Humanoid Robotics Mock Exam

Course: Humanoid Robotics

10th July 2025
Discussion: 17th July 2025

Instructions

This is the mock exam for the Humanoid Robotics course. Please treat this exam as a closed-book test and do not use any additional materials while taking it. Although calculators will not be allowed during the final exam, you may use them for this one. There are 12 questions, each based on one lecture, and the exam is designed to be completed in approximately 60 minutes. The final exam will last 90 minutes and be worth 90 points, meaning you can plan one minute of work time per point. Please provide concise, mathematically precise answers.

1 Perception Basics

Question: Consider the humanoid robot H1 operating in a household environment. To compute the pose of H1's head frame relative to its torso frame, the following sequence is applied: first rotate by 90° about the z-axis (positive, counter-clockwise when viewed from above), and then translate by $\mathbf{t} = (2, -1, 3)^T$.

(a) Write the 4×4 homogeneous transformation matrix T that describes this sequence of operations. (1 point)

(b) Apply your T to the point

$$P = \begin{bmatrix} 1 \\ 0 \\ 2 \\ 1 \end{bmatrix},$$

expressed in the torso frame, and compute the resulting homogeneous coordinates of P in the head frame. (1 point)

Question: Consider robot H2, which is equipped with stereo camera C2. Its intrinsic matrix is

$$K = \begin{pmatrix} 10 & 0 & 5 \\ 0 & 10 & 5 \\ 0 & 0 & 1 \end{pmatrix}$$

(a) Determine the focal lengths (f_x, f_y) and the principal point (c_x, c_y) of camera C2. (1 point)

(b) A 3D landmark in the camera frame is given by

$$P_{\text{cam}} = \begin{bmatrix} 2 \\ 3 \\ 5 \end{bmatrix}.$$

Compute its projected pixel coordinates (u, v) on camera C2's image plane. (1 point)

2 3D World Representations

Question: A quadruped robot uses a 2.5D height map representation to model hilly terrain. Three laser scan measurements fall into the same grid cell with height values $Z_1 = 1$, $Z_2 = 3$, and $Z_3 = 2$.

(a) Compute the height map value assigned to this cell as the average height. (1 point)

(b) State one advantage and one disadvantage of using 2.5D height maps for terrain representation. (2 point)

Question: A humanoid robot maps an urban scene using a voxel grid with resolution 0.1 m.

(a) Compute the TSDF value update and weight for a voxel at $(0.5, 0.5)$ given a measured depth of 0.6 m. (2 points)

(b) Discuss how changing the grid resolution from 0.1 m to 0.01 m affects mapping accuracy and computational cost. (2 points)

3 Neural Fields

Question: Consider a ray $r(t) = o + t\,d$, sampled at N points $p_i = o + t_i\,d$. The network predicts at each p_i a volume density σ_i and a color c_i.

(a) Write the expression for absorption weight α_i. (1 point)

(b) Write the expression for accumulated transmittance T_i. (1 point)

(c) Combine these to give the formula for the rendered pixel color $\widehat{C}(r)$. (1 point)

(d) Briefly explain, in one sentence, what the product T_i physically represents. (2 point)

4 Active Perception & Next-Best-View

Question: Our humanoid maps an indoor corridor using a binary occupancy map composed of five grid cells: two free, one occupied, and two unknown (with $p_{\text{unk}} = 0.5$, $p_{\text{free}} = 0.0$, $p_{\text{occ}} = 1.0$).

(a) Compute the Shannon entropy H of this map, using the binary entropy function

$$H(p) = -p \log p - (1 - p) \log(1 - p).$$

(1 point)

(b) After one additional exploration step, one of the previously free cells is discovered to be occupied. Explain how the entropy changes and what conclusion can be drawn. (2 points)

Question: A robot at position R considers three candidate view poses V_1, V_2, and V_3. Their expected information gains (in unknown cells) are $I_1 = 2$, $I_2 = 3$, $I_3 = 1$, and their motion costs (Euclidean distances from R) are $C_1 = 2$, $C_2 = 4$, $C_3 = 1$. The utility function is

$$U_i = I_i - \alpha\, C_i, \quad \alpha = 0.3.$$

(a) Compute the utilities U_1, U_2, and U_3. (1 point)

(b) Which view V_i should the robot select as the next best view? Justify your choice based on the computed utilities. (2 points)

5 Manipulation I: Kinematics

Question: A robotic arm has two degrees of freedom in the plane, with link lengths $L_1 = 1\,\text{m}$ and $L_2 = 1\,\text{m}$. The joint angles are set to $\theta_1 = 45°$ and $\theta_2 = 45°$.

