You must answer all questions. For full credit, an answer must be both correct and well-presented (clear and concise). If you feel a question is ambiguous, state any assumptions that you need to make. Also, some of the questions are “essay” questions. For those, there are many correct answers. It is more important that you provide a good argument for the answers you give, than that you give the “most correct” answer. Sometimes, we are particularly looking for your ability to make a clear and concise argument based on things you are aware of, rather than to see if you can find the best possible answer, or have seen all possible research on the topic.

Question 1: The modern paper

3 parts A, B and C. C has 3 subparts

You were asked to read the paper:

Englert, P., Vien, N. A., & Toussaint, M. (2017). Inverse KKT: Learning cost functions of manipulation tasks from demonstrations. The International Journal of Robotics Research, 36(13–14),

Based on your analysis of this paper and any of its citations that you have reviewed, but also using your general knowledge of the graphics, robotics and HCI principles in your reading list, answer the following questions:

A . One key piece of the method is that it needs to be provided a set of cost terms. The method learns a policy represented as a set of weights over the cost terms.

In the paper's formulation, learning the policy (the set of weights) is efficient (usually a convex program), while creating a trajectory from a set of learned weights requires solving a non-linear program. Explain why.

B. The authors propose a ``non-parametric'' variant of their approach that does not require pre-defined cost terms, but instead uses kernel functions. Provide an intuition for how such a non-parametric approach works.

> *Hint : explain what an example kernel function might be, and how the method would make use of examples. Consider: What would the method reduce to if there were only a single example?*

C. Consider using the method to teach a robot chef to pour salt on food.

The salt is always in the same starting position on the Table (and needs to be returned) to the same position.

In each demonstration: the salt shaker starts out on (in contact with) the table, is lifted up, rotated (so the salt pours out), rotate back to upright, and place back in contact with the table.

C.1 In the parametric formulation, there would probably be a cost term for contact with the table. Such a term could be written as either an equality or inequality constraint. Which would the paper use and why (was this chosen over the alternative)?

C.2 Suppose the demonstrations have the food in different places. The method could correctly generalize to create a policy that can generalize to food in many different places. How?

C.3 Suppose when the person performing the demonstrations does the task, they take different amounts of time for each phase of the task in each demonstration. (e.g., one time they wait before picking up the salt shaker, another time they do it quickly but move to the target location more slowly). Describe either:

How such variable timing would be handled with the method? (if it can be handled well)

Why variable timing wouldn't be a big issue given the way the method is used?
(if the method can't easily handle changes in timing)

Question 2: Linear algebra and transformations

Two parts A&B

A. Matrix multiplication, in general, does not commute. That is, for two square matrices A and B the matrix products AB and BA are generally not equal.

   When we represent linear transformations as matrices, it is common to represent the result of applying two transformations in sequence by the product of their respective matrices (i.e. a non-commutative operation).

   On the other hand, rotation vectors are often used to encode 3D rotations (also known in Graphics as exponential maps, they are also analogous to angular velocities). We may often compose rotation vectors by adding them, but of course vector addition does commute, in contrast to matrix multiplication.

   Which of the two approaches correctly captures composition of transformations? If both of them are in fact correct and accurate, explain the apparent disparity. If one of the two approaches is not fully accurate, explain why it still does something reasonable despite being different.

B. Consider the 2D linear transformation (represented by a 2x2 matrix) shown in the illustration below. After applying this transformation (let's call it T), the unit square on the left is transformed into the rhombus on the right. Conceptually, the transformation stretched the square along one of the diagonals (the line x=y) so that this diagonal now has length 2, while the length of the other diagonal is now exactly equal to 1.

   From the theory of the Singular Value Decomposition, we know that any 2x2 matrix T can be written as a product $T=U*D*V^T$, where U and V are rotation matrices, and D is a diagonal matrix (corresponding to an axis-aligned scaling). Write down what those matrices would be in this case, and explain how this decomposition would work out intuitively. You can use notation like Rotate(-30^o) or Scale(3.0,0.4) to denote rotation and scaling matrices if you like – no need to write down all their elements explicitly.

C.

