**COMPUTER SCIENCE TRIPOS Part IB – 2025 – Paper 7**

## 1 Artificial Intelligence (sbh11)

(*a*) In the context of algorithms for search in AI, what is a *heuristic* and how is it used? Give a definition of what it means for a heuristic to be *admissible*. [3 marks]

(*b*) You have a collection $h_1, h_2, \ldots, h_n$ of heuristics, each of which is admissible. Prove that the heuristic

$$h(s) = \sum_{i=1}^{n} a_i h_i(s) \text{ where } a_i > 0 \text{ and } \sum_{i=1}^{n} a_i = 1$$

is also admissible. [2 marks]

Your aim is to use supervised learning to attempt to infer a general admissible heuristic for a particular search problem. States are denoted as vectors, and you have managed to gather information on a collection $(s_1, s_2, \ldots, s_m)$ of $m$ states, and the corresponding values $(h'(s_1), h'(s_2), \ldots, h'(s_m))$ for the *exact* distances between states and goals. Your friend has some software called VectorDribble and suggests that you use its built-in function for linear regression to achieve this.

(*c*) How would you address this as a supervised learning problem? Is linear regression likely to be a good method? Explain your answer. [3 marks]

(*d*) You find that VectorDribble's built-in function uses the standard measure of error

$$E(\mathbf{w}) = \sum_{i=1}^{m} \left(h'(\mathbf{s}_i) - f(\mathbf{s}_i, \mathbf{w})\right)^2$$

where $\mathbf{w}$ is the learned vector of weights and $f(\mathbf{s}_i, \mathbf{w})$ is the output of the learned function. Your friend claims that this is the error you must use as it is built into the library. Explain why this function should not be used. [2 marks]

(*e*) Suggest a better function to use in place of the built-in $E(\mathbf{w})$. Explain how it improves on the built-in version. [4 marks]

(*f*) As you now can not use VectorDribble to solve the problem, derive a gradient descent-based learning algorithm that is suitable. You may continue to assume that the underlying linearity is appropriate. [6 marks]