

CST1+CST2
COMPUTER SCIENCE TRIPOS Part IB, Part II 50%

Thursday 9 June 2022 11:00 to 14:00 BST

COMPUTER SCIENCE Paper 7

Answer *five* questions.

Submit each question answer in a **separate** PDF. As the file name, use your candidate number, paper and question number (e.g., *1234A-p7-q6.pdf*). Also write your candidate number, paper and question number at the start of each PDF.

**You must follow the official form and
conduct instructions for this online
examination**

1 Concepts in Programming Languages

This question focuses on concepts – exact syntactic encoding is unimportant.

- (a) Explain two possible implementation behaviours of the following program

```
let test(n:int) = (let f = (λx:int.n+x); n++; return f(10))
```

Now, give a corresponding Java-like program with `test` being a method containing a lambda and explain how Java resolves this situation. [4 marks]

- (b) Explain two possible implementation behaviours of the following program

```
int n = 0; int g() = n; int f(int n) = g(); print(f(1));
```

briefly indicating how `g` would be implemented in each case. [4 marks]

- (c) ML-family languages generally restrict polymorphic exceptions. Consider:

```
exception E of ('a->'a);;
try raise (E(fun x->x)) with E(f) -> (f 1, f true);;
```

Giving reasons: (i) would this code type-check in an ML-family language? (ii) would it execute successfully? Give a modified version of the code *with similar structure* (e.g. retaining two separate function applications to `1` and `true`) but which both type-checks and successfully evaluates to `(1,true)`. [5 marks]

- (d) A blog proclaims “prototypes and virtual method tables are alternative implementations of inheritance”. Clarify what was intended. [3 marks]

- (e) Explain, with reasons, how much a Java compiler can optimise `CreateVec`, given that its definition and its calls may appear in separate Java source files

```
class IPair { final int x,y; IPair(int X,int Y) { x=X;y=Y; }
              /* other methods */    };
static IPair[] CreateVec(int n) { var V = new IPair[n];
  for (int i=0; i<n; i++) V[i] = new IPair(0,0);
  return V; }
```

[4 marks]

2 Economics, Law and Ethics

You work for a consumer electronics company that dominates the market for mobile telephone handsets and operates an ‘app store’ for the distribution of mobile phone applications.

- (a) How does the market for mobile phone handsets differ from the market for mobile phone applications? [4 marks]
- (b) Why might a company providing hardware also start distributing applications? [6 marks]
- (c) How might the company price the mobile phone handsets and applications? Provide arguments for and against the available options. [10 marks]

3 Economics, Law and Ethics

You lead the security team for a large webmail provider.

- (a) You are seeking additional resources to boost your team’s capabilities. How might you frame your request for an increased budget to increase the likelihood it will be approved? [8 marks]
- (b) A crime group is using accounts with the webmail provider to organise the importation of firearms into the UK. The police would like to know more about their operations.
 - (i) How might the police discover more about the plans? [4 marks]
 - (ii) Apply ethical philosophies to argue for and against your suggested approach. [8 marks]

4 Formal Models of Language

This question concerns a hypothesis space where each language is characterised by three binary parameters:

Parameter name	0	1
S_{right}	subject before verb	subject after verb
O_{right}	objects before verb	objects after verb
S_{drop}	subject is required	subject is optional

- (a) Complete the table of language patterns exhibited by all possible parameter combinations. Denote S-optional with (S). When the order of S and O is ambiguous, place O closest to the verb. The first two rows are given:

S_{right}	O_{right}	S_{drop}	language pattern
0	0	0	SOV
0	0	1	(S)OV

[3 marks]

- (b) A learning algorithm attempts to identify the parameter settings which characterise a *target language*. The algorithm records the *learner state* (the current hypothesised parameter settings). It proceeds by checking the compatibility of an input sentence with the learner state. If they are not compatible, the algorithm flips **one** bit in the learner state. Draw a diagram to represent the possible transitions between learner states. [2 marks]
- (c) The target language is represented by the parameters [011] ($S_{\text{right}} = 0$, $O_{\text{right}} = 1$, $S_{\text{drop}} = 1$). Sentences in a language may contain 0, 1 or 2 objects and therefore sentences in the target language are in one of the following forms:

1. SV
2. SVO
3. SVOO
4. V
5. VO
6. VOO

When presented to the learning algorithm, the probability distribution over sentence forms is uniform. Each bit has an equal probability of being selected for flipping when the input is not compatible. Provide a transition matrix showing transition probabilities between learner states. [4 marks]

- (d) If the algorithm is initialised with state [100], give an equation for the expected number of steps to converge on the target language parameters. [3 marks]
- (e) Given equal probability of any given state being used to initialise the algorithm, write an equation for the overall expected number of steps before convergence on the target language parameters. [2 marks]
- (f) Suggest an alternative learning algorithm within the same learning paradigm and discuss the assumptions of the learning paradigm with respect to human language learning. [6 marks]

5 Formal Models of Language

This question concerns the noisy channel framework for problem-solving in natural language processing.