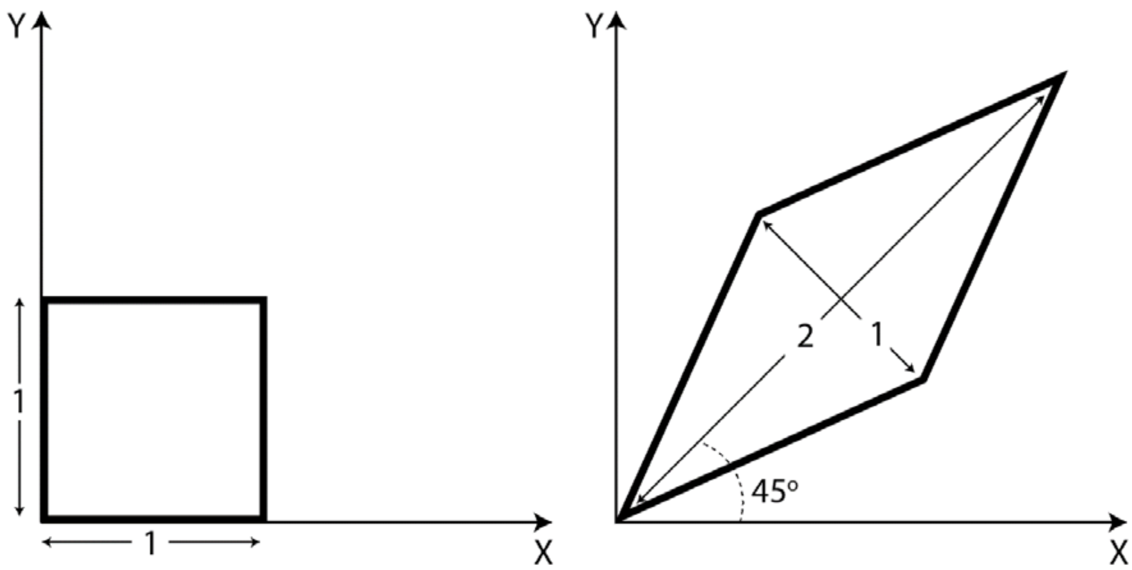

Question 3 : Color Theory

Two subparts A and B

The so-called “Rainbow Color Map” (RCM) is often used to encode a scalar quantity as a color for data visualization.

Note: while RCM was not part of the qual reading list, the ideas of visual encodings were, as well as perception, particularly of color. While you could answer this question by remembering some of the literature about RCM, you should be able to answer this question based on your understanding of visual encodings and perception

The rainbow color map encodes a continuous value as hue. Sometimes this is done in a truly spectral manner so that the value is mapped to the spectral frequency of the color, although usually the value is mapped around the color wheel. For example, the lowest value is mapped to blue, the highest value is mapped to red, and values in between are interpolated around the color wheel.

3.A There are many reasons why the RCM is a bad idea for use as a visual encoding of a continuous variable. Describe some of them.

Note: if you choose “it doesn’t work well for color blind people” as one of the reasons (and it is a good reason), you must explain how the RCM would be seen by someone with the common form(s) of color blindness.

3.B Consider showing the temperatures in Madison over the course of the year with a color with a color encoding. Give an alternative to the RCM. Discuss the pros and cons of each.

Question 4: Resampling Discrete Structures

Three subparts A, B and C.

The combination of low-pass pyramids and multi-way interpolation is a common solution for sampling discretely represented images (technically, this is re-sampling). For example, for surface texturing, MipMaps (a specific type of low-pass pyramid) combined with tri-linear interpolation is the de facto solution. Such solutions are desirable because they provide efficiency through pre-computation and quality tradeoffs.

For 4.A and 4.B, assume that the original image data is regularly sampled on a grid.

4.A: For the specific case of mip-mapped based, tri-linear interpolation in 2D, explain situations where these tradeoffs may not be acceptable. Give situations for both the performance assumptions failing to be valid, as well as where the quality assumptions may not be valid. For each, please explain the technical problem, but also give an example of a realistic application for which this may be a problem.

4.B: Consider the application of the approach to higher dimensional data (e.g. 3D volumes). How do these approaches extend? What new issues and tradeoffs arise as the dimension of the data goes up?

4.C: Consider what happens when the input data is not regularly sampled on the grid. Can you suggest ways of applying the pre-filter-pyramid plus interpolation approach to these cases? You can restrict your thinking to 2D if that is helpful.

Question 5: Measurement Design

Four subparts, A, B, C, and D.

Each semester, the Department of Computer Sciences uses a questionnaire to evaluate the courses offered by the Department. The Department is interested in redesigning this questionnaire to better capture student experience and to serve as a more valid and reliable measure of teaching quality. The questionnaire is intended to measure the course in dimensions of the instructor's teaching performance, the overall organization, quality, and usefulness of the course, and the overall student experience in the class. Based on Lazar et al.,[1] respond to the following considerations regarding the process of redesigning this measurement.

(a) Describe the types of questions and response types that such a questionnaire might include, discussing their advantages and disadvantages.
(b) List three of the common problems in survey questions that the designers of the new questionnaire must avoid.
(c) Provide the Department with recommendations on how they should assess the validity and reliability of the new questionnaire.
(d) Finally, discuss the advantages and disadvantages of administering the questionnaire online and on paper.

[1] Lazar, J., Feng, J. H., & Hochheiser, H. (2010). Research Methods in Human-Computer Interaction. Wiley.

Question 6: Quantitative Data Analysis

Four subparts, A, B, C, D

HCI researchers use a wide range of methods for statistical analysis, although the most frequent method of analysis is comparing two or more samples using a t- or an F-test. In employing this method, the researcher makes a number of assumptions regarding their data. Respond to the following according to Lazar et al.[1] and Hinton.[2]

(a) Describe three assumptions made in the use of this method.
(b) Provide an example of the violation of each assumption.
(c) Describe a precaution that the researcher can take to prevent each violation.
(d) Suggest a potential post-hoc solution to each violation that the researcher can explore after the data is collected and discuss the validity of each solution.

[1] Lazar, J., Feng, J. H., & Hochheiser, H. (2010). Research Methods in Human-Computer Interaction. Wiley.
[2] Hinton, P.R. (2004). Statistics Explained. 2nd Edition. Routledge.