *Single-view scene models, such as single-point perspective stage, ground + billboards*
44. Is it possible to solve for intrinsic matrix using vanishing points that correspond to three orthogonal directions? If so, how?
45. Suppose you know the camera intrinsics and vanishing line of the ground plane. What additional piece of information could be used to recover the height of the camera?
- The height of an object that is visibly connected to the ground
 - The height of an object that is not visibly connected to the ground
 - Knowledge of the sensor width
 - The height of a scene point that projects onto the horizon

46. True or False: Single-view 3D reconstruction can work because there is a unique solution to the 3D scene model that results in an observed image of a natural scene.

Image-based Lighting

Key ideas

- *How to create HDR image from multiple LDR images*
- *Equirectangular image representation*
- *Method to create light environment from photos of light probe*
- *Method to relight an object in an image using an environment map*

47. If you want to create a HDR image from multiple images, which of these should vary from image to image?
- a. Shutter time only
 - b. Camera position (translation) only
 - c. Camera orientation only
 - d. Camera position and orientation

Matching, Alignment, Stitching, RANSAC

Key ideas

- *Keypoint detection methods and advantages/disadvantages (DoG, Harris)*
- *Keypoint representations (intensity patch, SIFT)*
- *Keypoint matching criteria (SSD, Lowe's ratio, NCC)*
- *Solving for transformations given matching points (with or without outlier matches), inc. RANSAC and least squares*
- *Perspective and cylindrical image/canvas representations*

(For the two questions below) Suppose we think that two sets of 2D points are related by a homography. We want to estimate the transformation parameters using RANSAC given corresponding pairs of points (which may contain some bad correspondences).

For $t = 1$ to 100:

 Sample N pairs of points

 Solve for the homography transform from sampled pairs

 Get score for this transformation, check if it's the best so far

Return the best scoring transformation

48. What (integer) value should N have? _____

49. If you expect that a large fraction of the correspondences are incorrect (points in a pair don't correspond to the same scene point), what should you do to improve your result?
- Increase the number of iterations (e.g. from 100 to 5000)
 - Increase N, the number of sampled points
 - Decrease N, the number of sampled points
 - Instead of RANSAC, compute a least squares estimate of the transformation for all points
50. Suppose you have two sets of corresponding points. The point positions may not be exact, but the correspondences are correct (no outliers). Which method is most suitable to estimate the transformation between the points?
- Least squares estimate using all points
 - RANSAC
 - Neither can give a good estimate
51. Which of these is an advantage of the keypoint detection approach used by SIFT, compared to the Harris corner detector? (can select more than one)
- Faster to compute
 - More precisely repeatable detected positions
 - Provides an estimate of relative scale between two detected points

Object/image recognition and retrieval

Key ideas

- Techniques to improve speed and accuracy of keypoint-based recognition/retrieval (K-tree clustering, histograms of visual words, inverse index, tf-idf, geometric verification)*
52. Which of these techniques are primarily used for interest-point based image retrieval to improve the **speed** (ignoring effects on accuracy) of matching in large image sets?
- K-means clustering
 - Term frequency – inverse document frequency (tf-idf) score
 - Index table of which images contain each visual word
 - Geometric verification of matched keypoints

53. Which of the following increase the term-frequency inverse document frequency weight?
- The word appears often in the document
 - The word appears often in the database
 - The word appears rarely in the document
 - The word appears rarely in the database

Opportunities of Scale

Key ideas

- *Nearest neighbor retrieval as an inference/classification method*
- *Encoder/decoder framework (conceptual)*

54. True or False: With enough data, a simple nearest neighbor lookup can be used to solve complex image labeling and transformation problems.

