2224 – 6911M

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 sensors and **[1]** for calculating position and orientation*

Dead Reckoning is a technique to determine the current position of a robot;
Using the distance travelled and the direction of travel relative to known start location;

Dead reckoning measures wheel rotations/data from the odometry sensors;
to calculate the robot's change in position and orientation over time/determine the robot's current location;

Dead reckoning utilises internal sensors, such as accelerometers, gyroscopes, and (odometry sensors);
Changes in acceleration and angular velocity are used to calculate position and orientation;

Do not allow the exact wording "keeping track of its own location and orientation", which is taken directly from the case study.

(b) *Award **[2 max]***
*Award **[1]** for reason and **[1]** for expansion*
Changing environments/objects moving over time (accept example);
Making it difficult for the system to recognise a known place;

Different places look similar on video/difficulties in perceptual aliasing;
Leading to incorrect assumptions about the robot's location;

Low-texture environments (accept example such as empty corridor);
Leads to noise or distortion in images/poor feature detection;

High amount of features/large environment/heavily occluded;
Make searching a vast database of features too time-consuming;

Loop closure involves complex calculations in real-time;
Robots have limited computational resources;

As the robot moves, cumulative errors in motion result/scale drift;
Errors compound causing long term drift/earlier recorded positions and orientations have deviated;

A change in lighting conditions (of a previously visited area);
Hinders the ability to recognise a revisited place;

Low quality sensors (due to computational costs) may give inaccurate positions;
Leading to long term drift when remapping;

2. (a) *Award **[4 max]***
*Award **[1]** for refining structures, **[1]** for multiple views, **[1]** for error reduction, and [1] for specific enhancements*
Refines 3D coordinates (data points) of scene features/camera parameters that observed these features;
Optimises the same object from multiple views/readings (using Non-linear Least Squares);
Minimises the reprojection error/prevents error cumulation/adjusts for any biases or systematic errors/compares changes between different frames;
Improves the map terrain and objects/improves the camera trajectory estimate/improves feature tracking and mapping/reduces drift/robot's position within the environment/improves the pose estimation accuracy;

(b) *Award **[4 max]***
*Award **[1]** for concern and **[1]** for any reasonable expansion.*
Autonomy: The level of autonomy granted to rescue robots;
Some believe that critical decision-making should remain under human control/robots lack the necessary human abilities, like empathy or general intelligence;

Safety Concern: Rescue robots cannot identify injuries;
They may inadvertently injure a person they are trying to save;

Empathy: Robots do not have emotional intelligence;
A robot is incapable of using calming language or gestures might soothe a survivor/a survivor might become agitated (putting themselves at risk);

When to deploy robots: Human rescuers tire or may be limited in numbers;
It is ethical to deploy robots when humans need rest or are unavailable;

Privacy and Surveillance: Large amounts of visual and audio data are stored;
Strong safeguards against misuse/secure unauthorised access/post rescue data concerns;

Algorithmic/Dataset Bias: Algorithms/training data can potentially be biased;
If trained using non-representative data sets/if programmers are innately biased;

Fairness and discrimination: Are certain survivors prioritised over others;
Should it be first come first served/should children be prioritised;

Responsibility and Accountability: Liability if a robot fails causes harm;
Difficult to attribute responsibility to the manufacturer, programmer, operator or organisation;
Establishing clear lines of accountability is essential to address potential harm;

Job Displacement: Robots in rescue operations may displace human rescuers;
The social and economic impacts may need to be addressed;

Informed Consent Making: Unconscious survivors can’t give informed consent;
Respecting people’s rights to make informed decisions is an ethical consideration;

Resource Allocation: Decisions regarding where and how to deploy rescue robots;
Prioritising one region over another can lead to disparities in rescue services;

Digital Divide: Some countries or organisations cannot afford rescue robots;
Equality and fairness/poorer communities are disadvantaged in emergency situations;

Expectations; There should be an assessment about the capability and limitations of robot;
Their capabilities might be overestimated/they shouldn’t be used if they’re not safe;

Safety; Rescue conditions may be dangerous for humans;
May lead to mental trauma for rescuers/may be prioritising rescuers over potential survivors;

Mark as [2] + [2]

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

Environmental Challenges
Survivors may be difficult to distinguish due to poor lighting, smoke, or dust (affects the quality of sensory data);
Extreme dirt, mud, or blood can interfere with sensors identifying features or key points;
Survivors may be partially obscured behind debris/overlap with other survivors or objects;
Dataset may not have included occluded images;
Complex background scenes make it difficult to pick out survivors;

Technical Challenges
Insufficient joint segments/body parts caused by environment/occlusion leads to predictive errors;
3D pose estimation may be too computationally intensive for the robots to complete in real-time;
Implementing Edge computing may face its own reliability challenges;
Battery draining due to computation intensive operations in 3D pose estimation;
Connectivity issues/Loss of communications if remote storage/processing;
Challenge in deciding optimal image/video quality;
Resolution trade-off between accuracy (high resolution) and real-time response (low resolution);
Challenges in using multiple sensors due to limitations in visual-only data;
Training data may be different to the environment (forest vs office) in which the robot is operating;
Variations in sensor calibration across different devices can lead to inconsistent data;
Distance from the camera introduces scale adjustments, complicating pose estimation;

Human Factors
Dataset used must have variations in size, shape, flexibility (even amputees);
Survivor's position may change/they may be moving;

The challenges must relate directly to pose estimation.

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

Benefits:
- Reduces the risk for human rescuers in unstable or hazardous environments.
- Small robots can travel where humans can’t fit.
- Robots can be deployed quickly and operate continuously significantly increasing the chances of successful rescues.
- Pose estimation can identify and determine the positions of survivors in real time.
- Accurate mapping with new spaces where no previous map is available.
- Robots can gather comprehensive data from the disaster site, including visual recordings, temperature measurements, and structural assessments.

Costs in terms of:
- Development, deployment, and maintenance of robots is high.
- Expense limits the number of organisations that can deploy these robots.
- 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.

Description of pose estimation
- Pose estimation is a CV technique that can help track a person or an object in real-world spaces. It detects the position of an object or a human in an image or a video.
- Describe in detail top down and bottom-up approaches.
- Discuss specific algorithmic requirement for a rescue robot such as pose estimation algorithms reject outliers when matching with existing databases, however in rescue operations the algorithm needs to be modified to keep the outliers as survivors may be in many twisted body shapes.

Benefits of pose estimation
- Pose estimation models have an advantage over object detection models, which can locate objects in an image but provide only coarse-grained localization with a bounding box framing the object. Pose estimation models, in comparison, predict the precise location of the key points associated with a given object.

Challenges of pose estimation
- Differences in lighting conditions, weather, viewing angle, background context, etc.
- Partial occlusion (objects or other humans obstructing the subject of the analysis).
- The complexity of the human skeletal structure can make it difficult to identify exact joint coordinates, especially for small points which are barely visible in the image.
- High dimensionality of the pose.
- Loss of three-dimensional information as a result of observing the pose from a two-dimensional image. Obtaining and annotating 3D pose images is complex and expensive.

Evaluation of vSLAM effectiveness

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.
- Creates 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 initialization heavily 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 propagates 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 optimization 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.
- Relocalization 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 relocalization.

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 analyzing 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 gives 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 to discuss the benefits and costs of vSLAM-equipped rescue robots.

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 the 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.