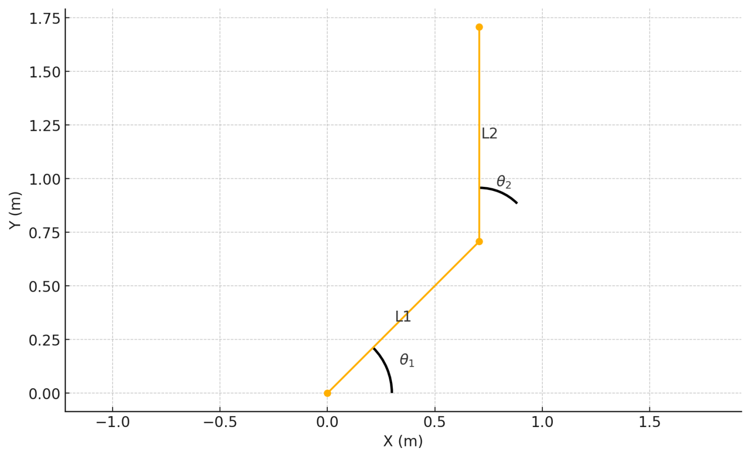

Using the Denavit–Hartenberg convention:

(a) Write the DH parameter table for this manipulator and the individual homogeneous transformation matrices T_{01} and T_{12}. (2 point)

(b) Compute the end-effector position $P_{\text{EE}} = \begin{pmatrix} x \\ y \end{pmatrix}$ in the base frame. (1 point)

6 Manipulation II: Motion Planning

Question: The arm's first two joints perform a coordinated move from

$$\mathbf{q}_{\text{start}} = \begin{pmatrix} 0° \\ 0° \end{pmatrix} \quad \text{to} \quad \mathbf{q}_{\text{goal}} = \begin{pmatrix} 90° \\ 45° \end{pmatrix}$$

using a constant-velocity profile over a total duration $\Delta t = 3\,\text{s}$.

(a) Compute the constant joint velocities $\dot{q}_1$ and $\dot{q}_2$. (1 point)

(b) Write the expression for the joint position $q_i(t)$ for $i = 1, 2$ as a function of time $t \in [0, 3]$. (1 point)

Question: Rapidly-exploring Random Tree (RRT)

(a) Provide a brief high-level pseudocode description of the RRT algorithm. (2 point)

(b) Identify and explain the roles of two key parameters in RRT (e.g., step size, maximum iterations). (2 point)

(c) Discuss one advantage and one limitation of using RRT for robotic motion planning. (2 points)

7 Manipulation III: Grasp & Force Analysis

Question: The arm's parallel-jaw gripper holds a cylindrical rod of radius r using two antipodal hard-finger contacts. Each finger applies a normal force F and the friction coefficient at the contacts is μ.

(a) Derive the expression for the maximum resisting torque $\tau_{\max}$ about the rod's central axis in terms of F, μ, and r. (1 point)

(b) Compute $\tau_{\max}$ for $F = 80\,\text{N}$, $\mu = 0.3$, and $r = 0.02\,\text{m}$. (2 point)

Question: The arm performs a planar push on a square block resting on a tabletop. The friction coefficient between the pusher and block is $\mu = 0.5$, yielding a friction-cone half-angle $\theta = \arctan(\mu)$. The pusher's velocity forms an angle $\phi = 45°$ relative to the surface normal.

(a) Compute the friction-cone half-angle θ in degrees. (1 point)

(b) Determine whether the contact will stick or slide given θ and ϕ, and briefly justify your answer. (1 point)

8 Interactive Scene Exploration

Question: Our general-purpose service robot (GPSR) perceives a partially occluded soup can on a pantry shelf. It can perform the following actions, each with its information gain (IG) and action length (L):

Action	IG	L
Reposition Whole Body and View (IGVb)	5	1
Push occluding object + Reposition and View (IGMb)	8	2
Push + Head Reposition and View (IGMh)	3	1
Head Reposition and View (IGVh)	2	1

(a) Compute the information gain per unit length, IG/L, for each action. (1 point)

(b) Which action should our GPSR choose to maximize information-gain efficiency? (1 point)

Question: Our GPSR utilizes two foundation models: RT-2 (a vision-language-action model) for high-level task planning, and SAM (a visual foundation model) for instance segmentation of cluttered scenes.

