2224 – 6916M

Diploma Programme
Programme du diplôme
Programa del Diploma

Markscheme

May 2024

Computer science

Higher level

Paper 3

10 pages

Subject details: **Computer science HL paper 3 markscheme**

Mark allocation

Candidates are required to answer **all** questions. Total 30 marks.

General

A markscheme often has more specific points worthy of a mark than the total allows. This is intentional. Do not award more than the maximum marks allowed for that part of a question.

When deciding upon alternative answers by candidates to those given in the markscheme, consider the following points:

- Each statement worth one point has a separate line and the end is signified by means of a semi-colon (;).
- An alternative answer or wording is indicated in the markscheme by a "/"; either wording can be accepted.
- Words in (…) in the markscheme are not necessary to gain the mark.
- If the candidate's answer has the same meaning or can be clearly interpreted as being the same as that in the markscheme then award the mark.
- Mark positively. Give candidates credit for what they have achieved and for what they have got correct, rather than penalizing them for what they have not achieved or what they have got wrong.
- Remember that many candidates are writing in a second language; be forgiving of minor linguistic slips. In this subject effective communication is more important than grammatical accuracy.
- Occasionally, a part of a question may require a calculation whose answer is required for subsequent parts. If an error is made in the first part then it should be penalized. However, if the incorrect answer is used correctly in subsequent parts then **follow through** marks should be awarded. Indicate this with "**FT**".
- Question 4 is marked against markbands. The markbands represent a single holistic criterion applied to the piece of work. Each markband level descriptor corresponds to a number of marks. When assessing with markbands, a "best fit" approach is used, with markers making a judgment about which particular mark to award from the possible range for each level descriptor, according to how well the candidate's work fits that descriptor.

General guidance

Issue	Guidance
Answering more than the quantity of responses prescribed in the questions	• In the case of an "identify" question read all answers and mark positively up to the maximum marks. Disregard incorrect answers. • In the case of a "describe" question, which asks for a certain number of facts *eg* "describe two kinds", mark the first two correct answers. This could include two descriptions, one description and one identification, or two identifications. • In the case of an "explain" question, which asks for a specified number of explanations *eg* "explain two reasons …", mark the first two correct answers. This could include two full explanations, one explanation, one partial explanation *etc*.

1. (a) *Award **[2 max]***
*Award **[1]** for definition/purpose and **[1]** for usage/example*

Machine learning/Artificial Intelligence subcategory that enables computers to derive meaning from visual information (simulate human perception);
Automates (analyse, process, and interpret) tasks requiring human visual intelligence;

System for acquiring, processing, analysing digital images/video/real world visual inputs;
To produce numerical or symbolic information used for decision-making;

System to interpret the visual environment using technology;
Then makes decisions upon the processed data (e.g. avoid an object in the robot's path/ determine its location);

System that detects and processes the environment;
Encompassing the subdomains: vSLAM/pose estimation/object recognition;

(b) *Award **[2 max]***
*Award **[1]** for sensor input and **[1]** for processing*
Uses moving parts such as a wheel or chain to determine speed/distance;
Processor calculates distance using starting point and sensor input;

The odometry sensor sends the number of wheel revolutions (for each wheel);
Multiply this number by the circumference of its wheels;

Accelerometers/wheel encoders contribute to distance;
Gyroscopes/magnetometers contribute to direction;

Processor receives odometry sensor data in terms of rotation number;
Distance = Number of Rotations × Circumference of Wheel/Circumference = π × distance;

Sensor input for each wheel is different for a robot moving in an arc;
The processor uses the difference in wheel distances to calculate the angle (and arc radius);
The robot updates its current position by determining distance and direction;

Encoders produce electrical signals (pulses) corresponding to wheel rotation;
Counting these pulses, the processor can determine how much the wheel has turned;

2. (a) *Award **[4 max]***
*Award **[1]** for any point and **[+1]** for further technical detail / example*

Sensing Difficulties
Identification relies on predetermined key points (e.g. joints) which may be missing;
Missing joint segments lead to erroneous predictions in pose data;
If the occluded object or the robot is moving, it may blur the image leading to inaccuracy;
The overlapping object may be included as part of the body structure/two people lying together may be interpreted as one object;

Computational Difficulties
More complex methods (e.g. sensor fusion) needed to infer occluded parts increases computational demands;
The need for real-time processing in rescue operations is compromised by occlusions; impacting system performance/increasing computational complexity;

Model and Algorithm Reliability
Estimations rely heavily on contextual features from the surrounding area/ limited light compounds the difficulties in pose estimation when there is occlusion;
Errors in the detection of visible parts can worsen inaccuracies in estimating occluded parts;
Extensive training on diverse datasets is essential to prepare models for various occlusion scenarios/Models may perform well in testing but not in the real world;

Allow marks to be award in any combination.