55. True or False: If you have millions of samples, the features used to represent the samples do not have much impact on the performance of nearest neighbor retrieval

Generating / detecting fakes

Key ideas

- *Understand the techniques behind the fake detection methods (i.e. based on image statistics, mosaicking regularity, JPEG details, lighting direction)*
- *Understand the basic concept of image generation (pix2pix and CycleGAN) – what do they do, what kinds of inputs, what objective or loss are they optimizing*

56. What do we know about the surface normal corresponding to the image boundary of a rounded object?

- The normal is parallel to the image plane ($N_z = 0$)
- N_x and N_y are determined by the orientation of the boundary
- Neither a nor b
- Both a and b

57. Suppose you have an image that was compressed with JPEG quality of 65. Which of these will lead to further degradation of the original image?
- Resaving the image with JPEG quality of 80
 - Resaving the image with JPEG quality of 65
 - Resaving the image with JPEG quality of 50
 - a, b, and c
 - a and c
 - b and c
 - none of these

Kinect sensor and its use for body pose estimation

Key ideas

- Use of stereo/projector pair for stereo depth estimation*
- In which cases will stereo likely succeed or fail*




58. Why is the Kinect sensor unable to produce depth values for the countertop, pointed to by “B”?
(black means Kinect did not return depth values)
- The surface is too slanted
 - The countertop is shiny (specular reflection), so the projected pattern is not received or matched by the camera
 - The surface doesn't have enough texture
 - The area of the surface in the image is too small

59. The Kinect stereo algorithm uses NCC, rather than SSD for a match score. Is NCC clearly a better choice than SSD in this case. If so, why?
- No, either SSD or NCC could be effective in this case
 - Yes, the projected image is known so we know what intensities the sensed image should match
 - Yes, the sensed image will be locally proportional to the projected image, but the overall intensity could be different

Computational approaches to cameras

Key ideas

- *Why the coded aperture enables better depth estimation and deblurring*
- *How the plenoptic camera works to capture both the position and angle of rays coming into camera*



60. Circle which camera has the advantage in each of the following ways. (3 pts)

C = conventional, **P** = plenoptic (shown above), **N** = neither

- | | | | |
|--|---|---|---|
| i. Can produce a higher resolution image | C | P | N |
| ii. Captures more light | C | P | N |
| iii. Encodes images with multiple camera centers | C | P | N |

Face detection, recognition, transformation

Key ideas

- *Basic idea of classic face detection (image pyramid, sliding window, extract features, apply trained classifier), and what some of the challenges are*
- *Face recognition via detection, alignment, feature extraction, and similarity comparison; what parts of the modern face recognizer provides the most improvement*
- *How to make a face look more masculine, happy, young, etc. through warping, averaging, and interpolation*

61. Which of these is **not** a step in the warping/transformation method of making an input young face look older? (can choose more than one)
- Obtain an average “young” face, including triangulated feature points
 - Obtain an average “old” face, including corresponding feature points
 - Define corresponding feature points on the input face
 - Define a new triangulation on the input face using the Delaunay algorithm
 - Create a target “shape” as $X_{target} = X_{input} + \alpha \cdot (X_{old} - X_{young})$, where α is a deformation factor greater than 0
 - Obtain the affine transformations for each triangle in X_{target} to X_{input} , X_{old} , and X_{young}
 - For each pixel in the target image, map to the input, young, and old faces and store the color as $C_{target} = C_{input} + \alpha \cdot (C_{old} - C_{young})$
 - Apply a sharpening filter to accentuate wrinkles

Video magnification

Key ideas

- Basic workings and advantages/disadvantages of the point tracking approach, the pixelwise intensity magnification, and the phase magnification approaches to motion magnification*

62. Which of these is an advantage of the point tracking method for magnification, compared to the Eulerian and Phase-based?
- Point tracking method is simpler and more robust
 - Point tracking method usually generates more compelling results
 - Point tracking method can be applied to large non-periodic motions
63. What are advantages of the phase magnification approach, compared to the pixel magnification approach?
- Performs translation more directly and so can apply to larger movements than intensity magnification
 - Works well with non-period motions
 - Robust to imaging noise