(a) Describe the primary function of RT-2 in our GPSR's manipulation pipeline. (2 point)

(b) Explain how SAM's output assists in interactive scene exploration and mechanical search. (2 point)

9 Locomotion: Kinematics & Dynamics

Question:

(a) Explain why it is necessary to normalize a quaternion to represent a valid 3D rotation. What is the physical meaning of the conjugate of a unit quaternion? (1 point)

(b) Suppose we want to rotate a 3D vector $\mathbf{v} \in \mathbb{R}^3$ using a unit quaternion $\mathbf{q}$. Provide pseudocode showing how to apply the rotation. (1 point)

Question: We wish to implement a linear Kalman Filter for estimating the floating-base position and linear velocity of a humanoid robot using foot-kinematics measurements. The state vector $\mathbf{x} \in \mathbb{R}^{12}$ is defined as

$$\mathbf{x} = \begin{bmatrix} \mathbf{p}_B \\ \mathbf{v}_B \\ \mathbf{p}_L \\ \mathbf{p}_R \end{bmatrix},$$

where $\mathbf{p}_B$ is the base position, $\mathbf{v}_B$ is the base linear velocity, and $\mathbf{p}_L, \mathbf{p}_R$ are the left and right foot positions in the world frame. The measurement is the relative position between the base and the contacting foot, computed via forward kinematics. At a given time, only the right foot contacts the ground, so

$$\mathbf{y} = \mathbf{p}_B - \mathbf{p}_R = H\,\mathbf{x}.$$

(a) Write the measurement matrix H explicitly and specify its dimensions. (1 point)

10 Locomotion II: Control and Motion Planning

Question: Given the nonlinear dynamics of the cart–pole model, where $\mathbf{q} = \begin{pmatrix} x \\ \theta \end{pmatrix}$,

$$M(\mathbf{q})\,\ddot{\mathbf{q}} + \mathbf{c}(\mathbf{q},\dot{\mathbf{q}}) + \mathbf{g}(\mathbf{q}) = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \mathbf{u},$$

$$M(\mathbf{q}) = \begin{pmatrix} m_c + m_p & m_p\,l\cos\theta \\ m_p\,l\cos\theta & m_p\,l^2 \end{pmatrix}, \quad \mathbf{c}(\mathbf{q},\dot{\mathbf{q}}) = \begin{pmatrix} -m_p\,l\,\dot{\theta}^2\sin\theta \\ 0 \end{pmatrix}, \quad \mathbf{g}(\mathbf{q}) = \begin{pmatrix} 0 \\ -m_p\,g\,l\sin\theta \end{pmatrix}.$$

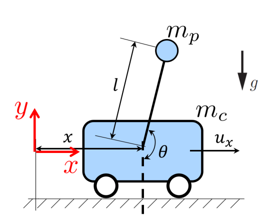

(a) Linearize the system about the upright pendulum pose ($x = 0,\ \theta = \pi$) and provide the continuous-time matrices A and B. (2 points)

(b) Describe how the choice of LQR weight matrices Q and R affects control effort and closed-loop stability. (1 point)

Question: We aim to design a Model Predictive Controller (MPC) for humanoid walking using the ZMP (Zero Moment Point) equation. For simplicity, we only consider motion in the x-direction, with fixed sampling time T. We define the system state and input as:

$$\hat{\mathbf{x}}_k = \begin{bmatrix} c_k \\ \dot{c}_k \\ \ddot{c}_k \end{bmatrix} \quad \text{and} \quad u_k = \dddot{c}_k,$$

and assume the linear inverted pendulum model (LIPM) with ZMP equation

$$z_k = c_k - \tfrac{h}{g}\,\ddot{c}_k.$$

where c_k is the Center of Mass (CoM) position in x-direction, $g = 9.81(m/s^2)$ is the gravity constant, and h is the CoM height.

(a) Explain why MPC is preferred over a PD controller for ZMP trajectory tracking. (1 point)

(b) Derive the discrete-time state-space model $\hat{\mathbf{x}}_{k+1} = A_d\,\hat{\mathbf{x}}_k + B_d\,u_k$. (1 point)

(c) Derive the ZMP output equation $z_k = D\,\hat{\mathbf{x}}_k$. (1 point)

11 Locomotion III: Footstep Planning and Realizing Motion Plans

Question: Describe how footstep planning can be formulated as a Mixed-Integer Quadratic Program (MIQP). What are the typical objectives and constraints? (2 points)

Question: Consider a 3-DoF planar robotic arm controlled using operational-space control. The control objectives are:

- Primary task (high priority): move the end-effector to a desired 2D position (x^*, y^*).
- Secondary task (low priority): maintain a desired joint posture $\mathbf{q}^*$.

Assume the joint accelerations $\ddot{\mathbf{q}}$ are the control inputs.

(a) Explain how task prioritization is achieved in this context. (1 point)

(b) Formulate the control problem as a hierarchical QP that ensures the secondary task does not interfere with the primary task. (1 point)

12 Locomotion IV: Robot Learning

Question: Compare model-free and model-based reinforcement learning.

(a) Explain the main differences between model-free and model-based RL. (1 point)

(b) Describe how supervised learning can be incorporated into model-based RL. (1 point)

Question: You are given a dataset of expert demonstrations (states and actions) for a robotic manipulation task.

(a) Which learning approach would be most appropriate to imitate this behavior? (1 point)

(b) Explain why this approach is more suitable than others. (1 point)