(b) *Award **[4 max]***
*Award **[1]** for any point and **[+1]** for further technical detail / example*

Rescue robots (e.g. VSLAM, pose estimation) require fast (real-time) processing;
Edge reduces the need for long-distance data transmission to central servers;
Edge computer allows fast local processing (either on the robot or nearby);
Allows for lightweight/manoeuvrable robots (offloading processing and storage tasks);
Reduces latency/improve the robot’s response time;
Allows for scalability for deploying more robots/create a mesh network between robots;
Ensures reliable operation (even in environments with poor or disrupted; connectivity)/minimises network congestion/allows for redundancy;
Minimises the risk of sensitive information being exposed during transmission;
Uses less power/conserving robot's battery life;
Enables a local human team / command centre to access all data in real time for decision making or intervention;

3. *Award **[6] max***

*Award **[1]** for each correct point up to **[6]** max.*

Concern that humans aren’t making these decisions/Trust issues with allowing MLs to make life and death decisions/MLs make black box decisions;

Human operators should be able to intervene, if necessary;

Fairness implies that algorithms should not create or reinforce unfair biases;

Awareness of predetermined fair criteria for rescue prioritisation (such as the urgency of medical need, rescue feasibility, or the potential for saving lives);

ML's decisions should not systematically disadvantage any group or individual/should not use prohibited characteristics to make decisions/datasets should not include prohibited features;

Characteristics (such as race, gender, age, etc.) must not influence rescue priority

A diverse team to ensure programmers natural biases towards people like them are eliminated;

Use inclusive training datasets to ensure all casualties have the same opportunity for rescue;

Concern that the ML algorithm has not been tested across diverse scenarios/that insufficient simulations have been run/that ML algorithms are making decisions on its own experiences;

Maintain transparency and accountability in how decisions are made/allow for audits and corrections of the system/screening of suitably diverse datasets;

4. *Award* ***[12 max]***
Candidates may discuss some of the following points:

General comments
- The effectiveness of rescue robots can vary depending on several factors, including the environment, the quality of the sensor data, and the complexity of the search area.
- The use of performance metrics for vSLAM-enabled robots in comparison with human rescuers.
- This reduces the time it takes to locate survivors, especially in an unsafe environment like a factory fire.
- The requirement for computing power, which reduces battery life and operational time.
- The inability of deployed robots to operate for extended periods without recharging.
- The ethics of relying on robots equipped with vSLAM algorithms instead of human rescuers.
- Consideration of rescuers who compromise their safety to enter a burning factory.
- Robots mapping the route to enable human rescuers to navigate quickly and with lower risk.
- The collaboration between robots and human rescue teams.
- The reliance of communications between robots and humans, and robots with other robots.
- The consequences of communication failure (e.g. edge computing).
- Combining LIDAR/thermal imaging data/multi-spectral Sensor Fusion with visual camera data can enhance object detection and localisation in poor lighting/smoky/dusty environments.
- The effectiveness of vSLAM and pose estimation in densely cluttered or dynamically changing environments is limited.
- Technical failures could lead to delays in rescue operations, potentially endangering lives.
- Rescue personnel must be trained to operate and work alongside these robots effectively.
- Integrating robotic systems into existing disaster response protocols requires adjustments in operational tactics.

Time-constraints
- If pre-existing maps (e.g. factory blueprint) exist this will speed up vSLAM process.
- Collaborative Robotics: Using VTV protocols, robots can share map data and scanning results in real-time.
- Sensor fusion with technologies like LIDAR, thermal imaging, stereo vision to map the area and identify survivors quickly.
- Enhanced Localisation Techniques: inertial measurement units (IMUs) and advanced feature-matching algorithms.
- Optimising SLAM algorithms.
- Task-specific customization: Customizing robots for specific environments (e.g. industrial setting) can reduce the need for broad-spectrum scanning.

