CST1
COMPUTER SCIENCE TRIPOS Part IB

Thursday 11 June 2026 14:00 to 17:00

COMPUTER SCIENCE Paper 7

*Answer **five** questions.*

*Submit the answers in five **separate** bundles, each with its own cover sheet. On each cover sheet, write the numbers of **all** attempted questions, and circle the number of the question attached.*

You may not start to read the questions printed on the subsequent pages of this question paper until instructed that you may do so by the Invigilator

STATIONERY REQUIREMENTS
Script paper
Blue cover sheets
Tags

SPECIAL REQUIREMENTS
Approved calculator permitted

1 Artificial Intelligence

Consider the following CSP. The nodes of the graph represent variables V_i where $i = 1, \ldots, 6$, each with domain $\{\downarrow, \uparrow, \oplus, \ominus\}$. Edges (V, V') are binary constraints, each requiring that V can take values in $\{\downarrow, \uparrow\}$ if and only if V' has values in $\{\oplus, \ominus\}$, and V can take values in $\{\oplus, \ominus\}$ if and only if V' has values in $\{\downarrow, \uparrow\}$. Let S be the sequence of assignments $(V_1 = \oplus, V_2 = \uparrow, V_3 = \downarrow, V_4 = \ominus, V_5 = \oplus)$.

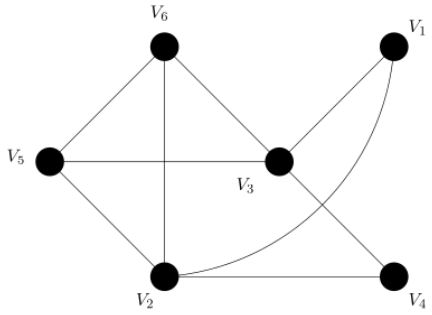

(a) (i) State the components of a *constraint satisfaction problem (CSP)* and describe the requirements for a solution to such a problem. [4 marks]

(ii) After performing the assignment S, can a valid assignment be made for V_6? Explain your answer. [1 mark]

(iii) Briefly explain how *forward checking* operates when solving a CSP. [2 marks]

(iv) Explain in detail how forward checking operates when applying the sequence S. At which assignment in S does the procedure indicate we can not continue? [4 marks]

(v) Would it have been possible in Part (a)(iv) to establish sooner that S could not lead to a solution, by performing a more detailed examination of the constraints? Explain your answer. [3 marks]

(b) Now assume that no forward checking or other form of constraint propagation is used.

(i) Assume that we are using *graph-based backjumping*, and have applied the sequence S of assignments. We are now attempting to assign V_6. Explain how this leads to the first backjump, and how we establish which variable the first backjump decides to re-assign. [2 marks]

(ii) After performing the first backjump we again exhaust the possible assignments for the variable. Explain how graph-based backjumping establishes the second backjump in this example. [4 marks]

2 Artificial Intelligence

(*a*) You wish to perform *linear regression* using a function

$$h_{\mathbf{w}}(\mathbf{x}) = w_0 + \sum_{i=1}^{n} w_i x_i$$

where $\mathbf{w}^T = \begin{bmatrix} w_0 & w_1 & \cdots & w_n \end{bmatrix}$ is a vector of weights including bias w_0, and $\mathbf{x}^T = \begin{bmatrix} x_1 & x_2 & \cdots & x_n \end{bmatrix}$ is a vector of features. You will do this by minimising the error function

$$E(\mathbf{w}) = \sum_{i=1}^{m} \left(h_{\mathbf{w}}(\mathbf{x}_i) - y_i\right)^2$$

for a training set $S = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_m, y_m)\}$. From an understanding of the underlying problem, you know that a given subset of the weights must be non-positive, and the rest must be non-negative. The bias is unconstrained.

