

APOLLO

A4 Machine learning HL

Computer Science
First assessment 2027
1 questions | mark schemes

4. (a) *Award [1] max*

Input layer;
Output layer;
Convolution layer/conv layer;
Pooling layer/subsampling layer/down-sampling layer;
Fully-connected layer/FC layer;
Normalization layer;
Activation layer.

(b) *Award [1] max*

To improve model accuracy/ensure the model is not trained on poor quality data;
To ensure inaccuracies/errors/inconsistencies in the data doesn't lead to false assumptions;
To removing duplicate data/irrelevant data points;
Reduce the training time/computational resources required;
To handle missing values by using imputation (e.g. inserting mean) or removing;
Ensures the CNN doesn't make predictions on incomplete data/have gaps in its knowledge;
To provide uniformity of data/ensure all images are the same size and format;
To ensure the model isn't giving undue importance to some features over others;
To balance the dataset (undersampling or oversampling);
To ensure the dataset doesn't have an uneven distribution of classes/become biased towards the majority class;
To avoid overfitting;
To remove outliers or extreme values in the training data that can cause overfitting.

(c) *Award [2] max*

Award [1] per point from one section

Deep learning involves many matrix operations (multiplications and convolutions);
GPUs have many smaller cores designed explicitly for parallel operations;

Large datasets increase the computational demand;
Training on a CPU only would be too time-consuming/GPUs speed up the training time;

Converges by minimizing the cost function (e.g., batch gradient descent) lending itself to simultaneous operations;
GPUs provide parallel processing, leading to faster convergence;

Accept any reasonable machine learning process that requires parallelization

– 12 –

(d) *Award [3] max**Award [1] per point*

Calculate the F1 score as the harmonic mean of precision and recall;
 A low F1 score implies inadequate performance of the CNN model;
 This may require training adjustments or data rebalancing;
 It is an iterative process of adjusting and testing to improve the F1 score.

Alternative answer

Produce a confusion matrix (TP, TN, FP, FN);
 Calculate precision using $TP / (TP + FP)$;
 Calculate recall/sensitivity using $TP / (TP + FN)$;
 Calculate F1 score = $2 * (Precision * Recall) / (Precision + Recall)$.

Note to examiners:

TP is true positive (correctly identified cases of pneumonia);
TN is true negative (correctly identified cases without pneumonia);
FP is false positive (wrongly identified cases of having pneumonia);
FN is false negative (cases of pneumonia that were not identified).

(e) *Award [3] max**Award [1] for identifying an **issue****Award [1] for expanding on the issue/stating the **context****Award [1] for a further **evidence** highlighting the choice of RL over SL*

RL works well in **dynamic environments** where layouts change/hospitals are dynamic environments (beds move, machinery is moved);
 RL involves the cleaner (agent) interacting with the environment (hospital), receiving continuous feedback based on its actions (movement decisions);
 RL allows the robot cleaner to update its learned policy through rewards and punishments;
 SL is trained in advance on a static dataset so would struggle to adapt to untrained layouts/changing environments;
 SL requires accurate labelled datasets, which is impractical for every possible situation or condition in a hospital;

Alternative answer:

RL allows for **real-time autonomous learning** (without interventions);
 The robot (agent) uses RL to interact with the environment making decisions;
 The feedback loop in RL (perform action, get reward, update policy) ensures the robot continuously improves its strategy based on real-time outcomes;
 SL is trained before use and does not learn in real-time;
 Changes in the environment would require external intervention (retraining);
 A dataset with all possible layouts and actions is implausible;

Alternative answer:

RL is goal driven/has an **ultimate goal** of cleanliness;
 RL is designed to allow the cleaner to discover an optimal sequence of actions using trial-and-error;
 RL, unlike SL, can handle the situation where there are delayed rewards;
 SL is not designed to find the best sequence of actions;
 SL is designed to predict outcomes based on provided input;
 SL is suited to classification and regression tasks, not optimization tasks.

(f) **Award [5] max**

The following points might be considered:

- Definition of online harassment - any form of bullying or intimidation that takes place through digital mediums
- Consideration of the context because much depends on the comments
- The sensitivity shown by hospital staff – only online harassment when they harm, threaten, or are offensive. They may only be asking for retractions in extreme circumstances
- The method of contacting public members (online, email, phone)
- How the hospital obtained members contact details (right to privacy)
- Freedom of speech and protection from unsubstantiated comments
- Potential impact on the public's trust in the hospital
- Alternative suggestions to the current practice to alleviate ethical concerns
- Legal implications such as GDPR
- Impacts of online harassment can be severe, leading to emotional distress, anxiety, depression, and even self-harm or suicide in extreme cases
- Is it ethical to use machine learning in this situation
- Consideration of the training data used for sentiment analysis
- The accuracy and potential bias (e.g. ethnicity, socio-economic status, gender) of the model

Marks	Level descriptor
No marks	No knowledge or understanding of the relevant issues and concepts.
1–2 marks	Limited knowledge and understanding of harassment. The answer may be little more than a list of points. (General)
3–4 marks	Detailed knowledge that links harassment to the scenario; Fairly balanced analysis. (In context)
5 marks	An appropriate conclusion, linked to the analysis.

Section B