

Review Questions on Introduction to AI

Part I: Definitions

- Q1.** Define *Artificial Intelligence* (AI). How do engineering and cognitive-science perspectives differ?
- Q2.** Distinguish between *narrow (weak) AI* and *general (strong) AI*; give one example of each.
- Q3.** Explain the relationship between AI, machine learning (ML), and deep learning (DL).
- Q4.** State the *Turing Test*. What does it evaluate, and what are two common criticisms?
- Q5.** Contrast *symbolic (good old-fashioned) AI* and *connectionist* approaches.
- Q6.** What are *knowledge representation* and *reasoning*? Give an example formalism for each.

Part II: History

- Q7.** What was the significance of the 1956 Dartmouth Workshop to the field of AI?
- Q8.** Briefly describe Alan Turing's contributions that influenced AI.
- Q9.** Define an *AI winter*. Identify two causes of the major AI winters and their approximate time periods.
- Q10.** What were expert systems? Name one influential expert system and its domain.
- Q11.** Describe the role of statistical learning (1990s–2000s) in revitalizing AI research and applications.
- Q12.** What factors enabled the deep learning surge in the 2010s?

Part III: Major Areas

- Q13.** List five major subfields of AI and give a real-world application for each.

- Q14.** Define *search* in AI. Contrast uninformed (e.g., BFS) and informed (e.g., A*) search.
- Q15.** What is *constraint satisfaction*? Provide a canonical example problem and typical techniques.
- Q16.** What is *planning* in AI? Contrast classical planning with probabilistic planning.
- Q17.** Define *reinforcement learning* (RL). Identify and explain the roles of state, action, policy, and reward.
- Q18.** What is *natural language processing* (NLP)? Name two core tasks and common model families.
- Q19.** What is *computer vision*? Give two tasks and typical model architectures.
- Q20.** Explain *multi-agent systems*. How do cooperation and competition change solution concepts?
- Q21.** What is *explainable AI* (XAI)? Give two techniques for interpretability or explanation.

Part IV: Major Milestones

- Q22.** Why was IBM Deep Blue's 1997 victory over Garry Kasparov historically significant for AI?
- Q23.** What did IBM Watson's 2011 *Jeopardy!* win demonstrate about NLP and knowledge integration?
- Q24.** Explain the importance of AlphaGo's 2016 victory over Lee Sedol for RL and search.
- Q25.** What role did ImageNet and the 2012 breakthrough play in advancing deep learning?
- Q26.** How have large language models (LLMs) changed AI capabilities and research focus? Mention two impacts and one limitation.
- Q27.** Describe one milestone in robotics (e.g., manipulation or locomotion) and its broader implications.

Part V: Social Impacts, Risks, and Governance

- Q28.** Identify three positive social impacts of AI in healthcare, science, or accessibility.
- Q29.** Discuss job displacement and transformation due to AI. What policies can mitigate negative effects?
- Q30.** Define *algorithmic bias*. Provide one example and a mitigation strategy.

- Q31.** What are key privacy and surveillance concerns with AI systems? Include at least one technical safeguard.
- Q32.** Explain the *AI alignment* and *safety* problems. Why are they challenging?
- Q33.** What is model evaluation (including robustness and reliability) and why is it crucial for deployment?
- Q34.** Describe the goals of AI governance and regulation (e.g., risk-based approaches). Give two concrete mechanisms.
- Q35.** Outline environmental impacts of AI (training/inference). Suggest two ways to measure or reduce them.