Environmental considerations
- Smoke or dust will make the cameras largely ineffective.
- Poor lighting may also reduce the effectiveness of vSLAM.
- Occlusions may make it difficult to interpret survivors.
- Pose estimation algorithms may assist with occlusions.
- Thermal cameras may be more effective in smoky environments, although heat from fires may interfere.

Candidates may evaluate vSLAM through the stages: **initialisation, Local mapping, Loop closure, Relocalization, Tracking**

Initialization
- This stage is critical as it sets the baseline for all subsequent mapping and navigation tasks, ensuring the robot has a reference point (position and orientation) from which to start its search.
- Creating the Initial set of features for mapping and tracking. The vSLAM system uses feature detection and matching techniques to identify distinctive points in the first frame of the camera feed.
- The accuracy of initialisation depends on the quality and distinctiveness of the initial features. integrating multiple sensory inputs can help improve the robustness of the initial map and pose estimation.
- Poor initial data can lead to inaccuracies in the map that propagate throughout the operation.

Local Mapping
- This dynamic updating allows vSLAM to maintain a detailed and accurate map, adapting to changes in the environment as they occur.
- As the robot moves and captures new frames of the environment, the vSLAM system incrementally adds new features to the map and refines the positions of existing features.
- Feature detection, description, matching, and optimisation techniques are employed during local mapping.

Loop Closure
- Loop closure is critical for the accuracy of long-term searches since it detects and corrects any accumulated drift or errors in the map.
- Detecting loop closures can be computationally intensive and prone to false positives, especially in environments with repetitive patterns.
- Once a loop closure is detected, the vSLAM system adjusts the map and refines the robot's pose to ensure consistency and alignment with the revisited areas.

Relocalization
- This capability is essential for recovery from tracking losses, allowing the robot to continue its task without manual intervention.
- Relocalisation occurs when the robot encounters a significant change in the environment (e.g. a collapsed ceiling) or the environment lacks distinctive features.
- Using a more diverse set of sensors, such as LIDAR combined with visual data, can provide more reliable cues for relocalisation.

Tracking the Robot's Motion
- Real-time tracking is pivotal for dynamic and autonomous navigation, enabling the robot to adjust its course instantly to obstacles or changes in the terrain.
- Tracking provides real-time feedback on the robot’s position and orientation, facilitating autonomous navigation and mapping in dynamic environments.
- The robot’s motion and pose relative to its surroundings are estimated by analysing the spatial displacements of matched feature points between frames.
- Motion blur, rapid movements, or similar impediments can disrupt feature tracking, leading to potential errors in pose estimation.
- Pose is the robot’s position and orientation. Position in 3D space is represented using x-axis, y-axis, and z-axis coordinates. Orientation (i.e. rotated or aligned) is represented by Euler angles (roll, pitch, and yaw) or quaternion representations.
- Enhancing frame rate, using motion prediction models, and incorporating inertial measurement units (IMUs) can mitigate these issues, providing smoother and more accurate tracking.

Robot drift
- Translational drift is a form of locomotion error.
- Increased rotational energy applied to spinning wheels give perfect symmetry.
- Works by modulating the power to the wheels that spin the robot.

Conclusion
A final measured conclusion is included in which the candidate links together the various points in evaluating the effectiveness of vSLAM.

Please see markbands.

Marks	Level descriptor
No marks	• No knowledge or understanding of the relevant issues and concepts. • No use of appropriate terminology.
Basic 1–3 marks	• Minimal knowledge and understanding of the relevant issues or concepts. • Minimal use of appropriate terminology. • The answer may be little more than a list. • No reference is made to the information in the case study or independent research.
Adequate 4–6 marks	• A descriptive response with limited knowledge and/or understanding of the relevant issues or concepts. • A limited use of appropriate terminology. • There is limited evidence of analysis. • There is evidence that limited research has been undertaken.
Competent 7–9 marks	• A response with knowledge and understanding of the related issues and/or concepts. • A response that uses terminology appropriately in places. • There is some evidence of analysis. • There is evidence that research has been undertaken.
Proficient 10–12 marks	• A response with a detailed knowledge and clear understanding of computer science. • A response that uses terminology appropriately throughout. • There is competent and balanced analysis. • Conclusions are drawn that are linked to the analysis. • There is clear evidence that extensive research has been undertaken.