(*i*) How can you influence the signs of the weights by modifying $E(\mathbf{w})$ only, by adding a term that is a function of $\mathbf{w}$ but not of the $\mathbf{x}_i$? In your answer it should be possible for the strength of the influence to be adjusted.
[4 marks]

(*ii*) How can you influence the signs of the weights by modifying $h_{\mathbf{w}}(\mathbf{x})$ only? In this part you do not need to adjust the strength of the influence. Discuss which changes to the usual learning algorithm are needed. It is not necessary to derive the learning algorithm in full. [4 marks]

(*iii*) Which of your suggestions from Parts (*a*)(*i*) and (*a*)(*ii*) would you expect to be more effective in practice? [2 marks]

(*b*) You have found that linear regression is not appropriate for the problem at hand. You have tried unsuccessfully to use a standard multi-layer perceptron (MLP) with a single layer of p hidden nodes. Instead you want to try to solve it using the function

$$h_W(\mathbf{x}) = w_0 + \sum_{i=1}^{p} w_i \exp\left(-\frac{1}{\sigma_i^2}||\mathbf{x} - \mathbf{c}_i||^2\right).$$

It is now no longer necessary to influence the signs of the weights. Here, W is a set containing all the weights w_i, σ_i and $\mathbf{c}_i$. The error function remains the sum of squared differences, but now depends on all the weights in W.

(*i*) How does the standard MLP differ from this system? [2 marks]

(*ii*) Assuming that you have already completed backpropagation to compute the gradient of the error function for the MLP case, explain in detail how the calculation of the gradient has to be modified for training the new system. [8 marks]

3 Economics, Law and Ethics

The new Master of (fictional) Porterhouse would like to modernise the college's room ballot. Previously, the Head Porter assigned rooms based on 'social standing' and family connections to the college, which the new Master deems opaque and inefficient. The Master proposes a new Centralised Allocation Algorithm, where students can bid on different rooms according to their preferences. As a member of the College Council, you are asked to provide input on the design. In your answers, explain and define all relevant concepts.

(*a*) It is suggested that students should be able to voluntarily swap rooms after the allocation. Does this meet the Pareto efficiency aims of the college? [3 marks]

(*b*) The student representative proposes that students are able to bid on a 'bundle' of room attributes, e.g. a north-facing room near the library with an ensuite. Argue for or against this proposal, outlining the type of auction that would be required if it were to be accepted. [5 marks]

(*c*) A 'group ballot' has been proposed to allow students to bid with their friends as a block on 'groups' (clusters of rooms). Model the decision of four friends to bid together as a Stag Hunt. Create an illustrative payoff table and identify the Nash Equilibria in this scenario. [7 marks]

(*d*) It is well-known that some rooms in Porterhouse are decrepit, whereas others are splendid. However, the students are only shown photographs of the best rooms during the application process. If students cannot distinguish good rooms from bad rooms, what impact might this have on the ballot system? [5 marks]

4 Economics, Law and Ethics

(*a*) Describe virtue ethics, Ubuntu ethics, consequentialism and deontology. Explain and define all relevant ethical concepts. [6 marks]

(*b*) Your company wants to develop a secure messaging application, but worries that the UK government could legally compel it to remove electronic protection (encryption) or to build functions that allow for the interception of messages. How could the Government do this, and under which legislation? [2 marks]

(*c*) The company needs to decide how to build a secure messaging application even in the scenario described in Part (*b*). To do so, they consider two design options:

Option A: Design the system using End-to-End Encryption (E2EE). This ensures maximum user privacy and security, but means the company cannot be compelled to decrypt data because it does not hold the keys.

Option B: Design a Key Escrow system. This allows the company to maintain high security for users while ensuring they can assist law enforcement if the application is used to cause harm.

Identify the key questions that arise under the four ethical frameworks. Given these considerations, what is your preferred design choice? Justify your answer. [12 marks]

5 Formal Models of Language

Alice and Bob have both installed an app that uses a language model exclusively trained on examples of their conversations with each other. In this app, Bob can watch words appear on his device's screen as Alice types into her device.

