

Your Name: \_\_\_\_\_ Your T-Square ID: \_\_\_\_\_

- **You must show all work to receive full credit.** Correct answers with no work shown will receive minimal partial credit, while incorrect answers with correct work shown will receive generous partial credit. Illegible answers are wrong answers.
- **Integrity:** By taking this quiz, you pledge that this is your work and you have neither given nor received inappropriate help during the taking of this quiz in compliance with the Academic Honor Code of Georgia Tech.
- **Academic Misconduct:** Academic misconduct will not be tolerated. You are to uphold the honor and integrity bestowed upon you by the Georgia Institute of Technology.
  - Keep your eyes on your own paper.
  - Do your best to prevent anyone else from seeing your work.
  - Do NOT communicate with anyone other than a proctor for ANY reason in ANY language in ANY manner.
  - Do NOT share ANYTHING during the quiz. (This includes no sharing of pencils, paper, erasers or calculators).
  - Do not use notes or books, etc during the quiz.

Problem	Points	Lost	Gained	Running Total	Grader
1	5				
2	5				
3	5				
4	3				
5	5				
6	3				
7	5				
8	3				
9	6				
Total:	40				

1. (5 points) When seam carving an image, you need to start with an “energy function”. Describe what a good energy function does, and how it is used by the seam carving algorithm.

A good energy function will determine which parts of the image are “important” so that they are not removed. It is used to find the “minimum energy” pixels of the image that are safe to remove.

Grading: +5 for idea that the energy function identifies important (or not important) pixels.  
+2/5 for mentioning specific example such as Gradient Magnitude or visual saliency, but not mentioning the above.

2. (5 points) What are the attributes that detected features (as detected in an image by a feature detector) should have to make them useful?

Well localized / Locality / Found Precisely  
Repeatability / Match-ability / Saliency  
Compactness / Efficiency (compact representation)  
Scale invariant  
Rotation Invariant  
Illumination Invariant

Grading: +1 for each/any of the above, maximum of 5.

3. (5 points) When transforming image points using a 3x3 transformation matrix, the matrix has three unique “regions”. For each of the (A, B, C) regions below, describe the transform (or transforms) that particular region encodes.

A	A	B
A	A	B
C	C	1

A : Main rotation/scale/shear matrix.

B: Translation (X/Y)

C: Projective vector

Grading:

+3 for identifying that A does any 2 or 3 of: rotation/scale/shear (+2 for only one, +1 if they add transform or projection to it.)

+1 for B doing transform

+1 for C being used only for projective transforms...

4. (3 points) What is the primary mathematical difference between an image transform (such as done with a 3x3 transformation matrix) and an image warp or morph, with respect to the lines and points in the image?

A transform will preserve lines, and affects the entire image in the same general way. A warp/morph maps individual points to other individual points, and will most likely not preserve lines, and may warp/morph different parts of the image differently.

+1 for transforms preserving lines, +1 for morphs being a point-to-point mapping, +1 for global nature of transforms vs more local nature of morphing.

5. (5 points) What are the five main steps to creating a panoramic image?
- Capture the image
  - Detect and match feature points
  - Transform or warp the images into alignment
  - Blend/fade/cut the images together
  - Crop images (possibly optional)

Grading: +1 for each step.

6. (3 points) What problems does a camera have in capturing a natural scene that has a wide variety of illumination levels?

The camera sensor or storage format does not typically have as wide a range of exposure as the human eye (only eight bits, for example), so can not capture details from both dark and light areas at the same time. (You can overcome this by taking multiple photos at different exposure levels and merging them together with HDR or exposure fusion. )

Grading: +3 for lack of range in exposure capture and storage +2 for one or the other, +1 for something that sort of shows they understand the issues.

7. (5 points) What is the epipolar constraint, and what applications does it have with respect to stereo vision?

The epipolar constraint states that a point in one image lies somewhere along a single ray in the real world, and that line can be projected onto a different image/view. This allows us to search for matching features to the original point along a single line, greatly speeding up stereo correspondence matching.

Grading: +3 for idea that a point in a 2D image lies along a ray in the world, or that it can be projected onto a line in a 2<sup>nd</sup> image of the same scene from a different viewpoint.  
+2 for the application (speeding up correspondence searching).

8. (3 points) Explain how a video is stored in computer memory when working on it with python code using the OpenCV library. (Specifically, what is different from the way a single image is stored in memory?)

The data structure in memory has one more dimension (specifically, time) and can be thought of as a 3D volume (or 4D with color channels).

+3/3 : Idea that time is it's own dimension.

+1/3 : Talks about the need to compress video to save space, or using OpenCV video tools.

9. (6 points) Given the following synthetic output from an “energy function”, calculate the intermediate result as the first part of a (vertical) seam carving dynamic programming solution. **After you do this, identify and mark the seam from the bottom to the top of the “image” which has the lowest energy.**

1	2	1	2	1
1	1	1	1	2
2	0	2	2	0
2	2	2	1	2

1	2	1	2	1
2	2	2	2	3
4	2	4	4	2
4	4	4	3	4

Grading: +2 for marking one of the (two possible) correct seams (only difference is in top row)

+1 for having 1,2,1,2,1 in the top row

+2 for having 2,2,2,2,3 in the 2<sup>nd</sup> row

+1 for getting the bottom row correct (4,4,4,3,4)