(a) You have the following sentences translated from English into Triposi.

English	Triposi
she drinks water	mwamni sileng
the rain soaks the teacher	sileng mworob sesesrakan
the teacher drinks here	sesesrakan mwamni mwabma
the teacher keeps drinking	sesesrakan mwatbo mwamni

(i) Translate the following English sentences into Triposi:

the rain keeps soaking the teacher

she keeps drinking water here

[2 marks]

(ii) Describe how we can calculate the likelihood of a translation using the noisy channel framework: you will need to give and explain the equation for decoding from one language into another, and explain how you can obtain the information needed to carry out the calculations. [4 marks]

(iii) What problem do the following underlined English words present, given the training data we have so far, and how can you still translate them into Triposi with a machine?

the teachers drink Perrier

[3 marks]

(b) You know that your optical character recognition model has made errors on every predicted instance of ‘e’. Explain what information you need in order to automatically correct the errors with a noisy channel approach. [3 marks]

(c) A signaller sends you Morse code messages, but you know that they send a dot when they should send a dash two times in five, and a dash instead of a dot three times in ten. You also know that they use the character M (– –) 3 times in 100, N (– ·) and I (· ·) 7/100, and A (· –) 8/100.

You receive the message “· · – · · –”. What is the likelihood that it represents MIN, MAN, NAN or AIM? [4 marks]

(d) Discuss the similarities and differences between the noisy channel and human processing of spoken English. [4 marks]

6 Further Graphics

(a) Which of the following is an implicit function for a closed curve? Briefly explain.


(i) $x^2 + (xy)^2$, $x, y > 0$ [1 mark]

(ii) $e^{f(x,y)} - 1$, where $f(x, y)$ is an implicit function for a closed curve. [1 mark]

(iii) $f(x, y) = 1$ if $g(x, y) > 1$ and $f(x, y) = g(x, y)$ otherwise, where $g(x, y)$ is an implicit function for a closed curve. [1 mark]

(iv) $f(x, y)g(x, y)$, where f and g are implicit functions for circles of radius 2, one centered at $(0, 0)$ and the other at $(1, 0)$. [2 marks]

(b) In this question, we will derive an implicit representation for a triangle.



(i) Write the implicit function for a line passing through $(0, 0)$ on the xy -plane. [1 mark]

(ii) Derive the implicit functions of the three lines in the figure on the left. [3 marks]

(iii) Derive an implicit function representing the triangle that is formed by the three lines. The function is 0 inside the triangle (shaded in the figure) and non-zero otherwise. [Hint: You may use the function $\max(0, x)$.] [3 marks]

(c) In this question, we represent rotations in the xy -plane with quaternions.

(i) Write the quaternion representing a rotation of angle θ around the z -axis. [1 mark]

(ii) Derive the quaternion for rotation by θ_1 and then by θ_2 around the z -axis. [Hint: $\cos(a)\cos(b) - \sin(a)\sin(b) = \cos(a+b)$, and $\sin(a)\cos(b) + \cos(a)\sin(b) = \sin(a+b)$.] [3 marks]

(iii) Starting from spherical blending of quaternions, prove that shortest path interpolation from the first to the second quaternion above (with $\theta_2 > \theta_1$) is given by: $\mathbf{q}(t) = \cos([(1-t)\theta_1 + t\theta_2]/2) + \hat{z} \sin([(1-t)\theta_1 + t\theta_2]/2)$, where $t \in [0, 1]$. [Hint: Recall that $\mathbf{q}^t = e^{t \log \mathbf{q}}$, $\log \mathbf{q} = \frac{\theta}{2} \mathbf{s}$, and $e^{\mathbf{q}} = \cos \|\mathbf{q}\| + \frac{\mathbf{q}}{\|\mathbf{q}\|} \sin \|\mathbf{q}\|$ for a quaternion $\mathbf{q} = \cos(\theta/2) + \mathbf{s} \sin(\theta/2)$.] [4 marks]

7 Further Graphics

- (a) Recall that the (local) rendering equation is $L_o(\mathbf{x}, \vec{\omega}_o) = L_e(\mathbf{x}, \vec{\omega}_o) + \int_{H^2} f_r(\mathbf{x}, \vec{\omega}_i, \vec{\omega}_o) L_i(\mathbf{x}, \vec{\omega}_i) \cos \theta_i d\vec{\omega}_i$. Assume we can importance sample according to a single function. For each of the following cases, explain what is the best function from the rendering equation to use for importance sampling of the incoming light directions.
- (i) Diffuse reflection, single light source with a small area, no occluders. [1 mark]
 - (ii) Diffuse reflection, incoming light is the same for all directions, occluders with complex geometry. [2 marks]
 - (iii) Diffuse reflection, incoming light is the same for all direction, no occluders. [1 mark]
- (b) We are given a scene with a single plane and light sources. The plane does not emit light, has the same BRDF at all points, and light is reflected equally for all outgoing directions.
- (i) Why can we use the local rendering equation even if we assume global illumination? [2 marks]
 - (ii) Write the simplest expression for the outgoing light at a point on the plane. [2 marks]
 - (iii) Write the expression that approximates the integral with n direction samples that are importance sampled according to the cosine term. You may ignore constants. [2 marks]
 - (iv) We know the incoming light radiance at each geometry point for each direction. We can take n outgoing radiance measurements on the plane. We use the same set of n direction samples at each geometry point to approximate the integral with importance sampling as above (ignoring constants). Compute a discrete approximation of the BRDF sampled at the same direction samples. [*Hint*: For $\mathbf{Ax} = \mathbf{b}$, assume \mathbf{A}^{-1} is known.] [5 marks]
- (c) We are given a scene where we know that the surface of the only object in the scene has positive Gaussian and mean curvature everywhere and that the surface is diffuse and does not emit light.
- (i) Derive lower bounds for the minimum and maximum curvatures of the surface at each point. [2 marks]
 - (ii) Assuming smooth BRDF and lighting functions, can we use differentiable rendering directly to estimate scene properties? Briefly explain. [3 marks]