(*a*) Today, Alice and Bob are arranging to meet. Alice is writing to Bob. Bob watches Alice's words appear on his screen. The words say *I'll see you at the* and then she gets interrupted. The language model's three most probable continuations are *lab*, *canteen*, and *plodge*. Calculate the surprisal (in bits) of all three given that: $P(lab \mid \text{context}) = 0.5$, $P(canteen \mid \text{context}) = 0.125$, $P(plodge \mid \text{context}) = 0.03125$. Provide relevant equations. [3 marks]

(*b*) At noon the next day, Bob is arranging to meet Alice for lunch. He is in a rush and messages her *c u @ cnTn*. Alice understands the message to mean *see you at the canteen*.

(*i*) With reference to information theory, explain how Alice was able to decode the message. [5 marks]

(*ii*) Alice sends a voice note reply. She says: *No, I'll meet you at, uhm, well, the chronophage*. With reference to ideas from information theory, explain why the inclusion of these fillers might help Bob process the high-surprisal word *chronophage*? [3 marks]

(*c*) Alice and Bob discover that the model assigns high surprisal to the word *arrived* in the sentence *The lecturer who the students like arrived late*. Given that their model is an n-gram model, explain why this might be. [3 marks]

(*d*) Alice wants to reduce the bandwidth used by the app. She proposes a probabilistic encoding scheme where the number of bits used to transmit a word is proportional to its surprisal.

(*i*) Discuss the efficiency of Alice's scheme with reference to the sentence *The lecturer who the students like arrived late*. [3 marks]

(*ii*) Make suggestions of ways to improve the language model to reduce bandwidth further, commenting on feasibility. [3 marks]

6 Formal Models of Language

(*a*) Construct a context-free grammar that generates each of the following languages (and justify each answer):

(*i*) $L = \{a^n b^{n+m} a^m \mid n, m \geq 0\}$ [4 marks]

(*ii*) $L = \{w \in \{a, b\}^* \mid |w|_a = 2k,\ k \geq 0\}$ [4 marks]

(*iii*) $L = \{w \in \{a, b\}^* \mid |w| = 2k + 1,\ w_1 = w_{k+1},\ k > 0\}$ [4 marks]

where $|w|$ denotes the length of string w, $|w|_a$ denotes the number of a symbols in w, and w_i denotes the i-th symbol in w.

(*b*) Consider the context-free grammar G with a single non-terminal symbol S and with the following production rules

$$S \to a\,S$$
$$S \to a\,S\,b\,S$$
$$S \to c$$

(*i*) Is the language generated by this grammar regular? Justify your answer. [5 marks]

(*ii*) Is this grammar ambiguous? Justify your answer. [3 marks]

7 Further Graphics

(*a*) (*i*) You are given a function $f(t) : \mathbb{R} \to \mathbb{R}$. Give the x and y coordinates of the corresponding planar curve $\mathbf{p}$ without re-parameterisation. [1 mark]

(*ii*) Is there any $\mathbf{p}$ as defined above that is arc-length parameterised? If so, what are the possible functions f? [3 marks]

(*iii*) Given an arc-length parameterised planar curve $[t^2, g(t)]$, give the expressions for the tangent and absolute curvature in as simple a form as possible. [5 marks]

(*b*) A signed distance field (SDF) f is an implicit function for a surface with the property that $f(\mathbf{x})$ is the distance to the closest point to $\mathbf{x}$ on the surface and the gradient of f has unit norm.

(*i*) If f is an SDF of a sphere of radius r and d is a scalar, what surface does the SDF $f - d$ represent? [2 marks]

(*ii*) Assume the sphere was reconstructed from sample points on the surface to get the SDF f. Derive the expression of the SDF of the sphere if the sample points were rotated with the matrix $\mathbf{R}$ and translated with the vector $\mathbf{t}$ before reconstruction. [5 marks]

(*iii*) If $(f+g)/2$ is a valid SDF for general surfaces, what is the relation between the surfaces represented by the SDFs f and g? [4 marks]

8 Further Graphics

(a) We want to render a single triangle with vertex locations $\mathbf{v}_1$, $\mathbf{v}_2$, $\mathbf{v}_3$, and the normal $\mathbf{n}$. Assume the outgoing light in a fixed direction has radiance L_1 at $\mathbf{v}_1$ and L_2 at $\mathbf{v}_2$. The BRDF is the same everywhere on the triangle.