8 Further HCI

Facebook founder Mark Zuckerberg, after changing the company name to Meta, has announced his intention to increase the market for Oculus VR headsets by replacing many current user interfaces with the “Metaverse” — a virtual reality 3D rendering. In this question, you are asked to imagine that you are responsible for designing a new interactive university, implemented in the Metaverse, called the Metaversity.

- (a) In the Further HCI course, we argue that complex interaction spaces are like programming languages, and that design of such spaces can draw on principles of programming language design. Briefly describe a design idea for the VR Metaversity, and then suggest *three* aspects of the Metaversity that might be related to aspects of a programming language, explaining for each one the nature of the analogy. [8 marks]
- (b) For *each* of these three aspects, describe how it might be represented visually, making reference to principles of correspondence in theories of visual representation. [6 marks]
- (c) For *each* of these three aspects, describe a design trade-off that should be considered in terms of the Cognitive Dimensions of Notations. [6 marks]

9 Further HCI

Many people start each day with the same set of actions on their phone, for example, checking the weather forecast, their calendar, looking at updates on Twitter, BBC News, email etc. This question asks you to imagine that you are leading the design team at a phone company, creating an interactive personal assistant that will automate this sequence of tasks. The app will be marketed as an “AI” assistant, but the company has not placed any constraint on what algorithms, representations or interaction techniques will be used.

- (a) The initial design budget will support a team of four people working for two months to carry out formative research. Explain i) what research methods the members of the team would use, ii) what empirical data they would collect, iii) what analytic activities they would carry out, and iv) how the results would be presented. [4 marks]
- (b) The company has allocated a prototyping budget that will support the development of two prototype apps using different technical approaches. Describe in general terms *two* approaches to interaction that would address this problem. Compare i) the elements of the visual representation that would be required for each, and ii) the cognitive implications for the user in terms of attention and decisions. [8 marks]
- (c) Describe how you would carry out a summative evaluation that compares the efficiency of the two prototypes in quantitative terms. [4 marks]
- (d) Which wave of theoretical development in HCI is best represented by this application concept? If you were to approach this market opportunity from the perspective of a different wave, which one would you choose and why? [4 marks]

10 Prolog

When answering this question ensure each relation has a comment giving a declarative reading of its behaviour. You should avoid unnecessary use of cut and not use extra-logical relations such as `findall`, `assertz` and `not (\+)`. Built-in library relations should not be assumed. The `notmember` relation given in the first part may be re-used if required.


- (a) Assume the built-in operator `\=` meaning *not unifiable with*, and a relation `notmember(+A,+L)` defined thus:

```
notmember(_, []).
notmember(A,[H|_]) :- A \= H, notmember(A,_).
```

Explain where *facts*, *rules*, *atoms*, *compound terms* have been used. Why does `notmember(A,[a,b,c])` fail? [2 marks]

- (b) Write a `reverse(+A,+B)` relation suitable for *last call optimisation*. What makes it suitable for LCO? [3 marks]

- (c) This small diagram represents our world map, with the arrows representing downhill lanes between places on the map.



Represent these downhill lanes with a `downhill(?A,+B)` relation. [2 marks]

- (d) Assuming `downhill(?A,+B)` is acyclic, define a relation `downhill_path(?A,+B)` which succeeds if place B can be reached from place A along downhill lanes. [2 marks]

- (e) Define a relation `linked(?A,+B)` which succeeds if a lane *directly* connects places A and B downhill or the reverse, e.g. `:- linked(cave,forest)` should succeed. [2 marks]

- (f) Define a relation `linked_path(+A,+B,+Path)` which finds a linked path between places A and B, reporting the ordered list of places visited from A to B in the `Path` argument. [6 marks]

- (g) Assume a relation `danger(?A,+D)` where D gives a numerical value for the danger at each place A, for example `:- danger(forest,X)` might succeed with `X=4`. Extend your `linked_path` relation so that it also returns the sum of the danger values along the path, i.e. `linked_path(+A,+B,+Path,+Danger)` [3 marks]

END OF PAPER