(i) What is the outgoing light radiance in the same fixed direction at some point on the triangle if we assume only a directional light source that emits light in a single direction? Explain the steps with equations. [4 marks]

(ii) What is L_2 in terms of L_1 if we assume diffuse reflection and only a point light source at the point $\mathbf{v}_1 + t\mathbf{n}$ that emits light of equal radiance in all directions? Explain the steps with equations. [5 marks]

(b) A mesh is attached to two bones via rigging with weight functions $w_1(\mathbf{x})$ and $w_2(\mathbf{x})$ and linear blend skinning.

(i) For a vertex with position $\mathbf{x} \in \mathbb{R}^3$ on the mesh, give the expression of the new position of this vertex after applying the transformations $\mathbf{T}_1$ and $\mathbf{T}_2$ to the two bones, where the $\mathbf{T}$s are 4×4 transformation matrices. [2 marks]

(ii) Assume rotations for both bones are around the same axis. Does linear blend skinning provide a valid rigid transformation in this case? [2 marks]

(iii) Assume both bones are rotated using the same matrix. Does linear blend skinning provide a valid rigid transformation in this case? [2 marks]

(iv) You now want to further translate the vertex by $\mathbf{t} \in \mathbb{R}^3$. Give the expression of the transformation matrix $\mathbf{T}$ to be added to $\mathbf{T}_1$ to achieve this. Assume that $\mathbf{T}$ represents a valid rigid transformation with no rotation. Explain the steps with equations. [5 marks]

9 Further Human–Computer Interaction

This question compares the approaches that might be taken to two different interaction design projects. The first project is the redesign of the dashboard user interface for a new generation of electric family car that includes self-driving features in addition to conventional manual controls. The second project is the redesign of the cockpit user interface for a new generation of passenger plane that includes autopilot features in addition to conventional manual controls.

Each section of this question should be answered in two parts, contrasting the two design problems. The first part of your answer should relate to the car design, and the second part should relate to the plane design.

(*a*) Which aspects of the three waves of HCI research and development are most relevant? Justify your answers. [4 marks]

(*b*) What kind of specialist researcher would you recruit to carry out formative (as opposed to summative) research? Include information about what kind of academic discipline they would be trained in, and what kind of methods they would use. [4 marks]

(*c*) What early prototyping methods would you use to collect information about potential alternative design approaches? What kinds of comparative analysis might be appropriate in order to inform design choices? [4 marks]

(*d*) What data might be collected after product launch to evaluate the success of the design and rectify any usability problems? [4 marks]

(*e*) Some issues might be missed if the HCI investigations focus only on the specific aspects that you highlighted as most relevant in Part (*a*). For each project, discuss the most serious issue missed. Is there a secondary investigation, using an approach from a different wave, that you could have done instead? Justify your answer. [4 marks]

10 Further Human–Computer Interaction

(*a*) Explain the aspects of goal-directed human behaviour involved in the phenomenon of bounded rationality, including the role of heuristics and satisficing. [2 marks]

(*b*) How do these principles apply when using an internet search engine to buy a product online? [2 marks]

(*c*) What will change, by reference to these principles, if the search interface is modified to include an AI-generated summary before the list of results? [2 marks]

(*d*) Sketch the design for a graphical browser add-on that would allow a user to specify a policy for their own future online search. The interface should allow the user to specify a condition and corresponding action. For instance, for any given new content, a user might want to default to using Bandcamp, where they can pay the artist directly, rather than using Spotify for streaming. Your add-on should be sufficiently general to support users in many different market domains. [4 marks]

(*e*) Give an example of a prompt that might be used when directing an LLM to automatically implement the same kinds of policy supported by the interface you sketched in Part (*d*). [2 marks]

(*f*) Use the cognitive dimensions of notations framework to compare four usability aspects of the two approaches in Part (*d*) and Part (*e*). [8 marks]

END OF PAPER